

Kleene Algebra Modulo Theories

RYAN BECKETT, Princeton University

ERIC CAMPBELL, Pomona College

MICHAEL GREENBERG, Pomona College

Kleene algebras with tests (KATs) offer sound, complete, and decidable equational reasoning about regularly structured programs. Interest in KATs has increased greatly since NetKAT demonstrated how well extensions of KATs with domain-specific primitives and extra axioms apply to computer networks. Unfortunately, extending a KAT to a new domain by adding custom primitives, proving its equational theory sound and complete, and coming up with efficient automata-theoretic implementations is still an expert’s task.

We present a general framework for deriving KATs we call *Kleene algebra modulo theories*: given primitives and a notion of state, we can automatically derive a corresponding KAT’s semantics, prove its equational theory sound and complete with respect to a tracing semantics, use term normalization from the completeness proof to create a decision procedure for equivalence checking, and formalize an automata-based equivalence checking procedure as well. Our framework is based on *pushback*, a generalization of weakest preconditions that specifies how predicates and actions interact. We offer several case studies, showing plain theories (natural numbers, bitvectors, NetKAT) along with compositional theories (products, temporal logic, and sets). We are able to derive several results from the literature. Finally, we provide an OCaml implementation of both decision procedures that closely matches the theory: with only a few declarations, users can automatically compose KATs with complete decision procedures. We offer a fast path to a “minimum viable model” for those wishing to explore KATs formally or in code.

1 INTRODUCTION

Kleene algebras with tests (KATs) provide a powerful framework for reasoning about regularly structured programs. Modeling simple programs with while loops, KATs can handle a variety of analysis tasks [2, 7, 12–14, 36] and typically enjoy sound, complete, and decidable equational theories. Interest in KATs has increased recently as they have been applied to the domain of computer networks: NetKAT, a language for programming and verifying Software Defined Networks (SDNs), was the first remarkably successful extension of KAT [1], followed by many other variations and extensions [4, 8, 23, 37, 38, 48].

Considering KAT’s success in networks, we believe other domains would benefit from programming languages where program equivalence is decidable. However, extending a KAT for a particular domain remains a challenging task even for experts familiar with KATs and their metatheory. To build a custom KAT, experts must craft custom domain primitives, derive a collection of new domain-specific axioms, prove the soundness and completeness of the resulting algebra, and implement a decision procedure. For example, NetKAT’s theory and implementation was developed over several papers [1, 24, 51], after a long series of papers that resembled, but did not use, the KAT framework [22, 30, 39, 44]. Yet another challenge is that much of the work on KATs applies only to abstract, purely propositional KATs, where the actions and predicates are not governed by a set of domain-specific equations but are left abstract [16, 34, 40, 43]. Propositional KATs have limited applicability for domain-specific reasoning because domain-specific knowledge must be encoded manually as additional equational assumptions. In the presence of such equational assumptions, program equivalence becomes undecidable in general [12]. As a result, decision procedures have limited support for reasoning over domain-specific primitives and axioms [12, 32].

We believe domain-specific KATs will find more general application when it becomes possible to cheaply build and experiment with them. Our goal in this paper is to democratize KATs, offering

Authors’ addresses: Ryan Beckett, Princeton University, rbeckett@cs.princeton.edu; Eric Campbell, Pomona College, ehc02013@mymail.pomona.edu; Michael Greenberg, Pomona College, michael@cs.pomona.edu.

a general framework for automatically deriving sound, complete, and decidable KATs for client theories. The proof obligations of our approach are relatively mild and our approach is *compositional*: a client can compose smaller theories to form larger, more interesting KATs than might be tractable by hand. In addition to the equivalence decision procedure that comes from our completeness proof’s normalization routine, our theoretical framework has an automata theory that we prove correct. Our OCaml implementation allows users to compose a KAT with both decision procedures from small theory specifications. The automata are not only for verification, of course, they are useful for a variety of tasks such as compiling KATs to different implementations [8, 51]. We offer a fast path to a “minimum viable model” for those wishing to explore KATs formally or in code.

1.1 What is a KAT?

From a bird’s-eye view, a Kleene algebra with tests is a first-order language with loops (the Kleene algebra) and interesting decision making (the tests). Formally, a KAT consists of two parts: a Kleene algebra $\langle 0, 1, +, \cdot, * \rangle$ of “actions” with an embedded Boolean algebra $\langle 0, 1, +, \cdot, \neg \rangle$ of “predicates”. KATs capture While programs: the 1 is interpreted as skip, \cdot as sequence, $+$ as branching, and $*$ for iteration. Simply adding opaque actions and predicates gives us a While-like language, where our domain is simply traces of the actions taken. For example, if α and β are predicates and π and ρ are actions, then the KAT term $\alpha \cdot \pi + \neg\alpha \cdot (\beta \cdot \rho)^* \cdot \neg\beta \cdot \pi$ defines a program denoting two kinds of traces: either α holds and we simply run π , or α doesn’t hold, and we run ρ until β no longer holds and then run π . i.e., the set of traces of the form $\{\pi, \rho^*\pi\}$. Translating the KAT term into a While program, we write: `if α then π else { while β do { ρ }; π }`. Moving from a While program to a KAT, consider the following program—a simple loop over two natural-valued variables i and j :

```

assume i < 50
while (i < 100) { i := i + 1; j := j + 2 }
assert j > 100

```

To model such a program in KAT, one replaces each concrete test or action with an abstract representation. Let the atomic test α represent the test $i < 50$, β represent $i < 100$, and γ represent $j > 100$; the atomic actions p and q represent the assignments $i := i + 1$ and $j := j + 2$, respectively. We can now write the program as the KAT expression $\alpha \cdot (\beta \cdot p \cdot q)^* \cdot \neg\beta \cdot \gamma$. The complete equational theory of KAT makes it possible to reason about program transformations and decide equivalence between KAT terms. For example, KAT’s theory can prove that the assertion $j > 100$ must hold after running the while loop by proving that the set of traces where this does not hold is empty:

$$\alpha \cdot (\beta \cdot p \cdot q)^* \cdot \neg\beta \cdot \neg\gamma \equiv 0$$

or that the original loop is equivalent to its unfolding:

$$\alpha \cdot (\beta \cdot p \cdot q)^* \cdot \neg\beta \cdot \gamma \equiv \alpha \cdot (1 + \beta \cdot p \cdot q) \cdot (\beta \cdot p \cdot q \cdot \beta \cdot p \cdot q)^* \cdot \neg\beta \cdot \gamma$$

Unfortunately, KATs are naïvely propositional: the algebra understands nothing of the underlying domain or the semantics of the abstract predicates and actions. For example, the fact that $(j := j + 2 \cdot j > 200) \equiv (j > 198 \cdot j := j + 2)$ does not follow from the KAT axioms and must be added manually to any proof as an equational assumption. Yet the ability to reason about the equivalence of programs in the presence of particular domains is important for many real programs and domain-specific languages. To allow for reasoning with respect to a particular domain (e.g., the domain of natural numbers with addition and comparison), one typically must extend KAT with additional axioms that capture the domain-specific behavior [1, 4, 8, 29, 35].

Unfortunately, it remains an expert’s task to extend the KAT with new domain-specific axioms, provide new proofs of soundness and completeness, and develop the corresponding implementation.

As an example of such a domain-specific KAT, NetKAT showed how packet forwarding in computer networks can be modeled as simple While programs. Devices in a network must drop or permit packets (tests), update packets by modifying their fields (actions), and iteratively pass packets to and from other devices (loops). NetKAT extends KAT with two actions and one predicate: an action to write to packet fields, $f \leftarrow v$, where we write value v to field f of the current packet; an action `dup`, which records a packet in a history log; and a field matching predicate, $f = v$, which determines whether the field f of the current packet is set to the value v . Each NetKAT program is denoted as a function from a packet history to a set of packet histories. For example, the program:

$$\text{dstIP} \leftarrow 192.168.0.1 \cdot \text{dstPort} \leftarrow 4747 \cdot \text{dup}$$

takes a packet history as input, updates the current packet to have a new destination IP address and port, and then records the current packet state. The original NetKAT paper defines a denotational semantics not just for its primitive parts, but for the various KAT operators; they explicitly restate the KAT equational theory along with custom axioms for the new primitive forms, prove the theory’s soundness, and then devise a novel normalization routine to reduce NetKAT to an existing KAT with a known completeness result. Later papers [24, 51] then developed the NetKAT automata theory used to compile NetKAT programs into forwarding tables and to verify networks. NetKAT’s power comes at a cost: one must prove metatheorems and develop an implementation—a high barrier to entry for those hoping to apply KAT in their domain.

We aim to make it easier to define new KATs. Our theoretical framework and its corresponding implementation allow for quick and easy composition of sound and complete KATs with normalization-based and automata-theoretic decision procedures when given arbitrary domain-specific theories. Our framework, which we call Kleene algebras modulo theories (KMTs), allows us to derive metatheory and implementation for KATs based on a given theory. KMTs obviate the need to deeply understand KAT metatheory and implementation for a large class of extensions; a variety of higher-order theories allow language designers to compose new KATs from existing ones, allowing them to rapidly prototype their KAT theories.

1.2 Using our framework: obligations for client theories

Our framework takes a *client theory* and produces a KAT, but what must one provide in order to know that the generated KAT is deductively complete, or to derive an implementation? We require, at a minimum, a description of the theory’s predicates and actions along with how these apply to some notion of state. We call these parts the *client theory*; the client theory’s predicates and actions are *primitive*, as opposed to those built with the KAT’s composition operators. We call the resulting KAT a *Kleene algebra modulo theory* (KMT). Deriving a trace-based semantics for the KMT and proving it sound isn’t particularly hard—it amounts to “turning the crank”. Proving the KMT is complete and decidable, however, can be much harder. We take care of much of the difficulty, lifting simple operations in the client theory generically to KAT.

Our framework hinges on an operation relating predicates and operations called *pushback*, first used to prove relative completeness for Temporal NetKAT [8]. Pushback is a generalization of weakest preconditions. Given a primitive action π and a primitive predicate α , the client theory must be able to compute weakest preconditions, telling us how to go from $\pi \cdot \alpha$ to some set of terms: $\sum_{i=0}^n \alpha_i \cdot \pi = \alpha_0 \cdot \pi + \alpha_1 \cdot \pi + \dots$. That is, the client theory must be able to take any of its primitive tests and “push it back” through any of its primitive actions. Given the client’s notion of weakest preconditions, we can alter programs to take an *arbitrary* term and normalize it into a form where

all of the predicates appear only at the front of the term, a convenient representation both for our completeness proof (Sec. 2.4) and our automata-theoretic implementation (Secs. 4 and 5).

The client theory's pushback should have two properties: it should be sound, (i.e., the resulting expression is equivalent to the original one); and none of the resulting predicates should be any bigger than the original predicates, by some measure (see Sec. 2). If the pushback has these two properties, we can use it to define a normal form for the KMT generated from the client theory—and we can use that normal form to prove that the resulting KMT is complete and decidable.

As an example, in NetKAT, for different fields f and f' , we can use the network axioms to derive the equivalence: $(f \leftarrow v \cdot f' = v') \equiv (f' = v' \cdot f \leftarrow v)$, which satisfies the pushback requirements. For Temporal NetKAT, which adds rich temporal predicates such as $\diamond \circ (\text{dstPort} = 4747)$ (the destination port was 4747 at some point before the previous state), we can use the domain axioms to prove the equivalence $(f \leftarrow v \cdot \diamond \circ a) \equiv (\diamond \circ a + a) \cdot f \leftarrow v$, which also satisfies the pushback requirements of equivalence and non-increasing measure.

Formally, the client must provide the following for our normalization routine (part of completeness): primitive tests and actions (α and π), semantics for those primitives (states σ and functions pred and act), a function identifying each primitive's subterms (sub), a weakest precondition relation (WP) justified by sound domain axioms (\equiv), and restrictions on WP term size growth.

The client's domain axioms extend the standard KAT equations to explain how the new primitives behave. In addition to these definitions, our client theory incurs a few proof obligations: \equiv must be sound with respect to the semantics; the pushback relation should never push back a term that's larger than the input; the pushback relation should be sound with respect to \equiv ; we need a satisfiability checking procedure for a Boolean algebra extended with the primitive predicates. Given these things, we can construct a sound and complete KAT with an automata-theoretic implementation.

1.3 Example: incrementing naturals

We can model programs like the While program over i and j from earlier by introducing a new client theory for natural numbers (Fig. 1). First, we extend the KAT syntax with actions $x := n$ and inc_x (increment x) and a new test $x > n$ for variables x and natural number constants n . First, we define the client semantics. We fix a set of variables, \mathcal{V} , which range over natural numbers, and the program state σ maps from variables to natural numbers. Primitive actions and predicates are interpreted over the state σ by the act and pred functions (where t is a trace of states).

Proof obligations. The WP relation provides a way to compute the weakest precondition for any primitive action and test. For example, the weakest precondition of $\text{inc}_x \cdot x > n$ is $x > n - 1$ when n is not zero. We must have domain axioms to justify the weakest precondition relation. For example, the domain axiom: $\text{inc}_x \cdot (x > n) \equiv (x > n - 1) \cdot \text{inc}_x$ ensures that weakest preconditions for inc_x are modeled by the equational theory. The other axioms are used to justify the remaining weakest preconditions that relate other actions and predicates. Additional axioms that do not involve actions, such as $\neg(x > n) \cdot (x > m) \equiv 0$, are included to ensure that the predicate fragment of IncNat is complete in isolation. The deductive completeness of the model shown here can be reduced to Presburger arithmetic.

For the relative ease of defining IncNat , we get real power—we've extended KAT with unbounded state! It is sound to add other operations to IncNat , like multiplication or addition by a scalar. So long as the operations are monotonically increasing and invertible, we can still define a WP and corresponding axioms. It is *not* possible, however, to compare two variables directly with tests like $x = y$ —to do so would not satisfy the requirement that weakest precondition does not grow the size of a test. It would be bad if it did: the test $x = y$ can encode context-free languages! The

197	Syntax	Semantics	
198	$\alpha ::= x > n$	$n \in \mathbb{N} \quad x \in \mathcal{V}$	
199	$\pi ::= \text{inc}_x \mid x := n$	State = $\mathcal{V} \rightarrow \mathbb{N}$	
200	$\text{sub}(x > n) = \{x > m \mid m \leq n\}$	$\text{pred}(x > n, t) = \text{last}(t)(x) > n$	
201		$\text{act}(\text{inc}_x, \sigma) = \sigma[x \mapsto \sigma(x) + 1]$	
202		$\text{act}(x := n, \sigma) = \sigma[x \mapsto n]$	
203	Weakest precondition	Axioms	
204	$x := n \cdot (x > m) \text{ WP } (n > m)$	$\neg(x > n) \cdot (x > m) \equiv 0 \text{ when } n \leq m$	GT-CONTRA
205	$\text{inc}_y \cdot (x > n) \text{ WP } (x > n)$	$x := n \cdot (x > m) \equiv (n > m) \cdot x := n$	ASGN-GT
206	$\text{inc}_x \cdot (x > n) \text{ WP } (x > n - 1)$	$(x > m) \cdot (x > n) \equiv (x > \max(m, n))$	GT-MIN
207	when $n \neq 0$	$\text{inc}_y \cdot (x > n) \equiv (x > n) \cdot \text{inc}_y$	GT-COMM
208	$\text{inc}_x \cdot (x > 0) \text{ WP } 1$	$\text{inc}_x \cdot (x > n) \equiv (x > n - 1) \cdot \text{inc}_x \text{ when } n > 0$	INC-GT
209		$\text{inc}_x \cdot (x > 0) \equiv \text{inc}_x$	INC-GT-Z

Fig. 1. IncNat, increasing naturals

(inadmissible!) term $x := 0 \cdot y := 0; (\text{inc}_x)^* \cdot (\text{inc}_y)^* \cdot x = y$ describes programs with balanced increments of x and y . For the same reason, we cannot safely add a decrement operation dec_x . Either of these would allow us to define counter machines, leading inevitably to undecidability.

Implementation. Users implement KMT's client theories by defining OCaml modules; users give the types of actions and tests along with functions for parsing, computing subterms, calculating weakest preconditions for primitives, mapping predicates to an SMT solver, and deciding predicate satisfiability (see Sec. 5 for more detail).

Our example implementation starts by defining a new, recursive module called IncNat. Recursive modules allow the author of the module to access the final KAT functions and types derived after instantiating KA with their theory within their theory's implementation. For example, the module K on the second line gives us a recursive reference to the resulting KMT instantiated with the IncNat theory; such self-reference is key for higher-order theories, which must embed KAT predicates inside of other kinds of predicates (Sec. 3). The user must define two types: a for tests and p for actions. Tests are of the form $x > n$ where variable names are represented with strings, and values with OCaml ints. Actions hold either the variable being incremented (inc_x) or the variable and value being assigned ($x := n$).

```

230 type a = Gt of string * int (* alpha ::= x > n *)
231 type p = Increment of string (* pi ::= inc x *)
232
233 module rec IncNat : THEORY with type A.t = a and type P.t = p = struct
234   (* generated KMT, for recursive use *)
235   module K = KAT (IncNat)
236   (* boilerplate necessary for recursive modules, hashconsing *)
237   module P : CollectionType with type t = p = struct ... end
238   module A : CollectionType with type t = a = struct ... end
239   (* extensible parser; pushback; subterms of predicates *)
240   let parse name es = ...
241   let push_back p a =
242     match (p,a) with
243     | (Increment x, Gt (_, j)) when 1 > j → PSet.singleton ~cmp:K.Test.compare (K.one ())

```

```

246 | (Increment x, Gt (y, j)) when x = y →
247   PSet.singleton ~cmp:K.Test.compare (K.theory (Gt (y, j - 1)))
248 | (Assign (x,i), Gt (y,j)) when x = y →
249   PSet.singleton ~cmp:K.Test.compare (if i > j then K.one () else K.zero ())
250 | _ → PSet.singleton ~cmp:K.Test.compare (K.theory a)
251 let rec subterms x =
252   match x with
253   | Gt (_, 0) → PSet.singleton ~cmp:K.Test.compare (K.theory x)
254   | Gt (v, i) → PSet.add (K.theory x) (subterms (Gt (v, i - 1)))
255   (* decision procedure for the predicate theory *)
256   let satisfiable (a: K.Test.t) = ...
257 end
258
259
260

```

261 The first function, `parse`, allows the library author to extend the KAT parser (if desired) to
 262 include new kinds of tests and actions in terms of infix and named operators. The other functions,
 263 `subterms` and `push_back`, follow from the KMT theory directly. Finally, the user must implement
 264 a function that decides satisfiability of theory tests.

265 The implementation obligations—syntactic extensions, `subterms` functions, WP on primitives, a
 266 satisfiability checker for the test fragment—mirror our formal development. We offer more client
 267 theories in Sec. 3 and more detail on the implementation in Sec. 5.

268 1.4 Contributions

269 We claim the following contributions:

- 270 • A compositional framework for defining KATs and proving their metatheory, with a novel
 271 development of the normalization procedure used in completeness (Sec. 2) and a new KAT
 272 theorem (PUSHBACK-NEG). Completeness yields a decision procedure based on normalization.
- 273 • Case studies of this framework (Sec. 3), several of which reproduce results from the literature,
 274 and several of which are new: base theories (e.g., naturals, bitvectors [29], networks), and more
 275 importantly, compositional, higher-order theories (e.g., sets and LTL_f). As an example, we
 276 define Temporal NetKAT compositionally [8] by applying the theory of LTL_f to a theory of
 277 NetKAT; doing so strengthens Temporal NetKAT’s completeness result.
- 278 • An automata-theoretic account of our proof technique, proven correct and applicable to
 279 compilation and equivalence checking for, e.g., NetKAT (Sec. 4).
- 280 • An implementation of KMTs (Sec. 5) mirroring our proofs; we derive two equivalence decision
 281 procedures for client theories from just a few definitions, one based on our normalization
 282 routine and one using automata. Our implementation is efficient enough for experimentation
 283 with small programs (Sec. 6).

284 Finally, our framework offers a new way in for those looking to work with KATs. Researchers
 285 comfortable with inductive relations from, e.g., type theory and semantics, will find a familiar friend
 286 in `pushback`, our generalization of weakest preconditions—we define it as an inductive relation. To
 287 restate our contributions for readers more deeply familiar with KAT: Our framework is similar to
 288 Schematic KAT, a KAT extended with first order theories. However, Schematic KAT is incomplete
 289 in general. Our framework shows that a subset of Schematic KATs is complete—those with tracing
 290 semantics and a monotonic `pushback`.

295	Predicates	$\mathcal{T}_{\text{pred}}^*$	Actions
296	a, b	$::= 0$	a <i>embedded predicates</i>
297		1 <i>additive identity</i>	$p + q$ <i>parallel composition</i>
298		$\neg a$ <i>multiplicative identity</i>	$p \cdot q$ <i>sequential composition</i>
299		$a + b$ <i>negation</i>	p^* <i>Kleene star</i>
300		$a \cdot b$ <i>disjunction</i>	π <i>primitive actions (\mathcal{T}_π)</i>
301		α <i>conjunction</i>	
302		<i>primitive predicates (\mathcal{T}_α)</i>	

Fig. 2. \mathcal{T}^* : generalized KAT syntax over a client theory \mathcal{T} (client parts highlighted)

2 THE KMT FRAMEWORK

The rest of our paper describes how our framework takes a client theory and generates a KAT. We emphasize that you need not understand the following mathematics to use our framework—we do it once and for all, so you don't have to. While we have striven to make this section accessible to non-expert readers, those completely new to KATs may do well to skip our discussion of pushback (Sec. 2.3.2 on) and read our case studies (Sec. 3). We **highlight** anything the client theory must provide.

We derive a KAT \mathcal{T}^* (Fig. 2) on top of a client theory \mathcal{T} where \mathcal{T} has two fundamental parts—predicates $\alpha \in \mathcal{T}_\alpha$ and actions $\pi \in \mathcal{T}_\pi$. These are the *primitives* of the client theory. We refer to the Boolean algebra over the client theory as $\mathcal{T}_{\text{pred}}^* \subseteq \mathcal{T}^*$.

Our framework can provide results for \mathcal{T}^* in a pay-as-you-go fashion: given a notion of state and an interpretation for the predicates and actions of \mathcal{T} , we derive a trace semantics for \mathcal{T}^* (Sec. 2.1); if \mathcal{T} has a sound equational theory with respect to our semantics, so does \mathcal{T}^* (Sec. 2.2); if \mathcal{T} has a complete equational theory with respect to our semantics, and satisfies certain weakest precondition requirements, then \mathcal{T}^* has a complete equational theory (Sec. 2.4); and finally, with just a few lines of code defining the structure of \mathcal{T} , we can provide two decision procedures for equivalence (Sec. 5): one using the normalization routine from completeness (Sec. 2.4) and one using automata (Sec. 4).

The key to our general, parameterized proof is a novel *pushback* operation that generalizes weakest preconditions (Sec. 2.3.2): given an understanding of how to push primitive predicates back to the front of a term, we can normalize terms for our completeness proof (Sec. 2.4).

2.1 Semantics

The first step in turning the client theory \mathcal{T} into a KAT is to define a semantics (Fig. 3). We can give any KAT a *trace semantics*: the meaning of a term is a trace t , which is a non-empty list of log entries l . Each *log entry* records a state σ and (in all but the initial state) a primitive action π . The client assigns meaning to predicates and actions by defining a set of states `State` and two functions: one to determine whether a predicate holds (`pred`) and another to determine an action's effects (`act`). To run a \mathcal{T}^* term on a state σ , we start with an initial state $\langle \sigma, \perp \rangle$; when we're done, we'll have a set of traces of the form $\langle \sigma_0, \perp \rangle \langle \sigma_1, \pi_1 \rangle \dots$, where $\sigma_i = \text{act}(\pi_i, \sigma_{i-1})$ for $i > 0$. (A similar semantics shows up in Kozen's application of KAT to static analysis [32].)

The client's `pred` function takes a primitive predicate α and a trace — predicates can examine the entire trace — returning true or false. When the `pred` function returns true, we return the singleton set holding our input trace; when `pred` returns false, we return the empty set. (Composite predicates follow this same pattern, always returning either a singleton set holding their input trace or the empty set.) It's acceptable for the `pred` function to recursively call the denotational

344 **Trace definitions**

345 $\sigma \in \text{State}$
 346 $l \in \text{Log} ::= \langle \sigma, \perp \rangle \mid \langle \sigma, \pi \rangle$
 347 $t \in \text{Trace} = \text{Log}^+$

$\text{pred} : \mathcal{T}_\alpha \times \text{Trace} \rightarrow \{\text{true}, \text{false}\}$
 $\text{act} : \mathcal{T}_\pi \times \text{State} \rightarrow \text{State}$

348 **Trace semantics**

349 $\llbracket 0 \rrbracket(t) = \emptyset$
 350 $\llbracket 1 \rrbracket(t) = \{t\}$
 351 $\llbracket \alpha \rrbracket(t) = \{t \mid \text{pred}(\alpha, t) = \text{true}\}$
 352 $\llbracket \neg a \rrbracket(t) = \{t \mid \llbracket a \rrbracket(t) = \emptyset\}$
 353 $\llbracket \pi \rrbracket(t) = \{t \langle \sigma', \pi \rangle \mid \sigma' = \text{act}(\pi, \text{last}(t))\}$
 354 $\llbracket p + q \rrbracket(t) = \llbracket p \rrbracket(t) \cup \llbracket q \rrbracket(t)$
 355 $\llbracket p \cdot q \rrbracket(t) = (\llbracket p \rrbracket \bullet \llbracket q \rrbracket)(t)$
 356 $\llbracket p^* \rrbracket(t) = \bigcup_{0 \leq i} \llbracket p \rrbracket^i(t)$

$\llbracket - \rrbracket : \mathcal{T}^* \rightarrow \text{Trace} \rightarrow \mathcal{P}(\text{Trace})$

$(f \bullet g)(t) = \bigcup_{t' \in f(t)} g(t')$
 $f^0(t) = \{t\} \quad f^{i+1}(t) = (f \bullet f^i)(t)$
 $\text{last}(\dots \langle \sigma, _ \rangle) = \sigma$

358 **Kleene Algebra axioms**

359 $p + (q + r) \equiv (p + q) + r$ KA-PLUS-ASSOC
 360 $p + q \equiv q + p$ KA-PLUS-COMM
 361 $p + 0 \equiv p$ KA-PLUS-ZERO
 362 $p + p \equiv p$ KA-PLUS-IDEM
 363 $p \cdot (q \cdot r) \equiv (p \cdot q) \cdot r$ KA-SEQ-ASSOC
 364 $1 \cdot p \equiv p$ KA-SEQ-ONE
 365 $p \cdot 1 \equiv p$ KA-ONE-SEQ
 366 $p \cdot (q + r) \equiv p \cdot q + p \cdot r$ KA-DIST-L
 367 $(p + q) \cdot r \equiv p \cdot r + q \cdot r$ KA-DIST-R
 368 $0 \cdot p \equiv 0$ KA-ZERO-SEQ
 369 $p \cdot 0 \equiv 0$ KA-SEQ-ZERO
 370 $1 + p \cdot p^* \equiv p^*$ KA-UNROLL-L
 371 $1 + p^* \cdot p \equiv p^*$ KA-UNROLL-R
 372 $q + p \cdot r \leq r \rightarrow p^* \cdot q \leq r$ KA-LFP-L
 373 $p + q \cdot r \leq q \rightarrow p \cdot r^* \leq q$ KA-LFP-R
 374 $p \leq q \Leftrightarrow p + q \equiv q$

358 **Boolean Algebra axioms**

$a + (b \cdot c) \equiv (a + b) \cdot (a + c)$ BA-PLUS-DIST
 $a + 1 \equiv 1$ BA-PLUS-ONE
 $a + \neg a \equiv 1$ BA-EXCL-MID
 $a \cdot b \equiv b \cdot a$ BA-SEQ-COMM
 $a \cdot \neg a \equiv 0$ BA-CONTRA
 $a \cdot a \equiv a$ BA-SEQ-IDEM

358 **Consequences**

$p \cdot a \equiv b \cdot p \rightarrow p \cdot \neg a \equiv \neg b \cdot p$ PUSHBACK-NEG
 $p \cdot (q \cdot p)^* \equiv (p \cdot q)^* \cdot p$ SLIDING
 $(p + q)^* \equiv p^* \cdot (q \cdot p^*)^*$ DENESTING
 $p \cdot a \equiv a \cdot q + r \rightarrow$
 $p^* \cdot a \equiv (a + p^* \cdot r) \cdot q^*$ STAR-INV
 $p \cdot a \equiv a \cdot q + r \rightarrow$
 $p \cdot a \cdot (p \cdot a)^* \equiv (a \cdot q + r) \cdot (q + r)^*$ STAR-EXPAND

375 Fig. 3. Semantics and equational theory for \mathcal{T}^*

376 semantics, though we have skipped the formal detail here. This way we can define composite
 377 primitive predicates as in, e.g., temporal logic (Sec. 3.6).

378 The client's act function takes a primitive action π and the last state in the trace, returning a new
 379 state. Whatever new state comes out is recorded in the trace, along with the action just performed.

382 **2.2 Soundness**

383 Proving that the equational theory is sound is relatively straightforward: we only depend on the
 384 client's act and pred functions, and none of our KAT axioms (Fig. 3) even mention the client's
 385 primitives. We believe the pushback negation theorem (PUSHBACK-NEG) is novel (though it holds
 386 in all KATs). Our soundness proof naturally enough requires that any equations the client theory
 387 adds need to be sound in our trace semantics. We do need to use several KAT theorems in our
 388 completeness proof (Fig. 3, Consequences), the most complex being star expansion (STAR-EXPAND),
 389 which we take from Temporal NetKAT [8]; we believe PUSHBACK-NEG is a novel theorem that holds
 390 in all KATs.

393 THEOREM 2.1 (SOUNDNESS OF \mathcal{T}^*). *If \mathcal{T} 's equational reasoning is sound ($p \equiv_{\mathcal{T}} q \Rightarrow \llbracket p \rrbracket = \llbracket q \rrbracket$)*
 394 *then \mathcal{T}^* 's equational reasoning is sound ($p \equiv q \Rightarrow \llbracket p \rrbracket = \llbracket q \rrbracket$).*

395 PROOF. By induction on the derivation of $p \equiv q$.¹ □

397 2.3 Normalization via pushback

398 In order to prove completeness (Sec. 2.4), we reduce our KAT terms to a more manageable subset
 399 of *normal forms*. Normalization happens via a generalization of weakest preconditions; we use
 400 a *pushback* operation to translate a term p into an equivalent term of the form $\sum a_i \cdot m_i$ where
 401 each m_i does not contain any tests. Once in this form, we can use the completeness result provided
 402 by the client theory to reduce the completeness of our language to an existing result for Kleene
 403 algebra.

404 In order to use our general normalization procedure, the client theory \mathcal{T} must define two things:

- 405 (1) a way to extract subterms from predicates, to define an ordering on predicates that serves as
 406 the termination measure on normalization (Sec. 2.3.1); and
- 407 (2) weakest preconditions for primitives (Sec. 2.3.2).

408 Once we've defined our normalization procedure, we can use it prove completeness (Sec. 2.4).

409
 410 **2.3.1 Normalization and the maximal subterm ordering.** Our normalization algorithm works by
 411 “pushing back” predicates to the front of a term until we reach a normal form with *all* predicates at
 412 the front. The pushback algorithm's termination measure is a complex one. For example, pushing a
 413 predicate back may not eliminate the predicate even though progress was made in getting predicates
 414 to the front. More trickily, it may be that pushing test a back through π yields $\sum a_i \cdot \pi$ where each
 415 of the a_i is a copy of some subterm of a —and there may be *many* such copies!

416 Let the set of *restricted actions* \mathcal{T}_{RA} be the subset of \mathcal{T}^* where the only test is 1. We will use
 417 metavariables m, n , and l to denote elements of \mathcal{T}_{RA} . Let the set of *normal forms* $\mathcal{T}_{\text{nf}}^*$ be a set of
 418 pairs of tests $a_i \in \mathcal{T}_{\text{pred}}^*$ and restricted actions $m_i \in \mathcal{T}_{\text{RA}}$. We will use metavariables t, u, v, w, x, y ,
 419 and z to denote elements of $\mathcal{T}_{\text{nf}}^*$; we typically write these sets not in set notation, but as sums, i.e.,
 420 $x = \sum_{i=1}^k a_i \cdot m_i$ means $x = \{(a_1, m_1), (a_2, m_2), \dots, (a_k, m_k)\}$. The sum notation is convenient, but
 421 it is important that normal forms really be treated as sets—there should be no duplicated terms
 422 in the sum. We write $\sum_i a_i$ to denote the normal form $\sum_i a_i \cdot 1$. The set of normal forms, $\mathcal{T}_{\text{nf}}^*$, is
 423 closed over parallel composition by simply joining the sums. The fundamental challenge in our
 424 normalization method is to define sequential composition and Kleene star on $\mathcal{T}_{\text{nf}}^*$.

425 Our normalization algorithm uses the *maximal subterm ordering* as its termination measure.
 426 Due to space constraints, we provide the formal definitions of maximal tests and subterms in the
 427 supplemental material. Here we simply give intuition for the two relevant high-level operations:
 428 $\text{mt}(x) \subseteq \mathcal{T}_{\text{pred}}^*$ computes the *maximal tests* of a normal form x , which are those tests that are not
 429 subterms of any other test; the maximal subterm ordering $x \leq y$ for normal forms holds when
 430 the subterms of x 's maximal tests are a subset of the subterms of y 's maximal tests. Informally,
 431 we have $x \leq y$ when every test in x is somehow “covered” by a test in y ; we have $x < y$ when
 432 $x \leq y$ and y has some maximal test x that does not. Our definition of subterms needs the client
 433 theory to identify the subterms of its primitives via a function $\text{sub}_{\mathcal{T}}$ such that (1) if $b \in \text{sub}_{\mathcal{T}}(a)$
 434 then $\text{sub}(b) \subseteq \text{sub}_{\mathcal{T}}(a)$ and (2) if $b \in \text{sub}_{\mathcal{T}}(a)$, then either $b \in \{0, 1, a\}$ or b precedes a in a global
 435 well ordering of predicates.

436
 437 LEMMA 2.2 (SPLITTING). *If $a \in \text{mt}(x)$, then there exist y and z such that $x \equiv a \cdot y + z$ and $y < x$*
 438 *and $z < x$.*

439
 440 ¹Full proofs with all necessary lemmas are available in an extended version of this paper in the supplementary material.

442 Splitting is the key lemma for making progress pushing tests back, allowing us to take a normal
 443 form and slowly push its maximal tests to the front; its proof follows from a chain of lemmas given
 444 in the supplementary material.

445 **2.3.2 Pushback.** In order to define normalization—necessary for completeness (Sec. 2.4)—the
 446 client theory must have a *weakest preconditions* operation that respects the subterm ordering.
 447

448 *Definition 2.3 (Weakest preconditions).* The *weakest precondition* operation of the client theory is
 449 a relation $WP \subseteq \mathcal{T}_\pi \times \mathcal{T}_\alpha \times \mathcal{P}(\mathcal{T}_{\text{pred}}^*)$, where \mathcal{T}_π are the primitive actions and \mathcal{T}_α are the primitive
 450 predicates of \mathcal{T} . We write the relation as $\pi \cdot \alpha \text{ WP } \sum a_i \cdot \pi$ and read it as “ α pushes back through π
 451 to yield $\sum a_i \cdot \pi$ ”; the second π is redundant but written for clarity. We require that if $\pi \cdot \alpha \text{ WP}$
 452 $\{a_1, \dots, a_k\} \cdot \pi$, then $\pi \cdot \alpha \equiv \sum_{i=1}^k a_i \cdot \pi$, and $a_i \leq \alpha$.
 453

454 Given the client theory’s weakest-precondition relation WP, we define a normalization procedure
 455 for \mathcal{T}^* by extending the client’s WP relation to a more general *pushback* relation, PB (Fig. 4). The
 456 client’s WP relation need not be a function, nor do the a_i need to be obviously related to α or π in
 457 any way. Even when the WP relation is a function, the PB relation will generally not be a function.
 458 While WP computes the classical weakest precondition, the PB relations do something different:
 459 when pushing back we have the freedom to *change the program itself*—not normally an option for
 460 weakest preconditions (see Sec. 7).

461 We define the top-level normalization routine with the p norm x relation (Fig. 4), a syntax
 462 directed relation that takes a term p and produces a normal form $x = \sum_i a_i m_i$. Most syntactic forms
 463 are easy to normalize: predicates are already normal forms (PRED); primitive actions π are normal
 464 forms where there’s just one summand and the predicate is 1 (ACT); and parallel composition of
 465 two normal forms means just joining the sums (PAR). But sequence and Kleene star are harder: we
 466 define judgments using PB to lift these operations to normal forms (SEQ, STAR).

467 For sequences, we can recursively take $p \cdot q$ and normalize p into $x = \sum a_i \cdot m_i$ and q into
 468 $y = \sum b_j \cdot n_j$. But how can we combine x and y into a new normal form? We can concatenate
 469 and rearrange the normal forms to get $\sum_{i,j} a_i \cdot m_i \cdot b_j \cdot n_j$. If we can push b_j back through m_i to
 470 find some new normal form $\sum c_k \cdot l_k$, then $\sum_{i,j,k} a_i \cdot c_k \cdot l_k \cdot n_j$ is a normal form (JOIN); we write
 471 $x \cdot y \text{ PB}^J z$ to mean that the concatenation of x and y is equivalent to the normal form z —the \cdot is
 472 suggestive notation.

473 For Kleene star, we can take p^* and normalize p into $x = \sum a_i \cdot m_i$, but x^* isn’t a normal form—we
 474 need to somehow move all of the tests out of the star and to the front. We do so with the PB^*
 475 relation (Fig. 4), writing $x^* \text{ PB}^* y$ to mean that the Kleene star of x is equivalent to the normal form
 476 y —the $*$ on the left is again suggestive notation. The PB^* relation is more subtle than PB^J . There are
 477 four possible ways to treat x , based on how it splits (Lemma 2.2): if $x = 0$, then our work is trivial
 478 since $0^* \equiv 1$ (STARZERO); if x splits into $a \cdot x'$ where a is a maximal test and there are no other
 479 summands, then we can either use the KAT sliding lemma to pull the test out when a is strictly the
 480 largest test in x (SLIDE) or by using the KAT expansion lemma (EXPAND); if x splits into $a \cdot x' + z$,
 481 we use the KAT denesting lemma to pull a out before recurring on what remains (DENEST).

482 The bulk of the pushback’s work happens in the PB^\bullet relation, which pushes a test back through
 483 a restricted action; PB^R and PB^T use PB^\bullet to push tests back through normal forms and normal
 484 forms back through restricted actions, respectively. To handle negation, the function nnf translates
 485 predicates to *negation normal form*, where negations only appear on primitive predicates (Fig. 4);
 486 PUSHBACK-NEG justifies this case.

487 We show that our notion of pushback is correct in two steps. First we prove that pushback is
 488 partially correct, i.e., if we can form a derivation in the pushback relations, the right-hand sides
 489 are equivalent to the left-hand-sides (Theorem 2.4). Once we’ve established that our pushback
 490

491
492
493
494
495
496
497

Normalization

$$\begin{array}{c}
 \boxed{p \text{ norm } x} \\
 \frac{}{a \text{ norm } a} \text{ PRED} \qquad \frac{}{\pi \text{ norm } 1 \cdot \pi} \text{ ACT} \qquad \frac{p \text{ norm } x \quad q \text{ norm } y}{p + q \text{ norm } x + y} \text{ PAR} \\
 \frac{p \text{ norm } x \quad q \text{ norm } y \quad x \cdot y \text{ PB}^\perp z}{p \cdot q \text{ norm } z} \text{ SEQ} \qquad \frac{p \text{ norm } x \quad x^* \text{ PB}^* y}{p^* \text{ norm } y} \text{ STAR}
 \end{array}$$

498
499
500
501
502

Sequential composition of normal forms

$$\frac{}{x \cdot y \text{ PB}^\perp z} \\
 \frac{m_i \cdot b_j \text{ PB}^* x_{ij}}{(\sum_i a_i \cdot m_i) \cdot (\sum_j b_j \cdot n_j) \text{ PB}^\perp \sum_i \sum_j a_i \cdot x_{ij} \cdot n_j} \text{ JOIN}$$

503
504
505
506
507
508
509

Normalization of star

$$\begin{array}{c}
 \boxed{x^* \text{ PB}^* y} \\
 \frac{}{0^* \text{ PB}^* 1} \text{ STARZERO} \qquad \frac{x < a \quad x \cdot a \text{ PB}^\top y \quad y^* \text{ PB}^* y' \quad y' \cdot x \text{ PB}^\perp z}{(a \cdot x)^* \text{ PB}^* 1 + a \cdot z} \text{ SLIDE} \\
 \frac{x \not< a \quad x \cdot a \text{ PB}^\top a \cdot t + u \quad (t + u)^* \text{ PB}^* y \quad y \cdot x \text{ PB}^\perp z}{(a \cdot x)^* \text{ PB}^* 1 + a \cdot z} \text{ EXPAND} \qquad \frac{a \notin \text{mt}(z) \quad y \neq 0 \quad y^* \text{ PB}^* y'}{x \cdot y' \text{ PB}^\perp x' \quad (a \cdot x')^* \text{ PB}^* z \quad y' \cdot z \text{ PB}^\perp z'} \text{ DENEST}
 \end{array}$$

510
511
512
513
514
515
516
517
518
519

Pushback

$$\begin{array}{c}
 \boxed{m \cdot a \text{ PB}^* y} \qquad \boxed{m \cdot x \text{ PB}^R y} \qquad \boxed{x \cdot a \text{ PB}^\top y} \\
 \frac{}{m \cdot 0 \text{ PB}^* 0} \text{ SEQZERO} \qquad \frac{}{m \cdot 1 \text{ PB}^* 1 \cdot m} \text{ SEQONE} \\
 \frac{m \cdot a \text{ PB}^* y \quad y \cdot b \text{ PB}^\top z}{m \cdot (a \cdot b) \text{ PB}^* z} \text{ SEQSEQTEST} \qquad \frac{n \cdot a \text{ PB}^* x \quad m \cdot x \text{ PB}^R y}{(m \cdot n) \cdot a \text{ PB}^* y} \text{ SEQSEQACTION} \\
 \frac{m \cdot a \text{ PB}^* x \quad m \cdot b \text{ PB}^* y}{m \cdot (a + b) \text{ PB}^* x + y} \text{ SEQPARTEST} \qquad \frac{m \cdot a \text{ PB}^* x \quad n \cdot a \text{ PB}^* y}{(m + n) \cdot a \text{ PB}^* x + y} \text{ SEQPARACTION}
 \end{array}$$

520
521
522
523
524
525
526
527
528
529

$$\begin{array}{c}
 \frac{\pi \cdot \alpha \text{ WP } \{a_1, \dots\}}{\pi \cdot \alpha \text{ PB}^* \sum_i a_i \cdot \pi} \text{ PRIM} \qquad \frac{\pi \cdot a \text{ PB}^* \sum_i a_i \cdot \pi \quad \text{nnf}(\neg(\sum_i a_i)) = b}{\pi \cdot \neg a \text{ PB}^* b \cdot \pi} \text{ PRIMNEG} \\
 \frac{m \cdot a \text{ PB}^* x \quad x < a \quad m^* \cdot x \text{ PB}^R y}{m^* \cdot a \text{ PB}^* a + y} \text{ SEQSTARSMALLER} \qquad \frac{m \cdot a \text{ PB}^* a \cdot t + u \quad m^* \cdot u \text{ PB}^R x \quad t^* \text{ PB}^* y \quad x \cdot y \text{ PB}^\perp z}{m^* \cdot a \text{ PB}^* a \cdot y + z} \text{ SEQSTARINV} \\
 \frac{m \cdot a_i \text{ PB}^* x_i}{m \cdot \sum_i a_i \cdot n_i \text{ PB}^R \sum_i x_i \cdot n_i} \text{ RESTRICTED} \qquad \frac{m_i \cdot a \text{ PB}^* \sum_j b_{ij} \cdot m_{ij}}{(\sum_i a_i \cdot m_i) \cdot a \text{ PB}^\top \sum_i \sum_j a_i \cdot b_{ij} \cdot m_{ij}} \text{ TEST}
 \end{array}$$

530
531
532
533
534
535
536
537

Negation normal form

$$\begin{array}{c}
 \boxed{\text{nnf} : \mathcal{T}^*_{\text{pred}} \rightarrow \mathcal{T}^*_{\text{pred}}} \\
 \begin{array}{ll}
 \text{nnf}(0) = 0 & \text{nnf}(\neg 0) = 1 \\
 \text{nnf}(1) = 1 & \text{nnf}(\neg 1) = 0 \\
 \text{nnf}(\alpha) = \alpha & \text{nnf}(\neg \alpha) = \neg \alpha \\
 \text{nnf}(a + b) = \text{nnf}(a) + \text{nnf}(b) & \text{nnf}(\neg \neg a) = \text{nnf}(a) \\
 \text{nnf}(a \cdot b) = \text{nnf}(a) \cdot \text{nnf}(b) & \text{nnf}(\neg(a + b)) = \text{nnf}(\neg a) \cdot \text{nnf}(\neg b) \\
 & \text{nnf}(\neg(a \cdot b)) = \text{nnf}(\neg a) + \text{nnf}(\neg b)
 \end{array}
 \end{array}$$

538
539

Fig. 4. Normalization for \mathcal{T}^*

relations' derivations mean what we want, we have to show that we can find such derivations; here we use our maximal subterm measure to show that the recursive tangle of our PB relations always terminates (Theorem 2.5).

THEOREM 2.4 (PUSHBACK SOUNDNESS). *For each of the PB relations, the left side is equivalent to the right side, e.g., if $x^* \text{ PB}^* y$ then $x^* \equiv y$.*

PROOF. By simultaneous induction on the derivations. Most cases follow by the IH and axioms, with a few relying on KAT theorems like sliding, denesting, star expansion [8], and pushback negation (Fig. 3, Consequences). \square

THEOREM 2.5 (PUSHBACK EXISTENCE). *For each of the PB relations, every left side relates to a right side that is no larger, e.g., for all x there exists $y \leq x$ such that $x^* \text{ PB}^* y$.*

PROOF. By induction on the lexicographical order of: the subterm ordering ($<$); the size of x ; the size of m ; and the size of a . Cases go by using splitting (Lemma 2.2) to show that derivations exist followed by subterm ordering congruence to find orderings to apply the IH. \square

Finally, to reiterate our discussion of PB^* , Theorem 2.5 shows that every left-hand side of the pushback relation has a corresponding right-hand side. We *haven't* proved that the pushback relation is functional— if a term has more than one maximal test, there could be many different choices of how we perform the pushback.

Now that we can push back tests, we can show that every term has an equivalent normal form.

COROLLARY 2.6 (NORMAL FORMS). *For all $p \in \mathcal{T}^*$, there exists a normal form x such $p \text{ norm } x$ and that $p \equiv x$.*

PROOF. By induction on p , using Theorems 2.5 and 2.4 in the SEQ and STAR case. \square

The PB relations and these two proofs are one of the contributions of this paper: we believe it is the first time that a KAT normalization procedure has been made so explicit, rather than hiding inside of completeness proofs. Temporal NetKAT, which introduced the idea of pushback, proved a concretization of Theorems 2.4 and 2.5 as a single theorem and without any explicit normalization or pushback relation.

2.4 Completeness

We prove completeness—if $\llbracket p \rrbracket = \llbracket q \rrbracket$ then $p \equiv q$ —by normalizing p and q and comparing the resulting terms. Our completeness proof uses the completeness of Kleene algebra (KA) as its foundation: the set of possible traces of actions performed for a restricted (test-free) action in our denotational semantics is a regular language, and so the KA axioms are sound and complete for it. In order to relate our denotational semantics to regular languages, we define the regular interpretation of restricted actions $m \in \mathcal{T}_{\text{RA}}$ in the conventional way and then relate our denotational semantics to the regular interpretation (Fig. 5). Readers familiar with NetKAT's completeness proof may notice that we've omitted the language model and gone straight to the regular interpretation. We're able to shorten our proof because our tracing semantics is more directly relatable to its regular interpretation, and because our completeness proof separately defers to the client theory's decision procedure for the predicates at the front. Our normalization routine—the essence of our proof—only uses the KAT axioms and doesn't rely on any property of our tracing semantics. We conjecture that one could prove a similar completeness result and derive a similar decision procedure with a merging, non-tracing semantics, like in NetKAT or KAT+B! [1, 29]. We haven't attempted the proof or an implementation.

589	$\mathcal{R} : \mathcal{T}_{RA} \rightarrow \mathcal{P}(\Pi_{\mathcal{T}}^*)$		$\text{label} : \text{Trace} \rightarrow \Pi_{\mathcal{T}}^*$
590	$\mathcal{R}(1) = \{\epsilon\}$		$\text{label}(\langle\sigma, \perp\rangle) = \epsilon$
591	$\mathcal{R}(\pi) = \{\pi\}$		$\text{label}(t\langle\sigma, \pi\rangle) = \text{label}(t)\pi$
592	$\mathcal{R}(m+n) = \mathcal{R}(m) \cup \mathcal{R}(n)$		
593	$\mathcal{R}(m \cdot n) = \{uv \mid u \in \mathcal{R}(m), v \in \mathcal{R}(n)\}$		$\mathcal{L}^0 = \{\epsilon\}$
594	$\mathcal{R}(m^*) = \bigcup_{0 \leq i} \mathcal{R}(m)^i$		$\mathcal{L}^{n+1} = \{uv \mid u \in \mathcal{L}, v \in \mathcal{L}^n\}$

Fig. 5. Regular interpretation of restricted actions

LEMMA 2.7 (LABELS ARE REGULAR). $\{\text{label}(\llbracket m \rrbracket(\langle\sigma, \perp\rangle)) \mid \sigma \in \text{State}\} = \mathcal{R}(m)$

PROOF. By induction on the restricted action m . □

THEOREM 2.8 (COMPLETENESS). *If the emptiness of \mathcal{T} predicates is decidable, then if $\llbracket p \rrbracket = \llbracket q \rrbracket$ then $p \equiv q$.*

PROOF. There must exist normal forms x and y such that p norm x and q norm y and $p \equiv x$ and $q \equiv y$ (Corollary 2.6); by soundness (Theorem 2.1), we can find that $\llbracket p \rrbracket = \llbracket x \rrbracket$ and $\llbracket q \rrbracket = \llbracket y \rrbracket$, so it must be the case that $\llbracket x \rrbracket = \llbracket y \rrbracket$. We will find a proof that $x \equiv y$; we can then transitively construct a proof that $p \equiv q$.

We have $x = \sum_i a_i \cdot m_i$ and $y = \sum_j b_j \cdot n_j$. In principle, we ought to be able to match up each of the a_i with one of the b_j and then check to see whether m_i is equivalent to n_j (by appealing to the completeness on Kleene algebra). But we can't simply do a syntactic matching—we could have a_i and b_j that are in effect equivalent, but not obviously so. Worse still, we could have a_i and $a_{i'}$ equivalent! We need to perform two steps of disambiguation: first each normal form must be unambiguous on its own, and then they must be pairwise unambiguous between the two normal forms.

To construct independently unambiguous normal forms, we explode our normal form x into a disjoint form \hat{x} , where we test each possible combination of a_i and run the actions corresponding to the true predicates, i.e., m_i gets run precisely when a_i is true:

$$\begin{aligned}
 \hat{x} = & a_1 \cdot a_2 \cdots a_n \cdot m_1 \cdot m_2 \cdots m_n \\
 & + \neg a_1 \cdot a_2 \cdots a_n \cdot m_2 \cdots m_n \\
 & + a_1 \cdot \neg a_2 \cdots a_n \cdot m_1 \cdots m_n \\
 & + \dots \\
 & + \neg a_1 \cdot \neg a_2 \cdots a_n \cdot m_n
 \end{aligned}$$

and similarly for \hat{y} . We can find $x \equiv \hat{x}$ via distributivity (BA-PLUS-DIST) and the excluded middle (BA-EXCL-MID).

Given normal forms with locally disjoint cases, we can take the Cartesian product of \hat{x} and \hat{y} , which allows us to do a *syntactic* comparison on each of the predicates. Let \check{x} and \check{y} be the extension of \hat{x} and \hat{y} with the tests from the other form, giving us $\check{x} = \sum_{i,j} c_i \cdot d_j \cdot l_i$ and $\check{y} = \sum_{i,j} c_i \cdot d_j \cdot m_j$. Extending the normal forms to be disjoint between the two normal forms is still provably equivalent using commutativity (BA-SEQ-COMM), distributivity (BA-PLUS-DIST), and the excluded middle (BA-EXCL-MID).

Now that each of the predicates are syntactically uniform and disjoint, we can proceed to compare the commands. But there is one final risk: what if the $c_i \cdot d_j \equiv 0$? Then l_i and o_j could safely be different. We therefore use the client's emptiness checker to eliminate those cases where the expanded tests at the front of \check{x} and \check{y} are equivalent to zero, which is sound by the client theory's completeness and zero-cancellation (KA-ZERO-SEQ and KA-SEQ-ZERO).

638	Syntax	638	Semantics
639	$\alpha ::= b = \text{true}$	639	$b \in \mathcal{B}$
640	$\pi ::= b := \text{true} \mid b := \text{false}$	640	State = $\mathcal{B} \rightarrow \{\text{true}, \text{false}\}$
641	$\text{sub}(\alpha) = \{\alpha\}$	641	$\text{pred}(b = \text{true}, t) = \text{last}(t)(b)$
642		642	$\text{act}(b := \text{true}, \sigma) = \sigma[b \mapsto \text{true}]$
643		643	$\text{act}(b := \text{false}, \sigma) = \sigma[b \mapsto \text{false}]$
644	Weakest precondition	644	Axioms
645	$b := \text{true} \cdot b = \text{true} \text{ WP } 1$	645	$(b := \text{true}) \cdot (b = \text{true}) \equiv (b := \text{true}) \quad \text{SET-TEST-TRUE-TRUE}$
646	$b := \text{false} \cdot b = \text{true} \text{ WP } 0$	646	$(b := \text{false}) \cdot (b = \text{true}) \equiv 0 \quad \text{SET-TEST-FALSE-TRUE}$

Fig. 6. BitVec, theory of bitvectors

650 Finally, we can defer to deductive completeness for KA to find proofs that the commands are equal.
 651 To use KA's completeness to get a proof over commands, we have to show that if our commands
 652 have equal denotations in our semantics, then they will also have equal denotations in the KA
 653 semantics. We've done exactly this by showing that restricted actions have regular interpretations:
 654 because the zero-canceled \dot{x} and \dot{y} are provably equal, soundness guarantees that their denotations
 655 are equal. Since their tests are pairwise disjoint, if their denotations are equal, it must be that
 656 any non-canceled commands are equal, which means that each label of these commands must be
 657 equal—and so $\mathcal{R}(l_i) = \mathcal{R}(o_j)$ (Lemma 2.7). By the deductive completeness of KA, we know that
 658 $\text{KA} \vdash l_i \equiv o_j$. Since we have the KA axioms in our system, then $l_i \equiv o_j$; by reflexivity, we know that
 659 $c_i \cdot d_j \equiv c_i \cdot d_j$, and we have proved that $\dot{x} \equiv \dot{y}$. By transitivity, we can see that $\hat{x} \equiv \hat{y}$ and so $x \equiv y$
 660 and $p \equiv q$, as desired. \square

662 3 CASE STUDIES

663 In this section, we define KAT client theories for bitvectors and networks, as well as higher-order
 664 theories for products of theories, sets over theories, and temporal logic over theories.

666 3.1 Bit vectors

667 The simplest KMT is bit vectors: we extend KAT with some finite number of bits, each of which
 668 can be set to true or false and tested for their current value (Fig. 6). The theory adds actions
 669 $b := \text{true}$ and $b := \text{false}$ for boolean variables b , and tests of the form $b = \text{true}$, where b is
 670 drawn from some set of names \mathcal{B} . Since our bit vectors are embedded in a KAT, we can use KAT
 671 operators to build up encodings on top of bits: $b = \text{false}$ desugars to $\neg(b = \text{true})$; flip b desugars to
 672 $(b = \text{true} \cdot b := \text{false}) + (b = \text{false} \cdot b := \text{true})$. We could go further and define numeric operators
 673 on collections of bits, at the cost of producing larger terms. We are not limited to just numbers, of
 674 course; once we have bits, we can encode any bounded data structure we like.

675 KAT+B! [29] develops a nearly identical theory, though our semantics admit different equations.
 676 We use a *trace* semantics, where we distinguish between $(b := \text{true} \cdot b := \text{true})$ and $(b := \text{true})$.
 677 Even though the final states are equivalent, they produce different traces because they run different
 678 actions. KAT+B!, on the other hand, doesn't distinguish based on the trace of actions, so they find
 679 that $(b := \text{true} \cdot b := \text{true}) \equiv (b := \text{true})$. It's difficult to say whether one model is better than the
 680 other—we imagine that either could be appropriate, depending on the setting. For example, our
 681 trace semantics is useful for answering model-checking-like questions (Sec. 3.4).

683 3.2 Disjoint products

684 Given two client theories, we can combine them into a disjoint product theory, $\text{Prod}(\mathcal{T}_1, \mathcal{T}_2)$, where
 685 the states are products of the two sub-theory's states and the predicates and actions from \mathcal{T}_1 can't
 686

687	Syntax	687	Semantics
688	$\alpha ::= \alpha_1 \mid \alpha_2$	688	State = State ₁ × State ₂
689	$\pi ::= \pi_1 \mid \pi_2$	689	pred(α_i, t) = pred _i (α_i, t_i)
690	sub(α_i) = sub _i (α_i)	690	act(π_i, σ) = $\sigma[\sigma_i \mapsto \text{act}_i(\pi_i, \sigma_i)]$
691	Weakest precondition extending $\overline{\mathcal{T}}_1$ and $\overline{\mathcal{T}}_2$		Axioms extending $\overline{\mathcal{T}}_1$ and $\overline{\mathcal{T}}_2$
692	$\pi_1 \cdot \alpha_2$ WP α_2	$\pi_2 \cdot \alpha_1$ WP α_1	$\pi_1 \cdot \alpha_2 \equiv \alpha_2 \cdot \pi_1$ L-R-COMM
693			$\pi_2 \cdot \alpha_1 \equiv \alpha_1 \cdot \pi_2$ R-L-COMM

Fig. 7. Prod($\overline{\mathcal{T}}_1, \overline{\mathcal{T}}_2$), products of two disjoint theories

697	Syntax	697	Semantics
698	$\alpha ::= \text{in}(x, c) \mid e = c \mid \alpha_e$	698	$c \in \mathcal{C}$
699	$\pi ::= \text{add}(x, e) \mid \pi_e$	699	$e \in \mathcal{E}$
700	sub(in(x, c)) = {in(x, c)} \cup sub($\neg(e = c)$)	700	$x \in \mathcal{V}$
701	sub($e = c$) = sub($e = c$)	701	State = ($\mathcal{V} \rightarrow \mathcal{P}(\mathcal{C})$) \times ($\mathcal{E} \rightarrow \mathcal{C}$)
702	sub(α_e) = sub(α_e)	702	pred(in(x, c), t) = last(t) ₂ (c) \in last(t) ₁ (x)
703		703	pred(α_e, t) = pred(α_e, t_2)
704		704	act(add(x, e), σ) = $\sigma[\sigma_1[x \mapsto \sigma_1(x) \cup \{\sigma(e)\}]]$
705		705	act(π_e, σ) = $\sigma[\sigma_2 \mapsto \text{act}(\pi_e, \sigma_2)]$
706	Weakest precondition extending \mathcal{E}		Axioms extending \mathcal{E}
707	add(y, e) \cdot in(x, c) WP in(x, c)	707	add(y, e) \cdot in(x, c) \equiv in(x, c) \cdot add(y, e) ADD-COMM
708	add(x, e) \cdot in(x, c) WP ($e = c$) + in(x, c)	708	add(x, e) \cdot in(x, c) \equiv ($(e = c) + \text{in}(x, c)$) \cdot add(x, e) ADD-IN
709	add(x, e) $\cdot \alpha_e$ WP α_e	709	add(x, e) $\cdot \alpha_e \equiv \alpha_e \cdot \text{add}(x, e)$ ADD-COMM2

Fig. 8. Set(\mathcal{E}), unbounded sets over expressions

affect $\overline{\mathcal{T}}_2$ and vice versa (Fig. 7). We explicitly give definitions for pred and act that defer to the corresponding sub-theory, using t_i to project the trace state to the i th component. It may seem that disjoint products don't give us much, but they in fact allow for us to simulate much more interesting languages in our derived KATs. For example, Prod(BitVec, IncNat) allows us to program with both variables valued as either booleans or (increasing) naturals; the product theory lets us directly express the sorts of programs that Kozen's early static analysis work had to encode manually, i.e., loops over boolean and numeric state [32].

3.3 Unbounded sets

We define a KMT for unbounded sets parameterized on a theory of expressions \mathcal{E} (Fig. 8). The set data type supports just one operation: add(x, e) adds the value of expression e to set x (we could add del(x, e), but we omit it to save space). It also supports a single test: in(x, c) checks if the constant c is contained in set x . The idea is that $e \in \mathcal{E}$ refers to expressions with, say, variables x and constants c . We allow arbitrary expressions e in some positions and constants c in others. (If we allowed expressions in all positions, WP wouldn't necessarily be non-increasing.)

To instantiate the Set theory, we need a few things: expressions \mathcal{E} , a subset of constants $C \subseteq \mathcal{E}$, and predicates for testing (in)equality between expressions and constants ($e = c$ and $e \neq c$). (We can not, in general, expect tests for equality of non-constant expressions, as it may cause us to accidentally define a counter machine.) We treat these two extra predicates as inputs, and expect that they have well behaved subterms. Our state has two parts: $\sigma_1 : \mathcal{V} \rightarrow \mathcal{P}(\mathcal{C})$ records the current sets for each set in \mathcal{V} , while $\sigma_2 : \mathcal{E} \rightarrow \mathcal{C}$ evaluates expressions in each state. Since each state has its own evaluation function, the expression language can have actions that update σ_2 .

Syntax	Semantics
$\alpha ::= \bigcirc a \mid a \mathcal{S} b \mid a$	State = $\text{State}_{\mathcal{T}}$
$\pi ::= \pi_{\mathcal{T}}$	$\text{pred}(\bigcirc a, \langle \sigma, l \rangle) = \text{false}$
$\text{sub}(\bigcirc a) = \{\bigcirc a\} \cup \text{sub}(a)$	$\text{pred}(\bigcirc a, t \langle \sigma, l \rangle) = \text{pred}(a, t)$
$\text{sub}(a \mathcal{S} b) = \{a \mathcal{S} b\} \cup \text{sub}(a) \cup \text{sub}(b)$	$\text{pred}(a \mathcal{S} b, \langle \sigma, l \rangle) = \text{pred}(b, \langle \sigma, l \rangle)$
$\text{act}(\pi, \sigma) = \text{act}(\pi, \sigma)$	$\text{pred}(a \mathcal{S} b, t \langle \sigma, l \rangle) = \text{pred}(b, t \langle \sigma, l \rangle) \vee$ $(\text{pred}(a, t \langle \sigma, l \rangle) \wedge \text{pred}(a \mathcal{S} b, t))$
$\bullet a = \neg \bigcirc \neg a \quad a \mathcal{B} b = a \mathcal{S} b + \square a$	
$\text{start} = \neg \bigcirc 1 \quad \diamond a = 1 \mathcal{S} a \quad \square a = \neg \diamond \neg a$	
Weakest precondition extending \mathcal{T}	Axioms extending \mathcal{T}
$\pi \cdot \bigcirc a \text{ WP } a$	inherited from \mathcal{T}
$\frac{\pi \cdot a \text{ PB}^{\bullet}_{\mathcal{T}} a' \cdot \pi \quad \pi \cdot b \text{ PB}^{\bullet}_{\mathcal{T}} b' \cdot \pi}{\pi \cdot (a \mathcal{S} b) \text{ WP } b' + a' \cdot (a \mathcal{S} b)}$	$\bigcirc(a \cdot b) \equiv \bigcirc a \cdot \bigcirc b$ LTL-LAST-DIST-SEQ
	$\bigcirc(a + b) \equiv \bigcirc a + \bigcirc b$ LTL-LAST-DIST-PLUS
	$\bullet 1 \equiv 1$ LTL-WLAST-ONE
	$a \mathcal{S} b \equiv b + a \cdot \bigcirc(a \mathcal{S} b)$ LTL-SINCE-UNROLL
	$\neg(a \mathcal{S} b) \equiv (\neg b) \mathcal{B} (\neg a \cdot \neg b)$ LTL-NOT-SINCE
	$a \leq \bullet a \cdot b \rightarrow a \leq \square b$ LTL-INDUCTION
	$\square a \leq \diamond(\text{start} \cdot a)$ LTL-FINITE

Fig. 9. $\text{LTL}_f(\mathcal{T})$, linear temporal logic on finite traces over an arbitrary theory

For example, we can have sets of naturals by setting $\mathcal{E} ::= n \in \mathbb{N} \mid i \in \mathcal{V}'$, where our constants $\mathcal{C} = \mathbb{N}$ and \mathcal{V}' is some set of variables distinct from those we use for sets. We can update the variables in \mathcal{V}' using IncNat 's actions while simultaneously using set actions to keep sets of naturals. Our KMT can then prove that the term $(\text{inc}_i \cdot \text{add}(x, i))^* \cdot (i > 100) \cdot \text{in}(x, 100)$ is non-empty by pushing tests back (and unrolling the loop 100 times). The set theory's sub function calls the client theory's sub function, so all $\text{in}(x, e)$ formulae must come *later* in the global well ordering than any of those generated by the client theory's $e = c$ or $e \neq c$.

3.4 Past-time linear temporal logic

Past-time linear temporal logic on finite traces (LTL_f) is a *higher-order theory*: LTL_f is itself parameterized on a theory \mathcal{T} , which introduces its own predicates and actions—any \mathcal{T} test can appear inside of LTL_f 's predicates (Fig. 9). For information on LTL_f , we refer the reader to work by Baier and McIlraith, De Giacomo and Vardi, Roşu, and Beckett et al., and Campbell and Greenberg [5, 8, 10, 11, 17, 18, 45].

LTL_f adds just two predicates: $\bigcirc a$, pronounced “last a ”, means a held in the prior state; and $a \mathcal{S} b$, pronounced “ a since b ”, means b held at some point in the past, and a has held since then. There is a slight subtlety around the beginning of time: we say that $\bigcirc a$ is false at the beginning (what can be true in a state that never happened?), and $a \mathcal{S} b$ degenerates to b at the beginning of time. The last and since predicates together are enough to encode the rest of LTL_f ; encodings are given below the syntax. Weakest preconditions uses inference rules: to push back \mathcal{S} , we unroll $a \mathcal{S} b$ into $a \cdot \bigcirc(a \mathcal{S} b) + b$; pushing last through an action is easy, but pushing back a or b recursively uses the PB^{\bullet} judgment. Adding these rules leaves our judgments monotonic, and if $\pi \cdot a \text{ PB}^{\bullet} x$, then $x = \sum a_i \pi$. In this case, our implementation's recursive modules are critical—they allow us to use the derived pushback inside our definition of weakest preconditions.

The equivalence axioms come from Temporal NetKAT [8]; the deductive completeness result for these axioms comes from Campbell and Greenberg's work, which proves deductive completeness

785	Syntax	785	Semantics
786	$\alpha ::= f = v$	786	$F =$ packet fields
787	$\pi ::= f \leftarrow v$	787	$V =$ packet field values
788	$\text{sub}(\alpha) = \{\alpha\}$	788	State = $F \rightarrow V$
789		789	$\text{pred}(f = v, t) = \text{last}(t).f = v$
790		790	$\text{act}(f \leftarrow v, \sigma) = \sigma[f \mapsto v]$
791	Weakest precondition	791	Axioms
792	$f \leftarrow v \cdot f = v$ WP 1	792	$f \leftarrow v \cdot f' = v' \equiv f' = v' \cdot f \leftarrow v$ PA-MOD-COMM
793	$f \leftarrow v \cdot f = v'$ WP 0 when $v \neq v'$	793	$f \leftarrow v \cdot f = v \equiv f \leftarrow v$ PA-MOD-FILTER
794	$f' \leftarrow v \cdot f = v$ WP $f = v$	794	$f = v \cdot f = v' \equiv 0$, if $v \neq v'$ PA-CONTRA
795		795	$\sum_v f = v \equiv 1$ PA-MATCH-ALL

Fig. 10. Tracing NetKAT a/k/a NetKAT without dup

for an axiomatic framing and then relates those axioms to our equations [10, 11]; we could have also used Roşu’s proof with coinductive axioms [45].

As a use of LTL_f , recall the simple While program from Sec. 1. We may want to check that, before the last state after the loop, the variable j was always less than or equal to 200. We can capture this with the test $\bigcirc \square(j \leq 200)$. We can use the LTL_f axioms to push tests back through actions; for example, we can rewrite terms using these LTL_f axioms alongside the natural number axioms:

$$\begin{aligned}
 j := j + 2 \cdot \bigcirc \square(j \leq 200) &\equiv j := j + 2 \cdot (j \leq 200 \cdot \bigcirc \square(j \leq 200)) \\
 &\equiv (j := j + 2 \cdot j \leq 200) \cdot \bigcirc \square(j \leq 200) \\
 &\equiv (j \leq 198) \cdot j := j + 2 \cdot \bigcirc \square(j \leq 200) \\
 &\equiv (j \leq 198) \cdot \bigcirc \square(j \leq 200) \cdot j := j + 2
 \end{aligned}$$

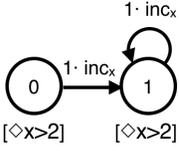
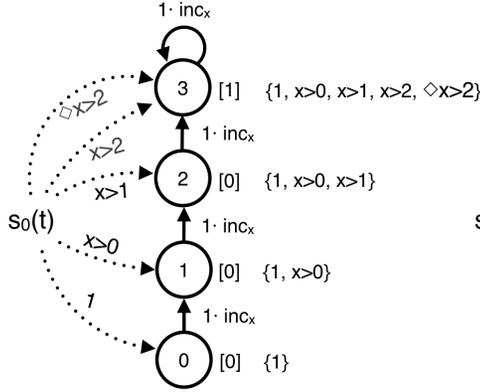
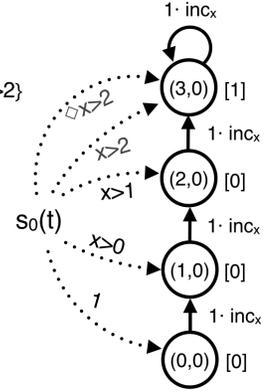
Pushing the temporal test back through the action reveals that j is never greater than 200 if before the action j was not greater than 198 in the previous state and j never exceeded 200 before the action as well. The final pushed back test $(j \leq 198) \cdot \bigcirc \square(j \leq 200)$ satisfies the theory requirements for pushback not yielding larger tests, since the resulting test is only in terms of the original test and its subterms. Note that we’ve embedded our theory of naturals into LTL_f : we can generate a complete equational theory for LTL_f over any other complete theory.

The ability to use temporal logic in KAT means that we can model check programs by phrasing model checking questions in terms of program equivalence. For example, for some program r , we can check if $r \equiv r \cdot \bigcirc \square(j \leq 200)$. In other words, if there exists some program trace that does not satisfy the test, then it will be filtered—resulting in non-equivalent terms. If the terms are equal, then every trace from r satisfies the test. Similarly, we can test whether $r \cdot \bigcirc \square(j \leq 200)$ is empty—if so, there are *no* satisfying traces.

In addition to model checking, temporal logic is a useful programming language feature: programs can make dynamic program decisions based on the past more concisely. Such a feature is useful for Temporal NetKAT (Sec. 3.6 below), but could also be used for, e.g., regular expressions with lookbehind or even a limited form of back-reference.

3.5 Tracing NetKAT

We define NetKAT as a KMT over packets, which we model as functions from packet fields to values (Fig. 10). KMT’s trace semantics diverge slightly from NetKAT’s: like KAT+B! (Sec. 3.1; [29]), NetKAT normally merges adjacent writes. If the policy analysis demands reasoning about the history of packets traversing the network—reasoning, for example, about which routes packets actually take—the programmer must insert dups to record relevant moments in time. From our

Term Automaton**Minimized****Theory Automaton****KMT Automaton**Fig. 11. Automata construction for $\text{inc}_x^* \cdot \diamond x > 2$ in the theory of $\text{LTL}_f(\text{IncNat})$.

perspective, NetKAT very nearly has a tracing semantics, but the traces are selective. If we put an implicit dup before *every* field update, NetKAT has our tracing semantics.

3.6 Temporal NetKAT

We derive Temporal NetKAT as $\text{LTL}_f(\text{NetKAT})$, i.e., LTL_f instantiated over tracing NetKAT; the combination yields precisely the system described in the Temporal NetKAT paper [8]. Our LTL_f theory can now rely on Campbell and Greenberg’s proof of deductive completeness for LTL_f [10, 11], we can automatically derive a stronger completeness result for Temporal NetKAT than that from the paper, which showed completeness only for “network-wide” policies, i.e., those with start at the front.

4 AUTOMATA

While the deductive completeness proof (Theorem 2.8 in Sec. 2) gives a way to determine equivalence of KAT terms through normalization, using such rewriting-based proofs as the basis of a decision procedure isn’t always practical. But just as pushback yields a novel completeness proof, it can also help provide an automata-theoretic account of equivalence. We compare performance in Sec. 6.

Our automata theory is heavily based on previous work on Antimirov partial derivatives [3] and NetKAT’s compiler [51]. We diverge their approach to account for client theory predicates that depend on more than the last state of the trace. Our solution is adapted from the Temporal NetKAT compiler [8]: to construct an automaton for a term in a KMT, we build *two* automata—one for the policy fragment of the term and one for each predicate that occurs therein—and combine the two in a specialized quasi-intersection operation.

A *KMT automaton* is a 4-tuple $(S, s_0, \epsilon, \delta)$, where: the set of automata states S identifies non-initial states (unrelated to *State*, the state space of the client theory); the *initial state selector* s_0 is a function that takes a trace and selects an initial state; the *acceptance function* $\epsilon : S \times \text{Trace} \rightarrow \mathcal{P}(\text{State})$ is a function identifying which theory states (in *State*) are accepted in each automaton state $s \in S$; the *transition function* $\delta : S \times \text{Trace} \rightarrow \mathcal{P}(\text{Log} \times S)$ identifies successor states given an automaton and a single KMT state. Intuitively, the automata works on traces, i.e., sequences of log entries: $\langle \sigma_0, \pi_1 \rangle \dots \langle \sigma_n, \pi_n \rangle$. While the acceptance and transition functions look at traces, that is an artifact of their construction: they will only actually look at the last state of the input.

883 Consider the KMT automaton (Fig. 11, rightmost) for the term $\text{inc}_x^* \cdot \diamond x > 2$ taken from the
 884 $\text{LTL}_f(\text{IncNat})$ theory. The automaton accepts a trace of the form: $\langle [x \mapsto 1, \perp] \rangle \langle [x \mapsto 2, \text{inc}_x] \rangle \langle [x \mapsto$
 885 $3], \text{inc}_x \rangle$. Informally, the initial state selector s_0 looks at the trace so far to determine where to begin
 886 a run. In our example, the state $(0,0)$ is used for a trace where x has never been greater than 2 and
 887 x is currently 0; we would start in state $(1,0)$ if x were 1. From state $(1,0)$, the automaton will move
 888 to state $(2,1)$ and then $(3,1)$ unconditionally for the inc_x action, which corresponds to actions in the
 889 log entries of the trace. The acceptance function, written in brackets alongside each state, assigns
 890 state $(3,1)$ the condition 1, meaning that all theory states are accepted; no other states are accepting,
 891 i.e., their acceptance condition is 0.

892 The transition function δ takes an automaton state S and a KMT trace and maps them to a set of
 893 new pairs of automaton state and and KMT log items (a KMT state/action pair). In the figure, we
 894 draw transitions as arcs between states with a pair of a KMT test and a primitive KMT action. For
 895 example, the transition from state $(1,0)$ to $(2,0)$ is captured by the term $1 \cdot \text{inc}_x$, i.e., the transition
 896 can always fire and increments the value of x .

897 Taken all together, our KMT automaton captures the fact that there are 4 interesting cases for
 898 the term $\text{inc}_x^* \cdot \diamond x > 2$. If the program trace already had $x > 2$ at some point in the past or has
 899 $x > 2$ in the current state, then we move to state $(3,0)$ and will accept the trace regardless of how
 900 many increment commands are executed in the future. If the initial trace has $x > 1$, then we move
 901 to state $(2,0)$. If we see at least one more increment command, then we move to state $(3,0)$ where the
 902 trace will be accepted no matter what. If the initial trace has $x > 0$, we move to state $(1,0)$ where
 903 we must see at least 2 more increment commands before accepting the trace. Finally, if the initial
 904 trace has any other value (here, only $x = 0$ is possible), then we move to state $(0,0)$ and must see at
 905 least 3 increment commands before accepting.

907 4.1 Constructing KMT automata

908 The KMT automaton for a given term p is constructed in two phases: we first construct a *term*
 909 *automaton* for a version of p where predicates are placed as transition and acceptance conditions.
 910 Such a symbolic automaton can be unwieldy—for example, the term automaton in (Fig. 11, top left)
 911 has a temporal predicate as an acceptance condition, which is challenging to reason about. We
 912 therefore find every predicate mentioned in the term automaton and construct a corresponding
 913 *theory automaton* (Fig. 11, middle), using pushback to move tests to the front of the automaton. We
 914 finally combine these two to form a KMT automaton with simple acceptance conditions (0 or 1).
 915

916 **4.1.1 Term automata.** The term automaton uses the Antimirov-derivative approach from the
 917 NetKAT compiler to construct an automaton for a given term. At this stage, we leave arbitrary
 918 predicates on the edges—we use theory automata (Sec. 4.1.2) to resolve those predicates. Formally,
 919 our automaton $\mathcal{A}_\pi(p)$ is defined in as a 4-tuple $(S, s_0, \epsilon, \delta)$, where S is a set of states, s_0 is an initial
 920 state, ϵ is an acceptance condition, and δ is a transition relation (Fig. 12). The automata's runs are
 921 described by the accepts relation, where we say $\mathcal{A}_\pi(p), \ell$ accepts $t; t'$ when the automaton $\mathcal{A}_\pi(p)$
 922 in state ℓ accepts the trace t' after having already seen the trace t . The semi-colon on the right-hand
 923 side of the accepts relation can be thought of as a 'cursor' indicating where we are in the trace so
 924 far. The NetKAT compiler's automaton doesn't bother keeping the trace, but our predicates can
 925 reflect on the entire trace—so we must be careful to keep track of it.

926 Given a KMT term p , we start constructing the term automaton $\mathcal{A}_\pi(p)$ by annotating each
 927 occurrence of each theory action π in p with a unique label ℓ ; these labels will form the states
 928 of $\mathcal{A}_\pi(p)$. Then we take the partial derivative of p by computing $\mathcal{D}(p)$ (Fig. 12). The derivative
 929 computes a set of *linear forms*—tuples of the form $\langle d, \pi^\ell, k \rangle$. There will be exactly one such tuple
 930 for each unique label ℓ , and each label will represent a single state in the automaton. We also
 931

932	Derivative	$\mathcal{D} : \mathcal{T}_\ell^* \rightarrow \mathcal{P}(\mathcal{T}_\ell^* \times \mathcal{T}_{\pi^\ell} \times \mathcal{T}_{\text{pred}}^*)$	Acceptance condition	$\mathcal{E} : \mathcal{T}_\ell^* \rightarrow \mathcal{T}_{\text{pred}}^*$
933	$\mathcal{D}(0)$	$= \emptyset$	$\mathcal{E}(0)$	$= 0$
934	$\mathcal{D}(1)$	$= \emptyset$	$\mathcal{E}(1)$	$= 1$
935	$\mathcal{D}(\alpha)$	$= \emptyset$	$\mathcal{E}(\alpha)$	$= \alpha$
936	$\mathcal{D}(\pi^\ell)$	$= \{\langle 1, \pi^\ell, 1 \rangle\}$	$\mathcal{E}(\pi^\ell)$	$= 0$
937	$\mathcal{D}(p + q)$	$= \mathcal{D}(p) \cup \mathcal{D}(q)$	$\mathcal{E}(p + q)$	$= \mathcal{E}(p) + \mathcal{E}(q)$
938	$\mathcal{D}(p \cdot q)$	$= \mathcal{D}(p) \odot q \cup \mathcal{E}(p) \odot \mathcal{D}(q)$	$\mathcal{E}(p \cdot q)$	$= \mathcal{E}(p) \cdot \mathcal{E}(q)$
939	$\mathcal{D}(p^*)$	$= \mathcal{D}(p) \odot p^*$	$\mathcal{E}(p^*)$	$= 1$
940				
941		$\mathcal{D}(p) \odot q = \{\langle d, \pi^\ell, k \cdot q \rangle \mid \langle d, \pi^\ell, k \rangle \in \mathcal{D}(p)\}$		$q \odot \mathcal{D}(p) = \{\langle q \cdot d, \pi^\ell, k \rangle \mid \langle d, \pi^\ell, k \rangle \in \mathcal{D}(p)\}$
942				
943	$\mathcal{A}_\pi(p)$	$= (S, s_0, \epsilon, \delta)$		Term automaton
944	S	$= \{0\} \cup \text{labels}(p)$		States
945	s_0	$= 0$		Initial state
946	$\epsilon \ell t$	$\Leftrightarrow t \in \llbracket \mathcal{E}(k_\ell) \rrbracket(t)$		Acceptance condition
947	$\delta \ell t$	$= \{(\sigma', \pi'^{\ell'}) \mid \langle d, \pi'^{\ell'}, k \rangle \in \mathcal{D}(k_\ell) \wedge t \in \llbracket d \rrbracket(t) \wedge t \langle \sigma', \pi'^{\ell'} \rangle \in \llbracket \pi'^{\ell'} \rrbracket(t)\}$		Transition relation
948	Term automaton trace acceptance		$\text{accepts} \subseteq \text{Automaton} \times S \times \text{Trace} \times (\text{State} \times \mathcal{T}_\pi)^*$	
949				
950	$\mathcal{A}_\pi(p), \ell$ accepts $t; \bullet$	$\Leftrightarrow \epsilon \ell t$		Accepting state
951	$\mathcal{A}_\pi(p), \ell$ accepts $t; \langle \sigma, \pi^{\ell'} \rangle t'$	$\Leftrightarrow (\sigma, \pi^{\ell'}) \in \delta \ell t \wedge \mathcal{A}_\pi(p), \ell'$ accepts $t \langle \sigma, \pi^{\ell'} \rangle; t'$		Taking a step

Fig. 12. KMT partial derivatives and automata

distinguish an initial state, 0. The acceptance function for state ℓ is given by $\mathcal{E}(k)$. To compute the transition relation, we compute $\mathcal{D}(k)$ for each such tuple, which yields another set of tuples. For each tuple $\langle d', \pi'^{\ell'}, k' \rangle \in \mathcal{D}(k)$, we add a transition from state π^ℓ to state $\pi'^{\ell'}$ labeled with the term $d' \cdot \pi'^{\ell'}$. The d part is a predicate identifying when the transition activates, while the k part is the “continuation”, i.e., what else in the term can be run. Since labelings are unique, we use k_ℓ to refer to the unique continuation of π^ℓ when constructing $\mathcal{A}_\pi(p)$ for a given p . We let k_0 be the continuation of the initial action, i.e., the original term p .

For example, the term $\text{inc}_x^* \cdot \diamond x > 2$, is first labeled as $(\text{inc}_x^1)^* \cdot \diamond x > 2$. We then compute $\mathcal{D}((\text{inc}_x^1)^* \cdot \diamond x > 2) = \{\langle 1, \text{inc}_x^1, (\text{inc}_x^1)^* \cdot \diamond x > 2 \rangle\}$. Hence, there is a transition from state 0 to state 1 with label $(1 \cdot \text{inc}_x)$. Taking the derivative of the resulting value, $(\text{inc}_x^1)^* \cdot \diamond x > 2$, results in the same tuple, so there is a single transition from state 1 to itself, also labeled with $1 \cdot \text{inc}_x^1$. The acceptance function for this state is given by $\mathcal{E}((\text{inc}_x^1)^* \cdot \diamond x > 2) = \diamond x > 2$. The resulting automaton, and its minimized form, are shown in Fig. 11 (left).

LEMMA 4.1 (DERIVATIVE CORRECT). *For all programs p where each primitive action π is augmented with a unique label ℓ ,*

$$(1) p \equiv \mathcal{E}(p) + \sum_{\langle d, \pi^\ell, k \rangle \in \mathcal{D}(p)} d \cdot \pi^\ell \cdot k, \text{ and}$$

(2) *For all labels ℓ in p , there exist unique d and k such that $\langle d, \pi^\ell, k \rangle \in \mathcal{D}(p)$.*

PROOF. For (1), we go by induction on p , using DENESTING in the star case. For (2), let π and ℓ be given; we go by induction on p . □

LEMMA 4.2 (TERM AUTOMATON CORRECT). *$tt' \in \llbracket k_\ell \rrbracket(t)$ iff $\mathcal{A}_\pi(p), \ell$ accepts $t; t'$.*

PROOF. By induction on the length of t' , leaving t general. □

981	$\mathcal{A}_\alpha(a)$	$=$	$(S, s_0, \epsilon, \delta)$	Theory automaton
982	S	$=$	$2^{\text{sub}(a)}$	States
983	$s_0(t)$	$=$	$\{b \in \text{sub}(a) \mid t \in \llbracket b \rrbracket(t)\}$	Initial state selector
984	$\text{serialize}(A)$	$=$	$\prod_{a \in A} a$	Serialization of predicate sets
985	$\epsilon A t$	\Leftrightarrow	$a \in A$	Acceptance condition
986	$\delta A t$	$=$	$\{(\sigma, \pi, \{c \mid \forall b \in A, \pi \cdot c \text{ PB}^\bullet b \cdot \pi\}) \mid t \langle \sigma, \pi \rangle \in \llbracket \pi \rrbracket(t)\}$	Transition relation
987	Theory automaton trace stepping			$\text{traces} \subseteq \text{Automaton}_\alpha \times S \times \text{Trace} \times (\text{State} \times \mathcal{T}_\pi)^*$
989	$\mathcal{A}_\alpha(a), A \text{ traces } t; \bullet$	\Leftrightarrow	$t \in \llbracket \text{serialize}(A) \rrbracket(t)$	Stopping
990	$\mathcal{A}_\alpha(a), A \text{ traces } t; \langle \sigma, \pi \rangle t'$	\Leftrightarrow	$(\sigma, \pi, A') \in \delta A t \wedge \mathcal{A}_\alpha(a), A' \text{ traces } t \langle \sigma, \pi \rangle; t'$	Taking a step

Fig. 13. Theory automata

The term automaton $\mathcal{A}_\pi(p)$ is equivalent to the original policy p , but we are not yet done. The term automaton makes use of arbitrary predicates in its transitions δ and its acceptance condition ϵ . For some client theories, predicates are immediately decidable, but predicates from a theory like LTL_f (Sec. 3.4) look at more than the last state of the trace. Depending on what the automata will be used for, these complex predicates may or may not be a problem. For our use here—deciding equivalence—we must simplify complex predicates: we define separate automata for tracking which predicates hold when (Sec. 4.1.2) and then construct a quasi-intersection automaton that implements predicates in the term automaton with theory automata.

4.1.2 Theory automata. Once we've constructed the term automaton, we construct theory automata for each predicate appearing anywhere in the term automaton, whether in an acceptance or a transition condition. The theory automaton for a predicate a , written $\mathcal{A}_\alpha(a)$, tracks whether a holds so far in a trace, given some initial trace and a sequence of primitive actions. Formally, $\mathcal{A}_\alpha(a)$ is a 4-tuple $(S, s_0, \epsilon, \delta)$ where S is a set of states, s_0 is an initial state selection function, ϵ is an acceptance condition, and δ is a transition relation. The states of the theory automaton are sets of subterms of the original predicate a ; when the automaton is in state $A \subseteq \text{sub}(a)$, then we expect every predicate $b \in A$ to hold. The runs of the theory automaton are characterized by the traces predicate. We say traces rather than accepts because we use the theory automaton to determine which predicates hold rather than to accept or reject a trace. (The KMT automaton will use the acceptance condition ϵ .) The initial state selector starts the theory automaton's run in the state identified by those subterms satisfied by the trace so far. The term automaton will use the theory automaton to implement its complex predicates by running each theory automaton in parallel: to determine whether to take an a transition, we consult the current state A of $\mathcal{A}_\alpha(a)$ and see whether $a \in A$, i.e., does a hold in the current state?

We use *pushback* (Sec. 2.3.2) to generate the transition relation of the theory automaton, since the pushback exactly characterizes the effect of a primitive action π on predicates a : to determine if a predicate α is true after some action a , we can instead check if b is true in the previous state when we know that $\pi \cdot a \text{ PB}^\bullet b \cdot \pi$.

While a KMT may include an infinite number of primitive actions (e.g., $x := n$ for $n \in \mathbb{N}$ in IncNat), any given term only has finitely many. For $\text{inc}_x^* \cdot \diamond x > 2$, there is only a single primitive action: inc_x . For each such action π that appears in the term and each subterm s of the test $\diamond x > 2$, we compute the pushback of π and s .

Continuing our example (Fig. 11 (middle)), there is a transition from state 2 to state 3 for action inc_x . State 3 is labeled with $\{1, x > 0, x > 1, x > 2, \diamond x > 2\}$ and state 2 is labeled with $\{1, x > 0, x > 1\}$. We compute $\text{inc}_x \cdot \diamond x > 2 \text{ WP } (\diamond x > 2 + x > 1)$. Therefore, $\diamond x > 2$ should be

1030	$\mathcal{A}_{\text{KMT}}(p) = (S, s_0, \epsilon, \delta)$	KMT automaton
1031	$S = S^{\mathcal{A}_\pi(p)} \times S^{\mathcal{A}_\alpha(a_1)} \times \dots \times S^{\mathcal{A}_\alpha(a_n)}$ where $a_i \in \mathcal{A}_\pi(p)$	States
1032	$s_0(t) = \lambda t. (s_0, s_0^{\mathcal{A}_\alpha(a_1)}(t), \dots, s_0^{\mathcal{A}_\alpha(a_n)}(t))$	Initial state selector
1033	$\epsilon s t \Leftrightarrow \epsilon^{\mathcal{A}_\alpha(a_i)} s.i t$ where $\epsilon^{\mathcal{A}_\pi(p)} s t = a_i$	Acceptance condition
1034	$\delta s t = \{(\sigma, \pi^{\ell'}, (\ell', \delta^{\mathcal{A}_\alpha(a_1)} s.1 t, \dots, \delta^{\mathcal{A}_\alpha(a_n)} s.n t)) \mid$	
1035	$\langle a_i, \pi^{\ell'}, k \rangle \in \mathcal{D}(k_\ell) \wedge \epsilon^{\mathcal{A}_\alpha(a_i)} s.i t \wedge t \langle \sigma', \pi^{\ell'} \rangle \in \llbracket \pi^{\ell'} \rrbracket(t)\}$	Transition relation
1036		
1037	KMT automaton acceptance	$\text{accepts} \subseteq \text{Automaton}_{\text{KMT}} \times S \times \text{Trace} \times (\text{State} \times \mathcal{T}_\pi)^*$
1038		
1039	$\mathcal{A}_{\text{KMT}}(p), s \text{ traces } t; \bullet \Leftrightarrow \epsilon s t$	Accepting state
1040	$\mathcal{A}_{\text{KMT}}(p), s \text{ traces } t; (\sigma, \pi)t' \Leftrightarrow (\sigma, \pi, s') \in \delta s t \wedge \mathcal{A}_{\text{KMT}}(p), s' \text{ accepts } t \langle \sigma, \pi \rangle; t'$	Taking a step
1041		
1042		
1043		
1044		
1045		
1046		
1047		
1048		
1049		
1050		
1051		
1052		
1053		
1054		
1055		
1056		
1057		
1058		
1059		
1060		
1061		
1062		
1063		
1064		
1065		
1066		
1067		
1068		
1069		
1070		
1071		
1072		
1073		
1074		
1075		
1076		
1077		
1078		

Fig. 14. Constructing KMT automata from term and theory automata

labeled in state 3 if and only if either $\diamond x > 2$ is labeled in state 2 or $x > 1$ is labeled in state 2. Since state 2 is labeled with $x > 1$, it follows that state 3 must be labeled with $\diamond x > 2$.

Finally, a state is accepting in the theory automaton if it is labeled with the top-level predicate for which the automaton was built. For example, state 3 is accepting (with acceptance function [1]), since it is labeled with $\diamond x > 2$. The acceptance condition is irrelevant for how the theory automaton itself steps—we use it in combination with the term automaton.

LEMMA 4.3 (THEORY AUTOMATON CORRECT). $t \in \llbracket \text{serialize}(A) \rrbracket(t) \iff \mathcal{A}_\alpha(a), A \text{ traces } t; t'$

PROOF. By induction on the length of t' , leaving t general. \square

4.2 KMT automata

We can combine the term and theory automata to create a KMT automaton, $\mathcal{A}_{\text{KMT}}(p)$. The idea is to run the term and theory automata in parallel, and replacing instances of theory tests in the acceptance and transition functions of the term automaton with the test on the current state in the theory automata. The states of the KMT automaton are of the form (ℓ, A_1, \dots, A_n) , where ℓ is a term automaton state and each A_i is a theory automaton state for some a occurring in the term automaton. In the product state, we refer to the underlying term automaton state with $s.0$ and each A_i as $s.i$. We use superscripts to disambiguate ϵ and δ , with the un-superscripted forms referring to the KMT automaton itself.

For example, in Fig. 11, the quasi-intersected automata (right) replaces instances of the $\diamond x > 2$ condition in state 0 of the term automaton, with the acceptance condition from the corresponding state in the theory automaton. In state (2,0) this is true, while in states (1,0) and (0,0) this is false. For transitions with the same action π , the quasi-intersection takes the conjunction of each edge's tests. Formally, we define the KMT automaton as a 4-tuple $(S, s_0, \epsilon, \delta)$, where the states are those of $\mathcal{A}_\pi(p)$ along with those of $\mathcal{A}_\alpha(a)$ for every predicate a that occurs in $\mathcal{A}_\pi(p)$. The initial state selector s_0 , acceptance condition ϵ , and transition relation δ are all defined as composites of the term and theory automata, using the appropriate theory automaton to implement the transition relation δ and acceptance condition ϵ .

The KMT automaton isn't, strictly speaking, an intersection automaton: we recapitulate the logic of the term automaton but use the theory automata where the term automaton would have consulted a complex predicate. As such, our proof follows the *logic* of Lemma 4.2, but we don't actually make use of that lemma at all.

LEMMA 4.4 (KMT AUTOMATON CORRECT).

$tt' \in \llbracket k_\ell \rrbracket(t)$ and $t \in \llbracket \text{serialize}(A_i) \rrbracket(t)$ iff $\mathcal{A}_{KMT}(p), (\ell, A_1, \dots, A_n)$ accepts $t; t'$.

PROOF. By induction on the length of t' , leaving t general and using Lemma 4.3. \square

4.3 Equivalence checking using automata

To check the equivalence of two KMT terms p and q , the implementation first converts both p and q to their respective (symbolic) automata. It then determinizes the automata to ensure that all transition predicates are disjoint (we use an algorithm based on minterms [15]). After combining the theory and term automata, we now have an automaton where the actions on transitions can be viewed as distinct characters. The implementation checks for a bisimulation between the two automata in a standard way by checking if, given any two bisimilar states, all transitions from the states lead to bisimilar states [9, 24, 43].

5 IMPLEMENTATION

We have implemented our ideas in an OCaml library; Sec. 1.3 summarizes the high-level idea and gives an example library implementation for the theory of increasing natural numbers. To use a higher-order theory such as that of product theories, one need only instantiate the appropriate modules in the library:

```

module P = Product(IncNat)(Boolean)
module A = Automata(P.K) (* automata-theoretic decision procedure *)
module D = Decide(P) (* normalization-based decision procedure *)
let a = P.K.parse "y<1; (a=F + a=T; inc(y)); y>0" in
let b = P.K.parse "y<1; a=T; inc(y)" in
assert (A.equivalent (A.of_term a) (A.of_term b));
assert (D.equivalent a b)

```

The module `P` instantiates `Product` over our theories of incrementing naturals and booleans; the module `A` gives us an automata theory for the KMT $(P.K)$ associated with `P`, and the module `D` gives a way to normalize terms based on the completeness proof. User's of the library can access these representations to perform any number of tasks such as compilation, verification, inference, and so on. For example, checking language equivalence is then simply a matter of reading in KMT terms and calling the appropriate equivalence function. Our implementation currently supports both a decision procedure based on automata and one based on the normalization term-rewriting from the completeness proof. In practice, our implementation uses several optimizations, with the two most prominent being (1) hash-consing all KAT terms to ensure fast set operations, and (2) lazy construction and exploration of automata during equivalence checking. Domain optimizations are possible, too: our satisfiability procedure for `IncNat` makes a heuristic decision between using our incomplete custom solver or `Z3` [19]—our solver is much faster on its restricted domain.

5.1 Optimizations

We've implemented smart constructors, which hash-cons and also automatically rewrite common identities (e.g., constructing $p \cdot 1$ will simply return p ; constructing $(p^*)^*$ will simply return p^*). Client theories can extend the smart constructors to witness theory-specific identities. Client theories can implement custom solvers or rely on `Z3` embeddings—custom solvers are typically faster. These optimizations are partly responsible for the speed of our normalization routine (when it avoids the costly `DENEST` case).

Benchmark	Decision Procedure	
	Automata	Normalization
test-in-loop	9.305 sec	0.001 sec
count-twice	0.012 sec	0.001 sec
loop-reorder-arith	6.166 sec	0.001 sec
loop-parity-swap	0.010 sec	TO
compute-bool-formula	2.659 sec	0.001 sec
population-count	21.451 sec	0.001 sec

Fig. 15. Implementation microbenchmarks

We haven't particularly optimized our automata implementation. Two particular opportunities for optimization stand out, both of which focus on reducing the state space of the theory automata. First, most client-theory predicates only consider the most recent state, in which case we need not generate a theory automaton at all. Second, the formal presentation of theory automata generates one automaton per predicate, the states of which are subsets of subterms of that predicate—an exponential blowup. While convenient for the proof, many predicates will share subterms—so we pay the cost of blowup more than once, tracking the same subterms in more than one theory automaton. We could instead generate a *single* theory automaton, where a state is a set drawn from subterms of all of the predicates in the term automaton, which would reduce some of the state-space blowup.

6 EVALUATION

We performed a few experiments to evaluate our tool on a collection of simple microbenchmarks. Fig. 15 shows the microbenchmarks, each of which performs a simple task. For instance, the `population-count` example initializes a collection of boolean variables and then counts how many are set to true using a natural number counter. It proves that, if the number is above a certain threshold, then all booleans must have been set to true. The figure also shows the time it takes to verify the equivalence of terms for each example using both the automata- and normalization-based decision procedures. We use a timeout of 5 minutes.

Interestingly, the normalization-based decision procedure is very fast in many cases. This is likely due to a combination of hash-consing and smart constructors that rewrite complex terms into simpler ones when possible, and the fact that, unlike previous KAT-based normalization proofs (e.g., [1, 32]) our normalization proof does not require splitting predicates into all possible “complete tests.” However, the normalization-based decision procedure does very poorly on examples where there is a sum nested inside of a Kleene star, i.e., $(p + q)^*$. The `loop-parity-swap` benchmark is one such example – it flips the parity of a boolean variables multiple times in a loop and verifies that the end value is always the same as the initial value. In this case the normalization-based decision procedure must repeatedly invoke the `DENEST` rewriting rule, which greatly increases the size of the term on each invocation.

On the other hand, the automata-based decision procedure easily handles the `loop-parity-swap`, terminating in all cases. It takes significantly longer on most examples due to the high cost of constructing and using theory automata for every theory predicate in the term.

7 RELATED WORK

Kozen and Mamouras's Kleene algebra with equations [35] is perhaps the most closely related work: they also devise a framework for proving extensions of KAT sound and complete. Both

1177 their work and ours use rewriting to find normal forms and prove deductive completeness. Their
 1178 rewriting systems work on mixed sequences of actions and predicates, but they can only delete
 1179 these sequences wholesale or replace them with a single primitive action or predicate; our rewriting
 1180 system’s pushback operation only works on predicates due to the trace semantics that preserves
 1181 the order of actions, but pushback isn’t restricted to producing at most a single primitive predicate.
 1182 Each framework can do things the other cannot. Kozen and Mamouras can accommodate equations
 1183 that combine actions, like those that eliminate redundant writes in KAT+B! and NetKAT [1, 29]; we
 1184 can accommodate more complex predicates and their interaction with actions, like those found in
 1185 Temporal NetKAT [8] or those produced by the compositional theories (Sec. 3). It may be possible
 1186 to build a hybrid framework, with ideas from both. A trace semantics occurs in previous work on
 1187 KAT as well [27, 32].

1188 Kozen studies KATs with arbitrary equations $x := e$ [33], also called Schematic KAT, where e
 1189 comes from arbitrary first-order structures over a fixed signature Σ . He has a pushback-like axiom
 1190 $x := e \cdot p \equiv \phi[x/e] \cdot x := e$. Arbitrary first-order structures over Σ ’s theory are much more expressive
 1191 than anything we can handle—the pushback may or may not decrease in size, depending on Σ ; KATs
 1192 over such theories are generally undecidable. We, on the other hand, are able to offer pay-as-you-
 1193 go results for soundness and completeness as well as an automata-theoretic implementation for
 1194 decidability—but only for first-order structures that admit a non-increasing weakest precondition.

1195 Larsen et al. [37] allow comparison of variables, but this of course leads to an incomplete theory.
 1196 They are, able, however, to decide emptiness of an entire expression.

1197 Coalgebra provides a general framework for reasoning about state-based systems [34, 46, 50],
 1198 which has proven useful in the development of automata theory for KAT extensions. Although
 1199 we do not explicitly develop the connection in this paper, KMT uses tools similar to those used
 1200 in coalgebraic approaches, and one could perhaps adapt our theory and implementation to that
 1201 setting. In principle, we ought to be able to combine ideas from the two schemes into a single, even
 1202 more general framework that supports complex actions *and* predicates.

1203 Our work is loosely related to Satisfiability Modulo Theories (SMT) [20]. The high-level motiva-
 1204 tion is the same—to create an extensible framework where custom theories can be combined [41]
 1205 and used to increase the expressiveness and power [52] of the underlying technique (SAT vs. KA).
 1206 Some of our KMT theories implement satisfiability checking by calling out to Z3 [19].

1207 The pushback requirement detailed in this paper generalizes the classical notion of weakest
 1208 precondition [6, 21, 47]. Automatic weakest precondition generation is generally limited in the
 1209 presence of loops in while-programs. While there has been much work on loop invariant infer-
 1210 ence [25, 26, 28, 31, 42, 49], the problem remains undecidable in most cases; however, the pushback
 1211 restrictions of “growth” of terms makes it possible for us to automatically lift the weakest pre-
 1212 condition generation to loops in KAT. In fact, this is exactly what the normalization proof does
 1213 when lifting tests out of the Kleene star operator. The pushback operation generalizes weakest
 1214 preconditions because the various PB relations can change the program itself.

1215 The automata representation described in Sec. 4 is based on prior work on symbolic automata [15,
 1216 43, 51]. In a departure from prior work, our automata construction must account for theories with
 1217 predicates that look arbitrarily far back into a trace. The separate theory and term automata we
 1218 use are based on ideas from Temporal NetKAT [8].

1219 8 CONCLUSION

1221 Kleene algebra modulo theories (KMT) is a new framework for extending Kleene algebra with tests
 1222 with the addition of actions and predicates in a custom domain. KMT uses an operation that pushes
 1223 tests back through actions to go from a decidable client theory to a domain-specific KMT. Derived
 1224 KMTs are sound and complete with respect to a trace semantics; we derive automata-theoretic

1225

1226 decision procedures for the KMT in an implementation that mirrors our formalism. The KMT
 1227 framework captures common use cases and can reproduce *by simple composition* several results
 1228 from the literature, some of which were challenging results in their own right, as well as several
 1229 new results: we offer theories for bitvectors [29], natural numbers, unbounded sets, networks [1],
 1230 and temporal logic [8].

1232 ACKNOWLEDGMENTS

1233 Dave Walker and Aarti Gupta provided valuable advice. Ryan Beckett was supported by NSF CNS
 1234 award 1703493.

1237 REFERENCES

- 1238 [1] Carolyn Jane Anderson, Nate Foster, Arjun Guha, Jean-Baptiste Jeannin, Dexter Kozen, Cole Schlesinger, and David
 1239 Walker. 2014. NetKAT: Semantic Foundations for Networks. In *Proceedings of the 41st ACM SIGPLAN-SIGACT Symposium
 1240 on Principles of Programming Languages (POPL '14)*. ACM, New York, NY, USA, 113–126.
- 1241 [2] Allegra Angus and Dexter Kozen. 2001. *Kleene Algebra with Tests and Program Schematology*. Technical Report. Cornell
 1242 University, Ithaca, NY, USA.
- 1243 [3] Valentin Antimirov. 1995. Partial Derivatives of Regular Expressions and Finite Automata Constructions. *Theoretical
 1244 Computer Science* 155 (1995), 291–319.
- 1245 [4] Mina Tahmasbi Arashloo, Yaron Koral, Michael Greenberg, Jennifer Rexford, and David Walker. 2016. SNAP: Stateful
 1246 Network-Wide Abstractions for Packet Processing. In *Proceedings of the 2016 ACM SIGCOMM Conference (SIGCOMM
 1247 '16)*. ACM, New York, NY, USA, 29–43.
- 1248 [5] Jorge A. Baier and Sheila A. McIlraith. 2006. Planning with First-order Temporally Extended Goals Using Heuristic
 1249 Search. In *National Conference on Artificial Intelligence (AAAI'06)*. AAAI Press, 788–795. <http://dl.acm.org/citation.cfm?id=1597538.1597664>
- 1250 [6] Mike Barnett and K. Rustan M. Leino. 2005. Weakest-precondition of Unstructured Programs. In *Proceedings of the 6th
 1251 ACM SIGPLAN-SIGSOFT Workshop on Program Analysis for Software Tools and Engineering (PASTE '05)*. ACM, New
 1252 York, NY, USA, 82–87.
- 1253 [7] Adam Barth and Dexter Kozen. 2002. *Equational verification of cache blocking in lu decomposition using kleene algebra
 1254 with tests*. Technical Report. Cornell University.
- 1255 [8] Ryan Beckett, Michael Greenberg, and David Walker. 2016. Temporal NetKAT. In *Proceedings of the 37th ACM SIGPLAN
 1256 Conference on Programming Language Design and Implementation (PLDI '16)*. ACM, New York, NY, USA, 386–401.
- 1257 [9] Filippo Bonchi and Damien Pous. 2013. Checking NFA Equivalence with Bisimulations Up to Congruence. *SIGPLAN
 1258 Not.* 48, 1 (Jan. 2013), 457–468.
- 1259 [10] Eric Hayden Campbell. 2017. *Infiniteness and Linear Temporal Logic: Soundness, Completeness, and Decidability*.
 1260 Undergraduate thesis. Pomona College.
- 1261 [11] Eric Hayden Campbell and Michael Greenberg. 2018. Injecting finiteness to prove completeness for finite linear
 1262 temporal logic. (2018). In submission.
- 1263 [12] Ernie Cohen. 1994. Hypotheses in Kleene Algebra. (1994). <http://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.56.6067>
- 1264 [13] Ernie Cohen. 1994. Lazy Caching in Kleene Algebra. (1994). <http://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.57.5074>
- 1265 [14] Ernie Cohen. 1994. *Using Kleene algebra to reason about concurrency control*. Technical Report. Telcordia.
- 1266 [15] Loris D'Antoni and Margus Veanes. 2014. Minimization of Symbolic Automata. *SIGPLAN Not.* 49, 1 (Jan. 2014), 541–553.
- 1267 [16] Anupam Das and Damien Pous. 2017. A Cut-Free Cyclic Proof System for Kleene Algebra. In *Automated Reasoning with
 1268 Analytic Tableaux and Related Methods*, Renate A. Schmidt and Cláudia Nalon (Eds.). Springer International Publishing,
 1269 Cham, 261–277.
- 1270 [17] Giuseppe De Giacomo, Riccardo De Masellis, and Marco Montali. 2014. Reasoning on LTL on Finite Traces: Insensitivity
 1271 to Infiniteness.. In *AAAI Citeseer*, 1027–1033.
- 1272 [18] Giuseppe De Giacomo and Moshe Y Vardi. 2013. Linear temporal logic and linear dynamic logic on finite traces.
 1273 In *IJCAI'13 Proceedings of the Twenty-Third international joint conference on Artificial Intelligence*. Association for
 1274 Computing Machinery, 854–860.
- 1275 [19] Leonardo De Moura and Nikolaj Bjørner. 2008. Z3: An Efficient SMT Solver. In *Proceedings of the Theory and
 1276 Practice of Software, 14th International Conference on Tools and Algorithms for the Construction and Analysis of Systems
 1277 (TACAS'08/ETAPS'08)*. Springer-Verlag, Berlin, Heidelberg, 337–340.

- 1275 [20] Leonardo De Moura and Nikolaj Bjørner. 2011. Satisfiability Modulo Theories: Introduction and Applications. *Commun.*
1276 *ACM* 54, 9 (Sept. 2011), 69–77.
- 1277 [21] Edsger W. Dijkstra. 1975. Guarded Commands, Nondeterminacy and Formal Derivation of Programs. *Commun. ACM*
1278 18, 8 (Aug. 1975), 453–457.
- 1279 [22] Nate Foster, Rob Harrison, Michael J. Freedman, Christopher Monsanto, Jennifer Rexford, Alec Story, and David Walker.
1280 2011. Frenetic: a network programming language. In *Proceeding of the 16th ACM SIGPLAN international conference on*
1281 *Functional Programming, ICFP 2011, Tokyo, Japan, September 19-21, 2011*. 279–291. <https://doi.org/10.1145/2034773.2034812>
- 1282 [23] Nate Foster, Dexter Kozen, Konstantinos Mamouras, Mark Reitblatt, and Alexandra Silva. 2016. Probabilistic NetKAT.
1283 In *Programming Languages and Systems: 25th European Symposium on Programming, ESOP 2016, Held as Part of the*
1284 *European Joint Conferences on Theory and Practice of Software, ETAPS 2016, Eindhoven, The Netherlands, April 2–8, 2016,*
1285 *Proceedings*, Peter Thiemann (Ed.). Springer Berlin Heidelberg, Berlin, Heidelberg, 282–309.
- 1286 [24] Nate Foster, Dexter Kozen, Matthew Milano, Alexandra Silva, and Laure Thompson. 2015. A Coalgebraic Decision Pro-
1287 cedure for NetKAT. In *Proceedings of the 42Nd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming*
1288 *Languages (POPL '15)*. ACM, New York, NY, USA, 343–355.
- 1289 [25] Carlo A. Furia and Bertrand Meyer. 2009. Inferring Loop Invariants using Postconditions. *CoRR* abs/0909.0884 (2009).
- 1290 [26] Carlo Alberto Furia and Bertrand Meyer. 2010. *Inferring Loop Invariants Using Postconditions*. Springer Berlin Heidelberg,
1291 Berlin, Heidelberg, 277–300.
- 1292 [27] Murdoch J. Gabbay and Vincenzo Ciancia. 2011. Freshness and Name-restriction in Sets of Traces with Names. In
1293 *Proceedings of the 14th International Conference on Foundations of Software Science and Computational Structures: Part of*
1294 *the Joint European Conferences on Theory and Practice of Software (FOSSACS'11/ETAPS'11)*. Berlin, Heidelberg, 365–380.
- 1295 [28] Juan P. Galeotti, Carlo A. Furia, Eva May, Gordon Fraser, and Andreas Zeller. 2014. Automating Full Functional
1296 Verification of Programs with Loops. *CoRR* abs/1407.5286 (2014). <http://arxiv.org/abs/1407.5286>
- 1297 [29] Niels Bjørn Bugge Grathwohl, Dexter Kozen, and Konstantinos Mamouras. 2014. KAT + B!. In *Proceedings of the*
1298 *Joint Meeting of the Twenty-Third EACSL Annual Conference on Computer Science Logic (CSL) and the Twenty-Ninth*
1299 *Annual ACM/IEEE Symposium on Logic in Computer Science (LICS) (CSL-LICS '14)*. ACM, New York, NY, USA, Article
1300 44, 44:1–44:10 pages.
- 1301 [30] Arjun Guha, Mark Reitblatt, and Nate Foster. 2013. Machine-verified network controllers. In *ACM SIGPLAN Conference*
1302 *on Programming Language Design and Implementation, PLDI '13, Seattle, WA, USA, June 16-19, 2013*. 483–494. <https://doi.org/10.1145/2462156.2462178>
- 1303 [31] Soonho Kong, Yungbum Jung, Cristina David, Bow-Yaw Wang, and Kwangkeun Yi. 2010. Automatically Inferring
1304 Quantified Loop Invariants by Algorithmic Learning from Simple Templates. In *Proceedings of the 8th Asian Conference*
1305 *on Programming Languages and Systems (APLAS'10)*. 328–343.
- 1306 [32] Dexter Kozen. 2003. *Kleene algebra with tests and the static analysis of programs*. Technical Report. Cornell University.
- 1307 [33] Dexter Kozen. 2004. Some results in dynamic model theory. *Science of Computer Programming* 51, 1 (2004), 3 – 22.
1308 <https://doi.org/10.1016/j.scico.2003.09.004> Mathematics of Program Construction (MPC 2002).
- 1309 [34] Dexter Kozen. 2017. On the Coalgebraic Theory of Kleene Algebra with Tests. In *Rohit Parikh on Logic, Language and*
1310 *Society*. Springer, 279–298.
- 1311 [35] Dexter Kozen and Konstantinos Mamouras. 2014. Kleene Algebra with Equations. In *Automata, Languages, and*
1312 *Programming: 41st International Colloquium, ICALP 2014, Copenhagen, Denmark, July 8-11, 2014, Proceedings, Part II*,
1313 Javier Esparza, Pierre Fraigniaud, Thore Husfeldt, and Elias Koutsoupias (Eds.). Springer Berlin Heidelberg, Berlin,
1314 Heidelberg, 280–292.
- 1315 [36] Dexter Kozen and Maria-Christina Patron. 2000. Certification of Compiler Optimizations Using Kleene Algebra with
1316 Tests. In *Proceedings of the First International Conference on Computational Logic (CL '00)*. Springer-Verlag, London, UK,
1317 UK, 568–582.
- 1318 [37] Kim G Larsen, Stefan Schmid, and Bingtian Xue. 2016. WNetKAT: Programming and Verifying Weighted Software-
1319 Defined Networks. In *OPODIS*.
- 1320 [38] Jedidiah McClurg, Hossein Hojjat, Nate Foster, and Pavol Černý. 2016. Event-driven Network Programming. In
1321 *Proceedings of the 37th ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI '16)*.
1322 ACM, New York, NY, USA, 369–385.
- 1323 [39] Christopher Monsanto, Nate Foster, Rob Harrison, and David Walker. 2012. A compiler and run-time system for network
1324 programming languages. In *Proceedings of the 39th ACM SIGPLAN-SIGACT Symposium on Principles of Programming*
1325 *Languages, POPL 2012, Philadelphia, Pennsylvania, USA, January 22-28, 2012*. 217–230. <https://doi.org/10.1145/2103656.2103685>
- 1326 [40] Yoshiki Nakamura. 2015. Decision Methods for Concurrent Kleene Algebra with Tests: Based on Derivative. *RAMiCS*
1327 *2015* (2015), 1.

- 1324 [41] Greg Nelson and Derek C. Oppen. 1979. Simplification by Cooperating Decision Procedures. *ACM Trans. Program.*
1325 *Lang. Syst.* 1, 2 (Oct. 1979), 245–257.
- 1326 [42] Saswat Padhi, Rahul Sharma, and Todd Millstein. 2016. Data-driven Precondition Inference with Learned Features. In
1327 *Proceedings of the 37th ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI '16)*.
1328 New York, NY, USA, 42–56.
- 1329 [43] Damien Pous. 2015. Symbolic Algorithms for Language Equivalence and Kleene Algebra with Tests. In *Proceedings of*
1330 *the 42Nd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL '15)*. New York,
1331 NY, USA, 357–368.
- 1332 [44] Mark Reitblatt, Marco Canini, Arjun Guha, and Nate Foster. 2013. FatTire: declarative fault tolerance for software-
1333 defined networks. In *Proceedings of the Second ACM SIGCOMM Workshop on Hot Topics in Software Defined Networking,*
1334 *HotSDN 2013, The Chinese University of Hong Kong, Hong Kong, China, Friday, August 16, 2013*. 109–114. <https://doi.org/10.1145/2491185.2491187>
- 1335 [45] Grigore Roşu. 2016. Finite-Trace Linear Temporal Logic: Coinductive Completeness. In *International Conference on*
1336 *Runtime Verification*. Springer, 333–350.
- 1337 [46] J. J.M.M. Rutten. 1996. *Universal Coalgebra: A Theory of Systems*. Technical Report. CWI (Centre for Mathematics and
1338 Computer Science), Amsterdam, The Netherlands, The Netherlands.
- 1339 [47] Andrew E. Santosa. 2015. Comparing Weakest Precondition and Weakest Liberal Precondition. *CoRR* abs/1512.04013
1340 (2015).
- 1341 [48] Cole Schlesinger, Michael Greenberg, and David Walker. 2014. Concurrent NetCore: From Policies to Pipelines. In
1342 *Proceedings of the 19th ACM SIGPLAN International Conference on Functional Programming (ICFP '14)*. ACM, New York,
1343 NY, USA, 11–24.
- 1344 [49] Rahul Sharma and Alex Aiken. 2014. From Invariant Checking to Invariant Inference Using Randomized Search. In
1345 *Proceedings of the 16th International Conference on Computer Aided Verification - Volume 8559*. New York, NY, USA,
1346 88–105.
- 1347 [50] Alexandra Silva. 2010. *Kleene Coalgebra*. PhD Thesis. University of Minho, Braga, Portugal.
- 1348 [51] Steffen Smolka, Spiridon Eliopoulos, Nate Foster, and Arjun Guha. 2015. A Fast Compiler for NetKAT. In *Proceedings*
1349 *of the 20th ACM SIGPLAN International Conference on Functional Programming (ICFP 2015)*. ACM, New York, NY, USA,
1350 328–341.
- 1351 [52] Aaron Stump, Clark W. Barrett, David L. Dill, and Jeremy R. Levitt. 2001. A Decision Procedure for an Extensional
1352 Theory of Arrays. In *LICS*.

1353
1354
1355
1356
1357
1358
1359
1360
1361
1362
1363
1364
1365
1366
1367
1368
1369
1370
1371
1372