

Kleene Algebra Modulo Theories

MICHAEL GREENBERG, Pomona College

RYAN BECKETT, Microsoft Research

ERIC CAMPBELL, Cornell University

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 an efficient implementation 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. Our framework is based on *pushback*, a generalization of weakest preconditions that specifies how predicates and actions interact. We offer several case studies, showing tracing variants of theories from the literature (bitvectors, NetKAT) along with novel compositional theories (products, temporal logic, and sets). We derive new results over *unbounded state*, reasoning about monotonically increasing, unbounded natural numbers. 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.

CCS Concepts: • **Software and its engineering** → **Formal language definitions**; *Frameworks*; *Formal software verification*; *Correctness*; *Automated static analysis*; • **Theory of computation** → **Regular languages**.

ACM Reference Format:

Michael Greenberg, Ryan Beckett, and Eric Campbell. 2020. Kleene Algebra Modulo Theories. *ACM Trans. Program. Lang. Syst.* 1, 1, Article 1 (January 2020), 47 pages.

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 [3, 7, 12–14, 41] 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, 22, 42, 44, 56].

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

Authors' addresses: Michael Greenberg, Pomona College, michael@cs.pomona.edu; Ryan Beckett, Microsoft Research, Ryan.Beckett@microsoft.com; Eric Campbell, Cornell University, ehc86@cornell.edu.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

© 2020 Copyright held by the owner/author(s).

0164-0925/2020/1-ART1

<https://doi.org/>

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, 23, 60], after a long series of papers that resembled, but did not use, the KAT framework [21, 31, 45, 51]. 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 [15, 39, 46, 50]. 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, 37].

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 a general framework for automatically deriving sound, complete, and decidable KATs with tracing semantics for client theories. To demonstrate the effectiveness of our approach, we not only reproduce results from the literature (e.g., tracing variants of bit vectors and NetKAT), but we also derive new KATs that go behind the existing, finite-state KATs to KATs using monotonically increasing, unbounded naturals. 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. Our completeness proof corresponds directly to an equivalence decision procedure. Our OCaml implementation allows users to compose a KAT with both decision procedures from small theory specifications. 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:

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

KATs are naively 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 critical for reasoning about 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, 30, 40]. 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 [23, 60] 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 decision procedures when given arbitrary domain-specific theories. Our framework, which we call Kleene algebras modulo theories (KMT) after the objects it produces, allows us to derive metatheory and implementation for KATs based on a given theory. The KMT framework obviates 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.

We offer some cartoons of KMTs here; see Sec. 2 for technical details.

Consider P_{set} (Fig. 1b), a program defined over both naturals and a *set* data structure with two operations: insertion and membership tests. The insertion action `insert(x, j)` inserts the value of an expression (j) into a given set (x); the membership test `in(x, c)` determines whether a constant

<pre> 148 assume i < 50 149 while (i < 100) do 150 i := i + 1 151 j := j + 2 152 end 153 assert j > 100 </pre>	<pre> assume 0 ≤ j < 4 while (i < 10) do i := i + 1 j := (j << 1) + 3 if i < 5 then insert(X, j) end assert in(X, 9) </pre>	<pre> i := 0 parity := false while (true) do odd[i] := parity i := i + 1 parity := !parity end assert odd[99] </pre>
(a) P_{nat}	(b) P_{set}	(c) P_{map}

Fig. 1. Example simple while programs.

(c) is included in a given set (x). An axiom characterizing pushback for this theory has the form:

$$\text{insert}(x, e) \cdot \text{in}(x, c) \equiv ((e = c) + \text{in}(x, c)) \cdot \text{insert}(x, e)$$

Our theory of sets works for expressions e taken from another theory, so long as the underlying theory supports tests of the form $e = c$. For example, this would work over the theory of naturals since a test like $j = 10$ can be encoded as $(j > 9) \cdot \neg(j > 10)$.

Finally, P_{map} (Fig. 1c) uses a combination of mutable *boolean* values and a *map* data structure. Just as before, we can craft custom theories for reasoning about each of these types of state. For booleans, we can add actions of the form $b := t$ and $b := \bar{t}$ and tests of the form $b = t$ and $b = \bar{t}$. The axioms are then simple equivalences like $(b := t \cdot b = \bar{t}) \equiv 0$ and $(b := t \cdot b = t) \equiv (b := t)$. To model map data structures, we add actions of the form $X[e] := e$ and tests of the form $X[c] = c$. Just as with the set theory, the map theory is parameterized over other theories, which can provide the type of keys and values—here, integers and booleans. In P_{map} , the odd map tracks whether certain natural numbers are odd or not by storing a boolean into the map’s index. A sound axiom characterizing pushback in the theory of maps has the form:

$$(X[e_1] := e_2 \cdot X[c_1] = c_2) \equiv (e_1 = c_1 \cdot e_2 = c_2 + X[c_1] = c_2) \cdot X[e_1] := e_2$$

Each of the theories we have described so far—naturals, sets, booleans, and maps—have tests that only examine the *current* state of the program. However, we need not restrict ourselves in this way. Primitive tests can make dynamic decisions or assertions based on any previous state of the program. As an example, consider the theory of past-time, finite-trace linear temporal logic (LTL_f) [16, 17]. Linear temporal logic introduces new operators such as: $\bigcirc a$ (in the last state a), $\diamond a$ (in some previous state a), and $\square a$ (in every state a); we use finite-time LTL because finite traces are a reasonable model in most domains modeling programs.

Finally, we can encode a tracing variant of NetKAT, a system that extends KAT with actions of the form $f \leftarrow v$, where some value v is assigned to one of a finite number of fields f , and tests of the form $f = v$ where field f is tested for value v . It also includes a number of axioms such as $f \leftarrow v \cdot f = v \equiv f \leftarrow v$. The NetKAT axioms can be captured in our framework with minor changes. Further extending NetKAT to Temporal NetKAT is captured trivially in our framework as an application of the LTL_f theory to NetKAT’s theory, deriving Beckett et al.’s [8] completeness result compositionally (in fact, we can strengthen it—see Sec. 2.5).

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

197 some notion of state. We call these parts the *client theory*; the client theory’s predicates and actions
 198 are *primitive*, as opposed to those built with the KAT’s composition operators. We call the resulting
 199 KAT a *Kleene algebra modulo theory* (KMT). Deriving a trace-based semantics for the KMT and
 200 proving it sound isn’t particularly hard—it amounts to “turning the crank”. Proving the KMT is
 201 complete and decidable, however, can be much harder. We take care of much of the difficulty, lifting
 202 simple operations in the client theory generically to KAT.

203 Our framework hinges on an operation relating predicates and operations called *pushback*, first
 204 used to prove relative completeness for Temporal NetKAT [8]. Pushback is a generalization of
 205 weakest preconditions: we translate programs to a normal form with all predicates at the front (i.e.,
 206 all predicates become pre-conditions). Pushback generalizes weakest preconditions because we
 207 alter the program as we go, possibly changing its commands or structure. Given a primitive action
 208 π and a primitive predicate α , the client theory must be able to compute weakest preconditions,
 209 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
 210 client theory must be able to take any of its primitive tests and “push it back” through any of its
 211 primitive actions. Given the client’s notion of weakest preconditions, we can alter programs to
 212 take an *arbitrary* term and normalize it into a form where *all* of the predicates appear only at the
 213 front of the term, a convenient representation both for our completeness proof (Sec. 3.4) and our
 214 implementation (Sec 4).

215 The client theory’s pushback should have two properties: it should be sound, (i.e., the resulting
 216 expression is equivalent to the original one); and none of the resulting predicates should be any
 217 bigger than the original predicates, by some measure (see Sec. 3). If the pushback has these two
 218 properties, we can use it to define a normal form for the KMT generated from the client theory—and
 219 we can use that normal form to prove that the resulting KMT is complete and decidable.

220 As an example, in NetKAT, for different fields f and f' , we can use the network axioms to derive
 221 the equivalence: $(f \leftarrow v \cdot f' = v') \equiv (f' = v' \cdot f \leftarrow v)$, which satisfies the pushback requirements.
 222 For Temporal NetKAT, which adds rich temporal predicates such as $\diamond \circ$ (dstPort = 4747) (the
 223 destination port was 4747 at some point before the previous state), we can use the domain axioms
 224 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
 225 requirements of equivalence and non-increasing measure (because a is a subterm of $\diamond \circ a$).

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

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

237 1.3 Example: incrementing naturals

238 We can model programs like the While program over i and j from earlier by introducing a new
 239 client theory for natural numbers (Fig. 2). First, we extend the KAT syntax with actions $x := n$ and
 240 inc_x (increment x) and a new test $x > n$ for variables x and natural number constants n . First, we
 241 define the client semantics. We fix a set of variables, \mathcal{V} , which range over natural numbers, and
 242 the program state σ maps from variables to natural numbers. Primitive actions and predicates are
 243 interpreted over the state σ by the act and pred functions (where t is a trace of states).
 244
 245

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

Fig. 2. IncNat, increasing naturals

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 (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. 4 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. 2). 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 ($x := n$).


```

295 type a = Gt of string * int    (* alpha ::= x > n *)
296 type p = Increment of string  (* pi    ::= inc x *)
297
298 module rec IncNat : THEORY with type A.t = a and type P.t = p = struct
299   (* generated KMT, for recursive use *)
300   module K = KAT (IncNat)
301   (* boilerplate necessary for recursive modules, hashconsing *)
302   module P : CollectionType with type t = p = struct ... end
303   module A : CollectionType with type t = a = struct ... end
304   (* extensible parser; pushback; subterms of predicates *)
305   let parse name es = ...
306   let push_back p a =
307     match (p,a) with
308     | (Increment _x, Gt (_, j)) when 1 > j → PSet.singleton ~cmp:K.Test.compare (K.one ())
309     | (Increment x, Gt (y, j)) when x = y →
310       PSet.singleton ~cmp:K.Test.compare (K.theory (Gt (y, j - 1)))
311     | (Assign (x,i), Gt (y,j)) when x = y → PSet.singleton ~cmp:K.Test.compare (if i > j then K.one () else
312     | _ → PSet.singleton ~cmp:K.Test.compare (K.theory a)
313   let rec subterms x =
314     match x with
315     | Gt (_, 0) → PSet.singleton ~cmp:K.Test.compare (K.theory x)
316     | Gt (v, i) → PSet.add (K.theory x) (subterms (Gt (v, i - 1)))
317   (* decision procedure for the predicate theory *)
318   let satisfiable (a: K.Test.t) = ...
319 end
320
321
322
323


```

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

The implementation obligations—syntactic extensions, `subterms` functions, WP on primitives, a satisfiability checker for the test fragment—mirror our formal development. We offer more client theories in Sec. 2 and more detail on the implementation in Sec. 4.

1.4 A case study: network routing protocols

As a final example demonstrating the kinds of theories supported by KMT, we turn our attention to modeling network routing protocols. While NetKAT uses Kleene algebra to define simple, stateless forwarding tables of networks, the most common network routing protocols are distributed algorithms that actually compute paths in a network by passing messages between devices. As an example the Border Gateway Protocol (BGP) [52], which allows users to define rich routing policy, has become the de facto internet routing protocol used to transport data between autonomous networks under the control of different entities (e.g. Verizon, Comcast). However, the combination of the distributed nature of BGP, the difficulty of writing policy per-device, and the fact that network devices can and often do fail [28] all contribute to the fact that network outages caused by BGP misconfiguration are common [2, 29, 33, 43, 48, 58, 63]. By encoding BGP policies



prefer neighbor B over A

if from D **then** block **else** accept

(a) Sample network with policy on C

A := 0

B := ∞

C := ∞

D := ∞

while (true) **do**

 B := min+(A, C, D)

 C := min+(A, B, D)

 D := min+(B, C)

end

(b) P_{SP}, default policy

A := (0, true)

B := (0, false)

C := (0, false)

D := (0, false)

while (true) **do**

 updateB

 updateC

 updateD

end

(c) P_{BGP}, local policies

Fig. 3. An example network and models of BGP routing.

in our framework, it immediately follows that we can decide properties about networks running BGP such as “will router A always be able to reach router B after at most 1 link failure”.

Fig. 3a shows an example network that is configured to run BGP. In BGP, devices exchange messages between neighbors to determine routes to a destination. In the figure, router A is connected to an end host (the line going to the left) and wants to tell other routers how to get to this destination.

In the default behavior of the BGP protocol, each router selects the shortest path among all of its neighbors and then informs each of its neighbors about this route (with the path length increased by one). In effect, the devices will compute the shortest paths through the network in a distributed fashion. We can model shortest paths routing in a KMT using the theory of natural numbers: in P_{SP} (Fig. 3b), each router maintains a distance to the destination. Since A knows about the destination, it will start with a distance of 0, while all other routers start with distance ∞. Then, iteratively, each other router updates its distance to be 1 more than the minimum of each of its peers, which is captured by the min+ operator. The behavior of min+ can be described by pushback equivalences like:

$$B := \text{min}+(A, C, D) \cdot B < 3 \equiv (A < 2 + C < 2 + D < 2) \cdot B := \text{min}+(A, C, D)$$

BGP gets interesting when users go beyond shortest path routing and also define router-local policy. In our example network, router C is configured with local policy (Fig. 3a): router C will block messages received from D and will prioritize paths received from neighbor B over those from A (using distance as a tie breaker). In order to accommodate this richer routing behavior, we must extend our model to P_{BGP} (Fig. 3c). Now, each router is associated with a variable storing a tuple of both the distance and whether or not the router has a path to the destination; we write C₁ for the “does C have a path” boolean and C₀ for the length of that path, if it exists. We can then create a separate update action for each device in the network to reflect the semantics of the device’s local policy (updateC, etc.). Further, suppose we have a boolean variable fail_{X,Y} for each link between routers X and Y indicating whether or not the link is failed. The update action for router C’s local policy can be captured with the following type of equivalence:

$$\text{updateC} \cdot C_0 < 3 \equiv (\neg \text{fail}_{A,C} \cdot (\neg B_1 + \text{fail}_{B,C}) \cdot A_1 \cdot (A_0 < 2) + \neg \text{fail}_{B,C} \cdot B_1 \cdot (B_0 < 2)) \cdot \text{updateC}$$

In order for router C to have a path length < 3 to the destination after applying the local update function, it must have either been the case that B did not have a route to the destination (or the B-C link is down) and A had a route with length < 2 and the A-C link is not down, or B had a route with length < 2 and the B-C link is not down. Similarly, we would need an axiom to capture when router C will have a path to the destination based on equivalences like: updateC · C₁ ≡

($A_1 \cdot \neg\text{fail}_{A,C} + B_1 \cdot \neg\text{fail}_{B,C}$) \cdot updateC—C has a path to the destination if any of its neighbors has a path to the destination and the corresponding link is not failed.

It is now possible to ask questions such as “if there is any single network link failure, will C ever have a path with length greater than 2?”. Assuming the network program is encoded as ρ , we can answer this question by checking language non-emptiness for

$$(\text{fail}_{A,C} \cdot \neg\text{fail}_{B,C} + \neg\text{fail}_{A,C} \cdot \text{fail}_{B,C}) \cdot \rho \cdot (C_0 > 2) \equiv 0$$

While we have in a sense come back to a per-program world— P_{BGP} requires definitions and axioms for each router’s local policy—we can reason in a *very* complex domain.

1.5 Contributions

We claim the following contributions:

- A compositional framework for defining KATs and proving their metatheory, with a novel development of the normalization procedure used in completeness (Sec. 3). Completeness yields a decision procedure based on normalization.
- Several case studies of this framework (Sec. 2), including a strengthening of Temporal NetKAT’s completeness result, theories for unbounded state (naturals, sets, maps), distributed routing protocols, and, most importantly, compositional theories that allow designers to experiment new, complex theories. Several of these theories use unbounded state (e.g., naturals, sets, and maps), going beyond what the state of the art in KAT metatheory is able to accommodate.
- An implementation of KMT (Sec. 4) mirroring our proofs; we derive a normalization-based equivalence decision procedure for client theories from just a few definitions. Our implementation is efficient enough for experimentation with small programs (Sec. 5).

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

The core technique we discuss here was first developed in Beckett et al.’s work on Temporal NetKAT [8]. Our work here is a significant extension of that work:

- We explicitly define the normalization routine using inference rules (Section 3.3); in Temporal NetKAT, normalization is implicit in its completeness proof.
- The Temporal NetKAT proof of completeness is a morass, simultaneously proving the correctness and termination of normalization. In our framework, we prove those theorems separately (Theorems 3.34 and 3.35).
- Our treatment of negation is improved; we prove a new KAT theorem (PUSHBACK-NEG).
- We present a general *framework* for proving completeness, while the Temporal NetKAT development is specialized to a particular instance—tracing NetKAT with LTL_f .
- The Temporal NetKAT proof achieves only network-wide completeness because of its limited understanding of LTL_f ; we are able to achieve completeness.

Beckett et al. handles compilation to forwarding decision diagrams [60], while our presentation doesn’t discuss compilation.

Syntax	Semantics
$\alpha ::= b = t$	$b \in \mathcal{B}$
$\pi ::= b := t \mid b := \bar{t}$	State = $\mathcal{B} \rightarrow \{t, \bar{t}\}$
$\text{sub}(\alpha) = \{\alpha\}$	$\text{pred}(b = t, t) = \text{last}(t)(b)$
	$\text{act}(b := t, \sigma) = \sigma[b \mapsto t]$
	$\text{act}(b := \bar{t}, \sigma) = \sigma[b \mapsto \bar{t}]$
Weakest precondition	Axioms
$b := t \cdot b = t \text{ WP } 1$	$(b := t) \cdot (b = t) \equiv (b := t) \quad \text{SET-TEST-TRUE-TRUE}$
$b := \bar{t} \cdot b = t \text{ WP } 0$	$(b := \bar{t}) \cdot (b = t) \equiv 0 \quad \text{SET-TEST-FALSE-TRUE}$

Fig. 4. BitVec, theory of bitvectors

2 CASE STUDIES

In this section, we define KAT client theories for bitvectors and networks, as well as higher-order theories for products of theories, sets over theories, and temporal logic over theories. To give a sense of the range and power of our framework, we offer these case studies before the formal details of the framework itself (Section 3). We start with a simple theory (bit vectors in Sec. 2.1), building up to unbounded state from naturals (Sec. 1.3) to sets and maps parameterized over a notion of value and variable (Sec. 2.3). As an example of a higher-order theory, we define LTL on finite traces (a/k/a LTL_F ; Sec. 2.5), extending the predicate language with temporal operators like $\bigcirc a$, meaning “the predicate a holds in the previous state of the trace”.

2.1 Bit vectors

The simplest KMT is bit vectors: we extend KAT with some finite number of bits, each of which can be set to true or false and tested for their current value (Fig. 4). The theory adds actions $b := t$ and $b := \bar{t}$ for boolean variables b , and tests of the form $b = t$, where b is drawn from some set of names \mathcal{B} . Since our bit vectors are embedded in a KAT, we can use KAT operators to build up encodings on top of bits: $b = \bar{t}$ desugars to $\neg(b = t)$; flip b desugars to $(b = t \cdot b := \bar{t}) + (b = \bar{t} \cdot b := t)$. We could go further and define numeric operators on collections of bits, at the cost of producing larger terms. We are not limited to just numbers, of course; once we have bits, we can encode any bounded data structure we like.

KAT+B! [30] develops a nearly identical theory, though our semantics admit different equations. We use a *trace* semantics, where we distinguish between $(b := t \cdot b := t)$ and $(b := t)$. Even though the final states are equivalent, they produce different traces because they run different actions. KAT+B!, on the other hand, doesn’t distinguish based on the trace of actions, so they find that $(b := t \cdot b := t) \equiv (b := t)$. It’s difficult to say whether one model is better than the other—we imagine that either could be appropriate, depending on the setting. For example, our trace semantics is useful for answering model-checking-like questions (Sec. 2.5).

2.2 Disjoint products

Given two client theories, we can combine them into a disjoint product theory, $\text{Prod}(\mathcal{T}_1, \mathcal{T}_2)$, where the states are products of the two sub-theory’s states and the predicates and actions from \mathcal{T}_1 can’t affect \mathcal{T}_2 and vice versa (Fig. 5). 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, $\text{Prod}(\text{BitVec}, \text{IncNat})$ allows us to program with both variables valued as either booleans or (increasing) naturals; the product theory lets us directly

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

Fig. 5. Prod($\mathcal{T}_1, \mathcal{T}_2$), products of two disjoint theories

Syntax	Semantics
$\alpha ::= x[e] = c \mid e = c \mid \alpha_e$	$c \in \mathcal{C}$
$\pi ::= x[c] := e \mid \pi_e$	$e \in \mathcal{E}$
pred($x[e] = c, t$) = last(t) ₁ (x), last(t) ₂ (e) = c	$x \in \mathcal{V}$
pred(α_e, t) = pred(α_e, t_2)	State = ($\mathcal{V} \rightarrow \mathcal{C} \rightarrow \mathcal{C}$) × ($\mathcal{E} \rightarrow \mathcal{C}$)
sub($x[e] = c$) = $\{x[e] = c\} \cup$ sub($\neg(e = c')$)	act($x[c] := e, \sigma$) = $\sigma[\sigma_1[x[c] \mapsto \sigma_2(e)]]$
sub($e = c$) = sub($e = c$)	act(π_e, σ) = $\sigma[\sigma_2 \mapsto \text{act}(\pi_e, \sigma_2)]$
sub(α_e) = sub(α_e)	
Pushback extending \mathcal{E}	
$x[c] := e \cdot \alpha_e \text{ WP } \alpha_e$	
$(y[c_1] := e_1) \cdot (x[e_2] = c_2) \text{ WP } x[e_2] = c_2$	
$(x[c_1] := e_1) \cdot (x[e_2] = c_2) \text{ WP } (e_2 = c_1 \cdot e_1 = c_2) + (\neg(e_2 = c_1) \cdot x[e_2] = c_2)$	
Axioms extending \mathcal{E}	
	$(x[c] := e \cdot \alpha_e) \equiv (\alpha_e \cdot x[c] := e)$ E-COMM
	$(y[c_1] := e_1 \cdot x[e_2] = c_2) \equiv (x[e_2] = c_2 \cdot y[c_1] := e_1)$ MAP-NEQ
$(x[c_1] := e_1 \cdot x[e_2] = c_2) \equiv ((e_2 = c_1 \cdot e_1 = c_2) + \neg(e_2 = c_1) \cdot x[e_2] = c_2) \cdot x[c_1] := e_1$	MAP-EQ

Fig. 6. Map(\mathcal{E}), unbounded maps over arbitrary expressions/constants

express the sorts of programs that Kozen's early static analysis work had to encode manually, i.e., loops over boolean and numeric state [37].

2.3 Unbounded sets

We define a KMT for unbounded sets parameterized on a theory of expressions \mathcal{E} (Fig. 7). 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: $\sigma_1 : \mathcal{V} \rightarrow \mathcal{P}(C)$ records the current sets for each set in \mathcal{V} , while $\sigma_2 : \mathcal{E} \rightarrow C$ 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
$\alpha ::= \text{in}(x, c) \mid e = c \mid \alpha_e$	$c \in \mathcal{C}$
$\pi ::= \text{add}(x, e) \mid \pi_e$	$e \in \mathcal{E}$
$\text{sub}(\text{in}(x, c)) = \{\text{in}(x, c)\} \cup \text{sub}(\neg(e = c))$	$x \in \mathcal{V}$
$\text{sub}(e = c) = \text{sub}(e = c)$	State $= (\mathcal{V} \rightarrow \mathcal{P}(\mathcal{C})) \times (\mathcal{E} \rightarrow \mathcal{C})$
$\text{sub}(\alpha_e) = \text{sub}(\alpha_e)$	$\text{pred}(\text{in}(x, c), t) = \text{last}(t)_2(c) \in \text{last}(t)_1(x)$
	$\text{pred}(\alpha_e, t) = \text{pred}(\alpha_e, t_2)$
	$\text{act}(\text{add}(x, e), \sigma) = \sigma[\sigma_1[x \mapsto \sigma_1(x) \cup \{\sigma(e)\}]]$
	$\text{act}(\pi_e, \sigma) = \sigma[\sigma_2 \mapsto \text{act}(\pi_e, \sigma_2)]$
Weakest precondition extending \mathcal{E}	Axioms extending \mathcal{E}
$\text{add}(y, e) \cdot \text{in}(x, c) \text{ WP } \text{in}(x, c)$	$\text{add}(y, e) \cdot \text{in}(x, c) \equiv \text{in}(x, c) \cdot \text{add}(y, e) \text{ ADD-COMM}$
$\text{add}(x, e) \cdot \text{in}(x, c) \text{ WP } (e = c) + \text{in}(x, c)$	$\text{add}(x, e) \cdot \text{in}(x, c) \equiv ((e = c) + \text{in}(x, c)) \cdot \text{add}(x, e) \text{ ADD-IN}$
$\text{add}(x, e) \cdot \alpha_e \text{ WP } \alpha_e$	$\text{add}(x, e) \cdot \alpha_e \equiv \alpha_e \cdot \text{add}(x, e) \text{ ADD-COMM2}$

Fig. 7. $\text{Set}(\mathcal{E})$, unbounded sets over expressions

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 $e \neq c$.

2.4 Unbounded maps

Maps aren't much different from sets; rather than having simple membership tests, we instead check to see whether a given key maps to a given constant (Fig. 6). Our writes use constant keys and expression values, while our reads use variable keys but constant values. We could have flipped this arrangement—writing to expression keys and reading from constant ones—but we cannot allow *both* reads and writes to expression keys. Doing so would allow us to compare variables, putting us in the realm of context-free languages and foreclosing on the possibility of a complete theory. We could add other operations (at the cost of even more equational rules/pushback entries), like the ability to remove keys from maps or to test whether a key is in the map or not. Just as for $\text{Set}(\mathcal{E})$, we must put all $x[e] = c$ and $x[e] \neq c$ formulae *later* in the global well ordering than any of those generated by the client theory's $e = c$ or $e \neq c$.

2.5 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. 8). 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, 16, 17, 53].

LTL_f adds just two predicates: $\bigcirc a$, pronounced “last a ”, means a held in the prior state; and $a S b$, 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 $a S b$ 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. The pred definitions mostly defer to the client theory's definition of pred (which may

Syntax	Semantics
$\alpha ::= \bigcirc a \mid a \mathcal{S} b \mid a$	State = $\text{State}_{\mathcal{T}}$
$\pi ::= \pi_{\mathcal{T}}$	$\text{pred}(\bigcirc a, \langle \sigma, l \rangle) = \dagger$
$\text{sub}(\bigcirc a) = \{\bigcirc a\} \cup \text{sub}(a)$	$\text{pred}(\bigcirc a, t \langle \sigma, l \rangle) = \text{pred}(a, t)$
$\text{sub}(a \mathcal{S} b) = \{a \mathcal{S} b\} \cup \text{sub}(a) \cup \text{sub}(b)$	$\text{pred}(a \mathcal{S} b, \langle \sigma, l \rangle) = \text{pred}(b, \langle \sigma, l \rangle)$
$\text{act}(\pi, \sigma) = \text{act}(\pi, \sigma)$	$\text{pred}(a \mathcal{S} b, t \langle \sigma, l \rangle) = \text{pred}(b, t \langle \sigma, l \rangle) \vee$ $(\text{pred}(a, t \langle \sigma, l \rangle) \wedge \text{pred}(a \mathcal{S} b, t))$
$\bullet a = \neg \bigcirc \neg a \quad a \mathcal{B} b = a \mathcal{S} b + \square a$	
$\text{start} = \neg \bigcirc 1 \quad \diamond a = 1 \mathcal{S} a \quad \square a = \neg \diamond \neg a$	
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. 8. $\text{LTL}_f(\mathcal{T})$, linear temporal logic on finite traces over an arbitrary theory

recursively reference the LTL_f pred function), unrolling \mathcal{S} as it goes (LTL-SINCE-UNROLL). 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 $x = \sum a_i \pi$ (Lemma 3.33). 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 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 [53].

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:

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

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 $(j \leq 198) \cdot \square(j \leq 200)$ 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

638	Syntax	638	Semantics
639	$\alpha ::= f = v$	639	$F =$ packet fields
640	$\pi ::= f \leftarrow v$	640	$V =$ packet field values
641	$\text{sub}(\alpha) = \{\alpha\}$	641	State = $F \rightarrow V$
642		642	$\text{pred}(f = v, t) = \text{last}(t).f = v$
643		643	$\text{act}(f \leftarrow v, \sigma) = \sigma[f \mapsto v]$
644	Weakest precondition	644	Axioms
645	$f \leftarrow v \cdot f = v \text{ WP } 1$	645	$f \leftarrow v \cdot f' = v' \equiv f' = v' \cdot f \leftarrow v$ PA-MOD-COMM
646	$f \leftarrow v \cdot f = v' \text{ WP } 0 \text{ when } v \neq v'$	646	$f \leftarrow v \cdot f = v \equiv f \leftarrow v$ PA-MOD-FILTER
647	$f' \leftarrow v \cdot f = v \text{ WP } f = v$	647	$f = v \cdot f = v' \equiv 0, \text{ if } v \neq v'$ PA-CONTRA
648		648	$\sum_v f = v \equiv 1$ PA-MATCH-ALL

Fig. 9. Tracing NetKAT a/k/a NetKAT without dup

651
652 can check if $r \equiv r \cdot \bigcirc \square (j \leq 200)$. In other words, if there exists some program trace that does not
653 satisfy the test, then it will be filtered—resulting in non-equivalent terms. If the terms are equal,
654 then every trace from r satisfies the test. Similarly, we can test whether $r \cdot \bigcirc \square (j \leq 200)$ is empty—if
655 so, there are *no* satisfying traces.

656 In addition to model checking, temporal logic is a useful programming language feature: programs
657 can make dynamic program decisions based on the past more concisely. Such a feature is useful
658 for Temporal NetKAT (Sec. 2.7 below), but could also be used for, e.g., regular expressions with
659 lookbehind or even a limited form of back-reference.

660 2.6 Tracing NetKAT

661 We define NetKAT as a KMT over packets, which we model as functions from packet fields to
662 values (Fig. 9). KMT's trace semantics diverge slightly from NetKAT's: like KAT+B! (Sec. 2.1; [30]),
663 NetKAT normally merges adjacent writes. If the policy analysis demands reasoning about the
664 history of packets traversing the network—reasoning, for example, about which routes packets
665 actually take—the programmer must insert dups to record relevant moments in time. Typically,
666 dups are automatically inserted at the topology level, i.e., before a packet enters a switch, we record
667 its state by running dup. From our perspective, NetKAT very nearly has a tracing semantics, but
668 the traces are selective. If we put an implicit dup before *every* field update, NetKAT has our tracing
669 semantics. The upshot is that our “tracing NetKAT” has a slightly different equational theory from
670 conventional NetKAT, rejecting the following NetKAT laws as unsound for trace semantics:
671
672

$$\begin{aligned}
 673 \quad & f = v \cdot f \leftarrow v \equiv f = v && \text{PA-FILTER-MOD} \\
 674 \quad & f \leftarrow v \cdot f \leftarrow v' \equiv f \leftarrow v' && \text{PA-MOD-MOD} \\
 675 \quad & f \leftarrow v \cdot f' \leftarrow v' \equiv f' \leftarrow v' \cdot f \leftarrow v && \text{PA-MOD-MOD-COMM}
 \end{aligned}$$

676
677 In principle, one can abstract our semantics' traces to find the more restricted NetKAT traces, but
678 we can't offer any formal support in our framework for abstracted reasoning. Just as for BitVec, It
679 is possible that ideas from Kozen and Mamouras could apply here [40]; see Sec. 6.

680 2.7 Temporal NetKAT

681 We derive Temporal NetKAT as $\text{LTL}_f(\text{NetKAT})$, i.e., LTL_f instantiated over tracing NetKAT; the
682 combination yields precisely the system described in the Temporal NetKAT paper [8]. Our LTL_f
683 theory can now rely on Campbell and Greenberg's proof of deductive completeness for LTL_f [10, 11],
684 we can automatically derive a stronger completeness result for Temporal NetKAT than that from
685
686

Syntax

$$\begin{aligned}
\alpha & ::= x < n \\
\pi & ::= x := \min+(\vec{x}) \\
\text{pred}(x < n, t) & = \text{last}(t)(x) < n \\
\text{sub}(x < n) & = \{x_i < m \mid m \leq n, x_i \in \mathcal{V}\} \\
\text{act}(x := \min+(\vec{x}), \sigma) & = \sigma[x \mapsto 1 + \min(\sigma(\vec{x}))]
\end{aligned}$$

Semantics

$$\begin{aligned}
n & \in \mathbb{N} \cup \{\infty\} \\
x & \in \mathcal{V} \\
\text{State} & = \mathcal{V} \rightarrow \mathbb{N}
\end{aligned}$$

Weakest precondition

axioms are identical to pushback

$$\begin{aligned}
x := \min+(\vec{x}) \cdot (x < \infty) & \text{ WP } \Sigma_i(x_i < \infty) \\
x := \min+(\vec{x}) \cdot (x < n) & \text{ WP } \Sigma_i(x_i < n - 1)
\end{aligned}$$

Fig. 10. SP, shortest paths in a graph

Syntax

$$\begin{aligned}
\alpha & ::= C_0 < n \mid C_1 \mid \text{fail}_{R_1, R_2} \\
\pi & ::= \text{updateC} \\
\text{pred}(C_1, t) & = \text{last}(t)_1(C)_1 \\
\text{pred}(C_0 < n, t) & = \text{last}(t)_1(C)_0 < n \\
\text{pred}(\text{fail}_{R_1, R_2}, t) & = \text{last}(t)_2(R_1, R_2) \\
\text{sub}(\text{fail}_{R_1, R_2}) & = \{\text{fail}_{R_1, R_2}\} \\
\text{sub}(C_1) & = \{A_1, B_1, \text{fail}_{A, C}, \text{fail}_{B, C}\} \\
\text{sub}(C_0 < n) & = \{A_0 < n - 1, B_0 < n - 1, A_1, B_1, \text{fail}_{A, C}, \text{fail}_{B, C}\} \\
\text{act}(\text{updateC}, \sigma) & = \sigma \left[C \mapsto \begin{cases} (1 + \sigma(B)_0, \text{true}) & \text{path}(B) \\ (1 + \sigma(A)_0, \text{true}) & \text{else if path}(A) \\ (\sigma(C)_0, \sigma(C)_1) & \text{otherwise} \end{cases} \right] \\
& \text{where path}(X) = \sigma(X)_1 \wedge \sigma((X, C))_1
\end{aligned}$$

Semantics

$$\begin{aligned}
R & = \text{Routers} \\
L & = R \times R \text{ Links} \\
n & \in \mathbb{N} \\
x & \in R \\
\text{State} & = R \rightarrow \mathbb{N} \times \{t, \bar{f}\} \\
& \times L \rightarrow \{t, \bar{f}\}
\end{aligned}$$

Weakest precondition

axioms are identical to pushback

$$\begin{aligned}
(\pi \cdot \text{fail}_{R_1, R_2}) & \text{ WP } \text{fail}_{R_1, R_2} \\
(\text{updateC} \cdot D_1) & \text{ WP } D_1 \\
(\text{updateC} \cdot C_1) & \text{ WP } \neg \text{fail}_{A, C} \cdot A_1 + \neg \text{fail}_{B, C} \cdot B_1 \\
(\text{updateC} \cdot D_0 < n) & \text{ WP } (D_0 < n) \\
(\text{updateC} \cdot C_0 < n) & \text{ WP } \neg \text{fail}_{A, C} \cdot (\neg B_1 + \text{fail}_{B, C}) \cdot A_1 \cdot (A_0 < n - 1) + \\
& \neg \text{fail}_{B, C} \cdot B_1 \cdot (B_0 < n - 1)
\end{aligned}$$

Fig. 11. BGP, protocol theory for router C from the network in Fig. 3

the paper, which showed completeness only for “network-wide” policies, i.e., those with start at the front.

2.8 Distributed routing protocols

The theory for naturals with the $\min+$ operator used for shortest path routing is shown in Fig. 10. The theory is similar to the IncNat theory but for some minor differences. First, the domain is now over $\mathbb{N} \cup \{\infty\}$. Second, there is a new axiom and pushback relation relating $\min+$ to a test of the form $x < n$. Third, the subterms function is now defined in terms of all other variables, which are infinite in principle but finite in any given term (e.g., the number of routers in a given network).

The theory for the BGP protocol instance with local router policy described in Fig. 3 is now shown in Fig. 11. For brevity, we only show the theory for router C in the network. The state has two parts: the first part maps each router to a pair of a natural number describing the path length to the destination for that router, and a boolean describing whether or not the router has a route

736	Predicates	$\mathcal{T}_{\text{pred}}^*$	Actions
737	a, b	$::= 0$	a <i>embedded predicates</i>
738		1 <i>multiplicative identity</i>	$p + q$ <i>parallel composition</i>
739		$\neg a$ <i>negation</i>	$p \cdot q$ <i>sequential composition</i>
740		$a + b$ <i>disjunction</i>	p^* <i>Kleene star</i>
741		$a \cdot b$ <i>conjunction</i>	π <i>primitive actions (\mathcal{T}_π)</i>
742		α <i>primitive predicates (\mathcal{T}_α)</i>	

Fig. 12. \mathcal{T}^* : generalized KAT syntax over a client theory \mathcal{T} (client parts highlighted)

to the destination; the second part maps links to a boolean representing whether the link is up or not. We require new axioms corresponding to each of the pushback operations shown. The action updateC commutes with unrelated tests, and otherwise behaves as described in Sec. 1.

750 3 THE KMT FRAMEWORK

751 The rest of our paper describes how our framework takes a client theory and generates a KAT. We
 752 emphasize that you need not understand the following mathematics to use our framework—we
 753 do it once and for all, so you don't have to. We first explain the structure of our framework for
 754 defining a KAT in terms of a client theory. While we have striven to make this section accessible to
 755 non-expert readers, those completely new to KATs may do well to skip our discussion of pushback
 756 (Sec. 3.3.2 on) and read our case studies (Sec. 2). We highlight anything the client theory must
 757 provide.
 758

759 We derive a KAT \mathcal{T}^* (Fig. 12) on top of a client theory \mathcal{T} where \mathcal{T} has two fundamental parts—
 760 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
 761 Boolean algebra over the client theory as $\mathcal{T}_{\text{pred}}^* \subseteq \mathcal{T}^*$.

762 Our framework can provide results for \mathcal{T}^* in a pay-as-you-go fashion: given a notion of state and
 763 an interpretation for the predicates and actions of \mathcal{T} , we derive a trace semantics for \mathcal{T}^* (Sec. 3.1);
 764 if \mathcal{T} has a sound equational theory with respect to our semantics, so does \mathcal{T}^* (Sec. 3.2); if \mathcal{T} has a
 765 complete equational theory with respect to our semantics, and satisfies certain weakest precondition
 766 requirements, then \mathcal{T}^* has a complete equational theory (Sec. 3.4); and finally, with just a few
 767 lines of code defining the structure of \mathcal{T} , we can provide a decision procedure for equivalence
 768 (Sec. 4) using the normalization routine from completeness (Sec. 3.4).

769 The key to our general, parameterized proof is a novel *pushback* operation that generalizes
 770 weakest preconditions (Sec. 3.3.2): given an understanding of how to push primitive predicates
 771 back to the front of a term, we can normalize terms for our completeness proof (Sec. 3.4).

772 3.1 Semantics

773 The first step in turning the client theory \mathcal{T} into a KAT is to define a semantics (Fig. 13). We can
 774 give any KAT a *trace semantics*: the meaning of a term is a trace t , which is a non-empty list of log
 775 entries l . Each *log entry* records a state σ and (in all but the initial state) a primitive action π . The
 776 client assigns meaning to predicates and actions by defining a set of states State and two functions:
 777 one to determine whether a predicate holds (pred) and another to determine an action's effects
 778 (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
 779 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
 780 semantics shows up in Kozen's application of KAT to static analysis [37].)
 781

782 A reader new to KATs should compare this definition with that of NetKAT or Temporal NetKAT [1,
 783 8]: defined recursively over the syntax, the denotation function collapses predicates and actions
 784

Trace definitions

$$\begin{aligned} \sigma &\in \text{State} \\ l &\in \text{Log} \quad ::= \langle \sigma, \perp \rangle \mid \langle \sigma, \pi \rangle \\ t &\in \text{Trace} = \text{Log}^+ \end{aligned}$$

$$\begin{aligned} \text{pred} &: \mathcal{T}_\alpha \times \text{Trace} \rightarrow \{\mathfrak{t}, \mathfrak{f}\} \\ \text{act} &: \mathcal{T}_\pi \times \text{State} \rightarrow \text{State} \end{aligned}$$

Trace semantics

$$\begin{aligned} \llbracket 0 \rrbracket(t) &= \emptyset & (f \bullet g)(t) &= \bigcup_{t' \in f(t)} g(t') \\ \llbracket 1 \rrbracket(t) &= \{t\} \\ \llbracket \alpha \rrbracket(t) &= \{t \mid \text{pred}(\alpha, t) = \mathfrak{t}\} & f^0(t) &= \{t\} & f^{i+1}(t) &= (f \bullet f^i)(t) \\ \llbracket \neg a \rrbracket(t) &= \{t \mid \llbracket a \rrbracket(t) = \emptyset\} \\ \llbracket \pi \rrbracket(t) &= \{t \langle \sigma', \pi \rangle \mid \sigma' = \text{act}(\pi, \text{last}(t))\} & \text{last}(\dots \langle \sigma, _ \rangle) &= \sigma \\ \llbracket p + q \rrbracket(t) &= \llbracket p \rrbracket(t) \cup \llbracket q \rrbracket(t) \\ \llbracket p \cdot q \rrbracket(t) &= (\llbracket p \rrbracket \bullet \llbracket q \rrbracket)(t) \\ \llbracket p^* \rrbracket(t) &= \bigcup_{0 \leq i} \llbracket p \rrbracket^i(t) \end{aligned}$$

Kleene Algebra axioms

$$\begin{aligned} p + (q + r) &\equiv (p + q) + r & \text{KA-PLUS-ASSOC} \\ p + q &\equiv q + p & \text{KA-PLUS-COMM} \\ p + 0 &\equiv p & \text{KA-PLUS-ZERO} \\ p + p &\equiv p & \text{KA-PLUS-IDEM} \\ p \cdot (q \cdot r) &\equiv (p \cdot q) \cdot r & \text{KA-SEQ-ASSOC} \\ 1 \cdot p &\equiv p & \text{KA-SEQ-ONE} \\ p \cdot 1 &\equiv p & \text{KA-ONE-SEQ} \\ p \cdot (q + r) &\equiv p \cdot q + p \cdot r & \text{KA-DIST-L} \\ (p + q) \cdot r &\equiv p \cdot r + q \cdot r & \text{KA-DIST-R} \\ 0 \cdot p &\equiv 0 & \text{KA-ZERO-SEQ} \\ p \cdot 0 &\equiv 0 & \text{KA-SEQ-ZERO} \\ 1 + p \cdot p^* &\equiv p^* & \text{KA-UNROLL-L} \\ 1 + p^* \cdot p &\equiv p^* & \text{KA-UNROLL-R} \\ q + p \cdot r \leq r &\rightarrow p^* \cdot q \leq r & \text{KA-LFP-L} \\ p + q \cdot r \leq q &\rightarrow p \cdot r^* \leq q & \text{KA-LFP-R} \end{aligned}$$

$$p \leq q \Leftrightarrow p + q \equiv q$$

Boolean Algebra axioms

$$\begin{aligned} a + (b \cdot c) &\equiv (a + b) \cdot (a + c) & \text{BA-PLUS-DIST} \\ a + 1 &\equiv 1 & \text{BA-PLUS-ONE} \\ a + \neg a &\equiv 1 & \text{BA-EXCL-MID} \\ a \cdot b &\equiv b \cdot a & \text{BA-SEQ-COMM} \\ a \cdot \neg a &\equiv 0 & \text{BA-CONTRA} \\ a \cdot a &\equiv a & \text{BA-SEQ-IDEM} \end{aligned}$$

Consequences

$$\begin{aligned} p \cdot a &\equiv b \cdot p \Leftrightarrow p \cdot \neg a \equiv \neg b \cdot p & \text{PUSHBACK-NEG} \\ p \cdot (q \cdot p)^* &\equiv (p \cdot q)^* \cdot p & \text{SLIDING} \\ (p + q)^* &\equiv p^* \cdot (q \cdot p^*)^* & \text{DENESTING} \\ p \cdot a &\equiv a \cdot q + r \rightarrow & \\ p^* \cdot a &\equiv (a + p^* \cdot r) \cdot q^* & \text{STAR-INV} \\ p \cdot a &\equiv a \cdot q + r \rightarrow & \\ p \cdot a \cdot (p \cdot a)^* &\equiv (a \cdot q + r) \cdot (q + r)^* & \text{STAR-EXPAND} \end{aligned}$$

Fig. 13. Semantics and equational theory for \mathcal{T}^*

into a single semantics, using Kleisli composition (written \bullet) to give meaning to sequence and an infinite union and exponentiation (written $-^*$) to give meaning to Kleene star. We've generalized the way that predicates and actions work, though, deferring to two functions that must be defined by the client theory: `pred` and `act`.

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 \mathfrak{t} , we return the singleton set holding our input trace; when `pred` returns \mathfrak{f} , 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 (Lemma 3.4).) It's acceptable for the `pred` function to recursively call the denotational semantics, though we have skipped the formal detail here. This way we can define composite primitive predicates as in, e.g., temporal logic (Sec. 2.7).

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

3.2 Soundness

Proving that the equational theory is sound is relatively straightforward: we only depend on the client's act and pred functions, and none of our KAT axioms (Fig. 13) even mention the client's primitives. Pushback negation is a novel KAT theorem (PUSHBACK-NEG); it generalizes the result that theorem that $b \cdot p \equiv p \cdot b \leftrightarrow b \cdot p \cdot \neg b + \neg b \cdot p \cdot b \equiv 0$ from Kozen [36].

LEMMA 3.1 (PUSHBACK NEGATION (PUSHBACK-NEG)). $p \cdot a \equiv b \cdot p$ iff $p \cdot \neg a \equiv \neg b \cdot p$.

PROOF. We show that both sides $p \cdot \neg a$ and $\neg b \cdot p$ are equivalent to $\neg b \cdot p \cdot \neg a$ by way of BA-EXCL-MID:

$$\begin{aligned}
 p \cdot \neg a &\equiv (b + \neg b) \cdot p \cdot \neg a && \text{(KA-SEQ-ONE, BA-EXCL-MID)} \\
 &\equiv b \cdot p \cdot \neg a + \neg b \cdot p \cdot \neg a && \text{(KA-DIST-L)} \\
 &\equiv p \cdot a \cdot \neg a + \neg b \cdot p \cdot \neg a && \text{(assumption)} \\
 &\equiv p \cdot 0 + \neg b \cdot p \cdot \neg a && \text{(BA-CONTRA)} \\
 &\equiv \neg b \cdot p \cdot \neg a && \text{(KA-PLUS-COMM, KA-PLUS-ZERO)} \\
 &\equiv 0 \cdot p + \neg b \cdot p \cdot \neg a && \text{(BA-CONTRA)} \\
 &\equiv \neg b \cdot b \cdot p + \neg b \cdot p \cdot \neg a && \text{(assumption)} \\
 &\equiv \neg b \cdot p \cdot a + \neg b \cdot p \cdot \neg a && \text{(KA-DIST-R)} \\
 &\equiv \neg b \cdot p \cdot (a + \neg a) && \text{(KA-ONE-SEQ, BA-EXCL-MID)} \\
 &\equiv \neg b \cdot p && \square
 \end{aligned}$$

The other direction of the proof is symmetric, with the two terms meeting at $b \cdot p \cdot a$.

Our soundness proof naturally enough requires that any equations the client theory adds need to be sound in our trace semantics. We do need to use several KAT theorems in our completeness proof (Fig. 13, Consequences), the most complex being star expansion (STAR-EXPAND), which we take from Temporal NetKAT [8]; we believe PUSHBACK-NEG is a novel theorem that holds in all KATs.

LEMMA 3.2 (KLEISLI COMPOSITION IS ASSOCIATIVE). $\llbracket p \rrbracket \bullet (\llbracket q \rrbracket \bullet \llbracket r \rrbracket) = (\llbracket p \rrbracket \bullet \llbracket q \rrbracket) \bullet \llbracket r \rrbracket$.

PROOF. By direct computation. □

LEMMA 3.3 (EXPONENTIATION COMMUTES). $\llbracket p \rrbracket^{i+1} = \llbracket p \rrbracket^i \bullet \llbracket p \rrbracket$

PROOF. By induction on i . When $i = 0$, both yield $\llbracket p \rrbracket$. In the inductive case, we compute:

$$\begin{aligned}
 \llbracket p \rrbracket^{i+2} &= \llbracket p \rrbracket \bullet \llbracket p \rrbracket^{i+1} \\
 &= \llbracket p \rrbracket \bullet (\llbracket p \rrbracket^i \bullet \llbracket p \rrbracket) \quad \text{by the IH} \\
 &= (\llbracket p \rrbracket \bullet \llbracket p \rrbracket^i) \bullet \llbracket p \rrbracket \quad \text{by Lemma 3.2} \\
 &= (\llbracket p \rrbracket^{i+1}) \bullet \llbracket p \rrbracket \quad \text{by Lemma 3.2}
 \end{aligned}$$

LEMMA 3.4 (PREDICATES PRODUCE SINGLETON OR EMPTY SETS). $\llbracket a \rrbracket(t) \subseteq \{t\}$.

PROOF. By induction on a , leaving t general. □

THEOREM 3.5 (SOUNDNESS OF \mathcal{T}^*). *If \mathcal{T} 's equational reasoning is sound ($p \equiv_{\mathcal{T}} q \Rightarrow \llbracket p \rrbracket = \llbracket q \rrbracket$) then \mathcal{T}^* 's equational reasoning is sound ($p \equiv q \Rightarrow \llbracket p \rrbracket = \llbracket q \rrbracket$).*

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

(KA-PLUS-ASSOC) We have $p + (q + r) \equiv (p + q) + r$; by associativity of union.

(KA-PLUS-COMM) We have $p + q \equiv q + p$; by commutativity of union.

883 (KA-PLUS-ZERO) We have $p + 0 \equiv p$; immediate, since $\llbracket 0 \rrbracket(t) = \emptyset$.

884 (KA-PLUS-IDEM) By idempotence of union $p + p \equiv p$.

885 (KA-SEQ-ASSOC) We have $p \cdot (q \cdot r) \equiv (p \cdot q) \cdot r$; by Lemma 3.2.

887 (KA-SEQ-ONE) We have $1 \cdot p \equiv p$; immediate, since $\llbracket 1 \rrbracket(t) = \{t\}$.

888 (KA-ONE-SEQ) We have $p \cdot 1 \equiv p$; immediate, since $\llbracket 1 \rrbracket(t) = \{t\}$.

889 (KA-DIST-L) We have $p \cdot (q + r) \equiv p \cdot q + p \cdot r$; we compute:

$$\begin{aligned}
 \llbracket p \cdot (q + r) \rrbracket(t) &= \bigcup_{t' \in \llbracket p \rrbracket(t)} \llbracket q + r \rrbracket(t') \\
 &= \bigcup_{t' \in \llbracket p \rrbracket(t)} (\llbracket q \rrbracket(t') \cup \llbracket r \rrbracket(t')) \\
 &= \bigcup_{t' \in \llbracket p \rrbracket(t)} \llbracket q \rrbracket(t') \cup \bigcup_{t' \in \llbracket p \rrbracket(t)} \llbracket r \rrbracket(t') \\
 &= \llbracket p \cdot q \rrbracket(t) \cup \llbracket p \cdot r \rrbracket(t) \\
 &= \llbracket p \cdot q + p \cdot r \rrbracket(t)
 \end{aligned}$$

897 (KA-DIST-R) As for KA-DIST-L.

898 (KA-ZERO-SEQ) We have $0 \cdot p \equiv 0$; immediate, since $\llbracket 0 \rrbracket(t) = \emptyset$.

899 (KA-SEQ-ZERO) We have $p \cdot 0 \equiv 0$; immediate, since $\llbracket 0 \rrbracket(t) = \emptyset$.

901 (KA-UNROLL-L) We have $p^* \equiv 1 + p \cdot p^*$. We compute:

$$\begin{aligned}
 \llbracket p^* \rrbracket(t) &= \bigcup_{0 \leq i} \llbracket p \rrbracket^i(t) \\
 &= \llbracket 1 \rrbracket(t) \cup \bigcup_{1 \leq i} \llbracket p \rrbracket^i(t) \\
 &= \llbracket 1 \rrbracket(t) \cup \llbracket p \rrbracket(t) \cup \bigcup_{2 \leq i} \llbracket p \rrbracket^i(t) \\
 &= \llbracket 1 \rrbracket(t) \cup (\llbracket p \rrbracket \bullet \llbracket 1 \rrbracket)(t) \cup \bigcup_{1 \leq i} (\llbracket p \rrbracket \bullet \llbracket p \rrbracket^i)(t) \\
 &= \llbracket 1 \rrbracket(t) \cup (\llbracket p \rrbracket \bullet \llbracket 1 \rrbracket)(t) \cup (\llbracket p \rrbracket \bullet \bigcup_{1 \leq i} \llbracket p \rrbracket^i)(t) \\
 &= \llbracket 1 \rrbracket(t) \cup (\llbracket p \rrbracket \bullet \bigcup_{0 \leq i} \llbracket p \rrbracket^i)(t) \\
 &= \llbracket 1 \rrbracket(t) \cup \llbracket p \cdot p^* \rrbracket(t) \\
 &= \llbracket 1 + p \cdot p^* \rrbracket(t)
 \end{aligned}$$

911 (KA-UNROLL-R) As for KA-UNROLL-L.

912 (KA-LFP-L) We have $p^* \cdot q \leq r$, i.e., $p^* \cdot q + r \equiv r$. By the IH, we know that $\llbracket q \rrbracket(t) \cup (\llbracket p \rrbracket \bullet \llbracket r \rrbracket)(t) \cup$
 913 $\llbracket r \rrbracket(t) = \llbracket r \rrbracket(t)$. We show, by induction on i , that $(\llbracket p \rrbracket^i \bullet \llbracket q \rrbracket)(t) \cup \llbracket r \rrbracket(t) = \llbracket r \rrbracket(t)$.

914 $(i = 0)$ We compute:

$$\begin{aligned}
 &(\llbracket p \rrbracket^0 \bullet \llbracket q \rrbracket)(t) \cup \llbracket r \rrbracket(t) \\
 &= (\llbracket 1 \rrbracket \bullet \llbracket q \rrbracket)(t) \cup \llbracket r \rrbracket(t) \\
 &= \llbracket q \rrbracket(t) \cup \llbracket r \rrbracket(t) \\
 &= \llbracket q \rrbracket(t) \cup (\llbracket q \rrbracket)(t) \cup (\llbracket p \rrbracket \bullet \llbracket r \rrbracket)(t) \cup \llbracket r \rrbracket(t) \quad \text{by the outer IH} \\
 &= \llbracket q \rrbracket(t) \cup (\llbracket p \rrbracket \bullet \llbracket r \rrbracket)(t) \cup \llbracket r \rrbracket(t) \\
 &= \llbracket r \rrbracket(t) \quad \text{by the outer IH again}
 \end{aligned}$$

923 $(i = i' + 1)$ We compute:

$$\begin{aligned}
 &(\llbracket p \rrbracket^{i'+1} \bullet \llbracket q \rrbracket)(t) \cup \llbracket r \rrbracket(t) \\
 &= (\llbracket p \rrbracket \bullet \llbracket p \rrbracket^{i'} \bullet \llbracket q \rrbracket)(t) \cup \llbracket r \rrbracket(t) \\
 &= (\llbracket p \rrbracket \bullet \llbracket p \rrbracket^{i'} \bullet \llbracket q \rrbracket)(t) \cup (\llbracket q \rrbracket)(t) \cup (\llbracket p \rrbracket \bullet \llbracket r \rrbracket)(t) \cup \llbracket r \rrbracket(t) \quad \text{by the outer IH} \\
 &= \bigcup_{t' \in \llbracket p \rrbracket(t)} (\bigcup_{t'' \in \llbracket p \rrbracket^{i'}(t')} \llbracket q \rrbracket(t'') \cup \llbracket r \rrbracket(t'')) \cup (\llbracket q \rrbracket)(t) \cup \llbracket r \rrbracket(t) \\
 &= (\llbracket p \rrbracket \bullet \llbracket r \rrbracket)(t) \cup (\llbracket q \rrbracket)(t) \cup \llbracket r \rrbracket(t) \quad \text{by the inner IH} \\
 &= \llbracket r \rrbracket(t) \quad \text{by the outer IH again}
 \end{aligned}$$

931

So, finally, we have:

$$\llbracket p^* \cdot q + r \rrbracket(t) = \left(\bigcup_{0 \leq i} \llbracket p \rrbracket^i \bullet \llbracket q \rrbracket \right)(t) \cup \llbracket r \rrbracket(t) = \bigcup_{0 \leq i} (\llbracket p \rrbracket^i \bullet \llbracket q \rrbracket)(t) \cup \llbracket r \rrbracket(t) = \bigcup_{0 \leq i} \llbracket r \rrbracket(t) = \llbracket r \rrbracket(t)$$

(KA-LFP-R) As for KA-LFP-L.

(BA-PLUS-DIST) We have $a + (b \cdot c) \equiv (a + b) \cdot (a + c)$. We have $\llbracket a + (b \cdot c) \rrbracket(t) = \llbracket a \rrbracket(t) \cup (\llbracket b \rrbracket \bullet \llbracket c \rrbracket)(t)$. By Lemma 3.4, we know that each of these denotations produces either $\{t\}$ or \emptyset , where \cup is disjunction and \bullet is conjunction. By distributivity of these operations.

(BA-PLUS-ONE) We have $a + 1 \equiv 1$; we have this directly by Lemma 3.4.

(BA-EXCL-MID) We have $a + \neg a \equiv 1$; we have this directly by Lemma 3.4 and the definition of negation.

(BA-SEQ-COMM) $a \cdot b \equiv b \cdot a$; we have this directly by Lemma 3.4 and unfolding the union.

(BA-CONTRA) We have $a \cdot \neg a \equiv 0$; we have this directly by Lemma 3.4 and the definition of negation.

(BA-SEQ-IDEM) $a \cdot a \equiv a$; we have this directly by Lemma 3.4 and unfolding the union. □

For the duration of Sec. 3, we assume that any equations \mathcal{T} adds are sound and, so, \mathcal{T}^* is sound by Theorem 3.5.

3.3 Normalization via pushback

In order to prove completeness (Sec. 3.4), we reduce our KAT terms to a more manageable subset of *normal forms*. Normalization happens via a generalization of weakest preconditions; we use a *pushback* operation to translate a term p into an equivalent term of the form $\sum a_i \cdot m_i$ where each m_i does not contain any tests. Once in this form, we can use the completeness result provided by the client theory to reduce the completeness of our language to an existing result for Kleene algebra.

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

- (1) a way to extract subterms from predicates, to define an ordering on predicates that serves as the termination measure on normalization (Sec. 3.3.1); and
- (2) weakest preconditions for primitives (Sec. 3.3.2).

Once we've defined our normalization procedure, we can use it prove completeness (Sec. 3.4).

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

Let the set of *restricted actions* \mathcal{T}_{RA} be the subset of \mathcal{T}^* where the only test is 1. We will use 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 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 , and z to denote elements of $\mathcal{T}_{\text{nf}}^*$; we typically write these sets not in set notation, but as sums, i.e., $x = \sum_{i=1}^k a_i \cdot m_i$ means $x = \{(a_1, m_1), (a_2, m_2), \dots, (a_k, m_k)\}$. The sum notation is convenient, but it is important that normal forms really be treated as sets—there should be no duplicated terms in the sum. We write $\sum_i a_i$ to denote the normal form $\sum_i a_i \cdot 1$. We will call a normal form *vacuous* when it is the empty set (i.e., the empty sum, which we interpret conventionally as 0) or when

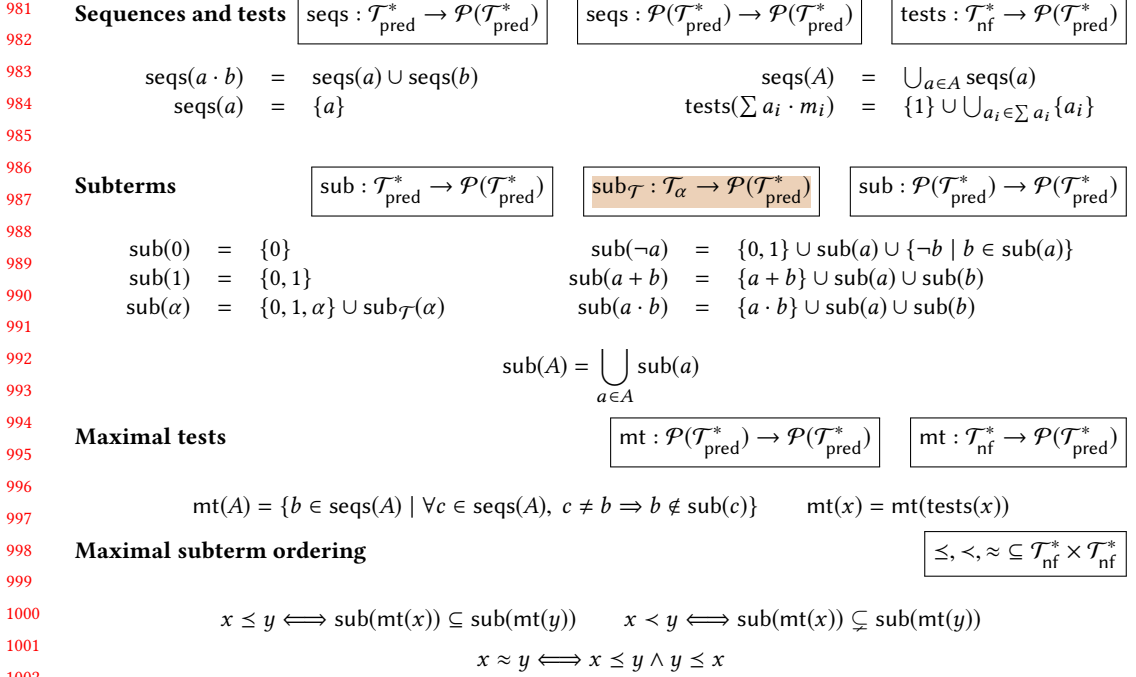


Fig. 14. Maximal tests and the maximal subterm ordering

all of its tests are 0. The set of normal forms, $\mathcal{T}_{\text{nf}}^*$, is closed over parallel composition by simply joining the sums. The fundamental challenge in our normalization method is to define sequential composition and Kleene star on $\mathcal{T}_{\text{nf}}^*$.

The definitions for the maximal subterm ordering are complex (Fig. 14), but the intuition is: seqs gets all the tests out of a predicate; tests gets all the predicates out of a normal form; sub gets subterms; mt gets “maximal” tests that cover a whole set of tests; we lift mt to work on normal forms by extracting all possible tests; the relation $x \leq y$ means that y ’s maximal tests include all of x ’s maximal tests. Maximal tests indicate which test to push back next in order to make progress towards normalization. For example, the subterms of $\diamond x > 1$ are defined by the client theory (Sec. 2.5) as $\{\diamond x > 1, x > 1, x > 0, 1, 0\}$, which represents the possible tests that might be generated pushing back $\diamond x > 1$; the maximal tests of $\diamond x > 1$ are just $\{\diamond x > 1\}$; the maximal tests of the set $\{\diamond x > 1, x > 0, y > 6\}$ are $\{\diamond x > 1, y > 6\}$ since these tests are not subterms of any other test. Therefore, we can choose to push back either of $\diamond x > 1$ or $y > 6$ next and know that we will continue making progress towards normalization.

LEMMA 3.6 (TERMS ARE SUBTERMS OF THEMSELVES). $a \in \text{sub}(a)$

PROOF. By induction on a . All cases are immediate except for $\neg a$, which uses the IH. \square

LEMMA 3.7 (0 IS A SUBTERM OF ALL TERMS). $0 \in \text{sub}(a)$

PROOF. By induction on a . The cases for 0, 1, and α are immediate; the rest of the cases follow by the IH. \square

LEMMA 3.8 (MAXIMAL TESTS ARE TESTS). $\text{mt}(A) \subseteq \text{seqs}(A)$ for all sets of tests A .

LEMMA 3.15 (SUBTERMS DECREASE IN SIZE). *If $a \in \text{sub}(b)$, then either $a \in \{0, 1, b\}$ or a comes before b in the global well ordering.*

PROOF. By induction on b . □

LEMMA 3.16 (MAXIMAL TESTS ALWAYS EXIST). *If A is a non-empty set of tests, then $\text{mt}(A) \neq \emptyset$.*

PROOF. We must show there exists at least one term in $\text{mt}(A)$.

If $\text{seqs}(A) = \{a\}$, then a is a maximal test. If $\text{seqs}(A) = \{0, 1\}$, then 1 is a maximal test. If $\text{seqs}(A) = \{0, 1, \alpha\}$, then α is a maximal test. If $\text{seqs}(A)$ isn't any of those, then let $a \in \text{seqs}(A)$ be the term that comes last in the well ordering on predicates.

To see why $a \in \text{mt}(A)$, suppose (for a contradiction) we have $b \in \text{mt}(A)$ such $b \neq a$ and $a \in \text{sub}(b)$. By Lemma 3.15, either $a \in \{0, 1, b\}$ or a comes before b in the global well ordering. We've ruled out the first two cases above. If $a = b$, then we're fine— a is a maximal test. But if a comes before b in the well ordering, we've reached a contradiction, since we selected a as the term which comes *latest* in the well ordering. □

As a corollary, note that a maximal test exists even for vacuous normal forms, where $\text{mt}(x) = \{0\}$ when x is vacuous.

LEMMA 3.17 (MAXIMAL TESTS GENERATE SUBTERMS). $\text{sub}(\text{mt}(A)) = \bigcup_{a \in \text{seqs}(A)} \text{sub}(a)$

PROOF. Since $\text{mt}(A) \subseteq \text{seqs}(A)$ (Lemma 3.8), we can restate our goal as:

$$\text{sub}(\text{mt}(A)) = \bigcup_{a \in \text{mt}(A)} \text{sub}(a) \cup \bigcup_{a \in \text{seqs}(A) \setminus \text{mt}(A)} \text{sub}(a)$$

We have $\text{sub}(\text{mt}(A)) = \bigcup_{a \in \text{mt}(A)} \text{sub}(a)$ by definition; it remains to see that the latter union is subsumed by the former; but we have $\text{seqs}(A) \subseteq \text{sub}(\text{mt}(A))$ by Lemma 3.9. □

LEMMA 3.18 (UNION DISTRIBUTES OVER MAXIMAL TESTS).

$\text{sub}(\text{mt}(A \cup B)) = \text{sub}(\text{mt}(A)) \cup \text{sub}(\text{mt}(B))$

PROOF. We compute:

$$\begin{aligned} \text{sub}(\text{mt}(A \cup B)) &= \bigcup_{a \in \text{seqs}(A \cup B)} \text{sub}(a) && \text{(Lemma 3.17)} \\ &= \bigcup_{a \in \text{seqs}(A) \cup \text{seqs}(B)} \text{sub}(a) \\ &= \left[\bigcup_{a \in \text{seqs}(A)} a \right] \cup \left[\bigcup_{b \in \text{seqs}(B)} \text{sub}(b) \right] \\ &= \text{sub}(\text{mt}(A)) \cup \text{sub}(\text{mt}(B)) && \text{(Lemma 3.17)} \end{aligned}$$

LEMMA 3.19 (MAXIMAL TESTS ARE MONOTONIC). *If $A \subseteq B$ then $\text{sub}(\text{mt}(A)) \subseteq \text{sub}(\text{mt}(B))$.*

PROOF. We have $\text{sub}(\text{mt}(B)) = \text{sub}(\text{mt}(A \cup B)) = \text{sub}(\text{mt}(A)) \cup \text{sub}(\text{mt}(B))$ (by Lemma 3.18). □

COROLLARY 3.20 (SEQUENCES OF MAXIMAL TESTS). $\text{sub}(\text{mt}(a \cdot b)) = \text{sub}(\text{mt}(a)) \cup \text{sub}(\text{mt}(b))$

PROOF.

$$\begin{aligned} &\text{sub}(\text{mt}(c \cdot d)) \\ &= \text{sub}(\text{mt}(\text{seqs}(c \cdot d))) && \text{(Corollary 3.13)} \\ &= \text{sub}(\text{mt}(\text{seqs}(c) \cup \text{seqs}(d))) \\ &= \text{sub}(\text{mt}(\text{seqs}(c))) \cup \text{sub}(\text{mt}(\text{seqs}(d))) && \text{(distributivity; Lemma 3.18)} \\ &= \text{sub}(\text{mt}(c)) \cup \text{sub}(\text{mt}(d)) && \text{(Corollary 3.13)} \end{aligned}$$

		$\text{nnf} : \mathcal{T}_{\text{pred}}^* \rightarrow \mathcal{T}_{\text{pred}}^*$
1128		
1129		
1130	$\text{nnf}(0) = 0$	$\text{nnf}(\neg 0) = 1$
1131	$\text{nnf}(1) = 1$	$\text{nnf}(\neg 1) = 0$
1132	$\text{nnf}(\alpha) = \alpha$	$\text{nnf}(\neg \alpha) = \neg \alpha$
1133	$\text{nnf}(a + b) = \text{nnf}(a) + \text{nnf}(b)$	$\text{nnf}(\neg \neg a) = \text{nnf}(a)$
1134	$\text{nnf}(a \cdot b) = \text{nnf}(a) \cdot \text{nnf}(b)$	$\text{nnf}(\neg(a + b)) = \text{nnf}(\neg a) \cdot \text{nnf}(\neg b)$
1135		$\text{nnf}(\neg(a \cdot b)) = \text{nnf}(\neg a) + \text{nnf}(\neg b)$
1136		

Fig. 15. Negation normal form

To handle negation, we translate predicates into a *negation normal form* where only primitive predicates α can be negated (Figure 15). The translation nnf uses De Morgan's laws to push negations inwards. These possibly negated predicates are commonly called “atoms”. In our setting, it is important that negation normal form is monotonic in the maximal subterm ordering (\leq).

LEMMA 3.21 (NEGATION NORMAL FORM IS MONOTONIC). *If $a \leq b$ then $\text{nnf}(\neg a) \leq \neg b$.*

PROOF. By induction on a .

($a = 0$) We have $\text{nnf}(\neg 0) = 1$ and $1 \leq \neg b$ by definition.

($a = 1$) We have $\text{nnf}(\neg 1) = 0$ and $0 \leq \neg b$ by definition.

($a = \alpha$) We have $\text{nnf}(\neg \alpha) = \neg \alpha$; since $a \leq b$, it must be that $\alpha \in \text{sub}(\text{mt}(b))$, so $\neg \alpha \in \text{sub}(\text{mt}(\neg b))$. We have $\alpha \in \text{sub}(\neg b)$, since $\alpha \in \text{sub}(b)$.

($a = \neg c$) We have $\text{nnf}(\neg \neg c) = \text{nnf}(c)$; since $c \in \text{sub}(a)$ and $a \leq b$, it must be that $c \in \text{sub}(\text{mt}(b))$, so $\text{nnf}(c) \leq \neg b$ by the IH.

($a = c + d$) We have $\text{nnf}(\neg(c + d)) = \text{nnf}(\neg c) \cdot \text{nnf}(\neg d)$; since c and d are subterms of a and $a \leq b$, $\neg c$ and $\neg d$ must be in $\text{sub}(\text{mt}(\neg b))$, and we are done by the IHs.

($a = c \cdot d$) We have $\text{nnf}(\neg(c \cdot d)) = \text{nnf}(\neg c) + \text{nnf}(\neg d)$; since c and d are subterms of a and $a \leq b$, $\neg c$ and $\neg d$ must be in $\text{sub}(\text{mt}(\neg b))$, and we are done by the IHs.

□

LEMMA 3.22 (NORMAL FORM ORDERING). *For all tests a, b, c and normal forms x, y, z , the following inequalities hold:*

(1) $a \leq a \cdot b$ (extension);

(2) if $a \in \text{tests}(x)$, then $a \leq x$ (subsumption);

(3) $x \approx \sum_{a \in \text{tests}(x)} a$ (equivalence);

(4) if $x \leq x'$ and $y \leq y'$, then $x + y \leq x' + y'$ (normal-form parallel congruence);

(5) if $x + y \leq z$, then $x \leq z$ and $y \leq z$ (inversion);

(6) if $a \leq a'$ and $b \leq b'$, then $a \cdot b \leq a' \cdot b'$ (test sequence congruence);

(7) if $a \leq x$ and $b \leq x$ then $a \cdot b \leq x$ (test bounding);

(8) if $a \leq b$ and $x \leq c$ then $a \cdot x \leq b \cdot c$ (mixed sequence congruence);

(9) if $a \leq b$ then $\text{nnf}(\neg a) \leq \neg b$ (negation normal-form monotonic).

Each of the above equalities also hold replacing \leq with $<$, excluding the equivalence (3).

PROOF. We prove each properly independently and in turn. Each property can be proved using the foregoing lemmas and set-theoretic reasoning. □

LEMMA 3.23 (TEST SEQUENCE SPLIT). *If $a \in \text{mt}(c)$ then $c \equiv a \cdot b$ for some $b < c$.*

PROOF. We have $a \in \text{seqs}(c)$ by definition. Suppose $\text{seqs}(c) = \{a, c_1, \dots, c_k\}$. By sequence extraction, we have $c \equiv a \cdot c_1 \cdot \dots \cdot c_k$ (Lemma 3.12). So let $b = c_1 \cdot \dots \cdot c_k$; we must show $b < c$, i.e., $\text{sub}(\text{mt}(b)) \subsetneq \text{sub}(\text{mt}(c))$. Note that $\{c_1, \dots, c_k\} = \text{seqs}(b)$. We find:

$$\begin{array}{ccc}
 \text{sub}(\text{mt}(b)) & \subsetneq & \text{sub}(\text{mt}(c)) \\
 & \Downarrow & \\
 \text{sub}(\text{mt}(\text{seqs}(b))) & \subsetneq & \text{sub}(\text{mt}(\text{seqs}(c))) \\
 & \Downarrow & \\
 \text{sub}(\text{mt}(\{c_1, \dots, c_k\})) & \subsetneq & \text{sub}(\text{mt}(\{a, c_1, \dots, c_k\})) \\
 & \Downarrow & \text{(distributivity; Lemma 3.18)} \\
 \bigcup_{i=1}^k \text{sub}(\text{mt}(\{c_i\})) & \subsetneq & \text{sub}(\text{mt}(a)) \cup \bigcup_{i=1}^k \text{sub}(\text{mt}(\{c_i\}))
 \end{array}
 \tag{Corollary 3.13}$$

Since $a \in \text{mt}(c)$, we know that $a \notin \text{sub}(\text{mt}(c_i))$ for all i . But terms are subterms of themselves (Lemma 3.6), so $a \in \text{sub}(a) = \text{sub}(\text{mt}(a))$. \square

LEMMA 3.24 (MAXIMAL TEST INEQUALITY). *If $a \in \text{mt}(y)$ and $x \leq y$ then either $a \in \text{mt}(x)$ or $x < y$.*

PROOF. Since $a \in \text{mt}(y)$, we have $a \in \text{sub}(\text{mt}(y))$. Since $x \leq y$, we know that $\text{sub}(\text{mt}(x)) \subseteq \text{sub}(\text{mt}(y))$. We go by cases on whether or not $a \in \text{mt}(x)$:

($a \in \text{mt}(x)$) We are done immediately.

($a \notin \text{mt}(x)$) In this case, we show that $a \notin \text{sub}(\text{mt}(x))$ and therefore $x < y$. Suppose, for a contradiction, that $a \in \text{sub}(\text{mt}(x))$. Since $a \notin \text{mt}(x)$, there must exist some $b \in \text{sub}(\text{mt}(x))$ where $a \in \text{sub}(b)$. But since $x \leq y$, we must also have $b \in \text{sub}(\text{mt}(y))$... and so it couldn't be that case that $a \in \text{mt}(y)$. We can conclude that it must, then, be the case that $a \notin \text{sub}(\text{mt}(x))$ and so $x < y$. \square

We can take a normal form x and *split* it around a maximal test $a \in \text{mt}(x)$ such that we have a pair of normal forms: $a \cdot y + z$, where both y and z are smaller than x in our ordering, because a (1) appears at the front of y and (2) doesn't appear in z at all.

LEMMA 3.25 (SPLITTING). *If $a \in \text{mt}(x)$, then there exist y and z such that $x \equiv a \cdot y + z$ and $y < x$ and $z < x$.*

PROOF. Suppose $x = \sum_{i=1}^k c_i \cdot m_i$. We have $a \in \text{mt}(x)$, so, in particular:

$$a \in \text{seqs}(\text{tests}(x)) = \text{seqs}(\text{tests}(\sum_{i=1}^k c_i \cdot m_i)) = \text{seqs}(\{c_1, \dots, c_k\}) = \bigcup_{i=1}^k \text{seqs}(c_i).$$

That is, $a \in \text{seqs}(c_i)$ for at least one i . We can, without loss of generality, rearrange x into two sums, where the first j elements have a in them but the rest don't, i.e., $x \equiv \sum_{i=1}^j c_i \cdot m_i + \sum_{i=j+1}^k c_i \cdot m_i$ where $a \in \text{seqs}(c_i)$ for $1 \leq i \leq j$ but $a \notin \text{seqs}(c_i)$ for $j+1 \leq i \leq k$. By subsumption (Lemma 3.22), we have $c_i \leq x$. Since $a \in \text{mt}(x)$, it must be that $a \in \text{mt}(c_i)$ for $1 \leq i \leq j$ (instantiating Lemma 3.24 with the normal form $c_i \cdot 1$). By test sequence splitting (Lemma 3.23), we find that $c_i \equiv a \cdot b_i$ with $b_i < c_i \leq x$ for $1 \leq i \leq j$, as well.

We are finally ready to produce y and z : they are the first j tests with a removed and the remaining tests which never had a , respectively. Formally, let $y = \sum_{i=1}^j b_i \cdot m_i$; we immediately have that $a \cdot y \equiv \sum_{i=1}^j c_i \cdot m_i$; let $z = \sum_{i=j+1}^k c_i \cdot m_i$. We can conclude that $x \equiv a \cdot y + z$.

It remains to be seen that $y < x$ and $z < x$. The argument is the same for both; presenting it for y , we have $a \notin \text{seqs}(y)$ (because of sequence splitting), so $a \notin \text{sub}(\text{mt}(y))$. But we assumed

1226 $a \in \text{mt}(x)$, so $a \in \text{sub}(\text{mt}(x))$, and therefore $y < x$. The argument for z is nearly identical but needs
 1227 no recourse to sequence splitting—we never had any $a \in \text{seqs}(c_i)$ for $j + 1 \leq i \leq k$. \square

1228 Splitting is the key lemma for making progress pushing tests back, allowing us to take a normal
 1229 form and slowly push its maximal tests to the front; its proof follows from a chain of lemmas given
 1230 in the supplementary material.
 1231

1232 3.3.2 *Pushback*. In order to define normalization—necessary for completeness (Sec. 3.4)—the client
 1233 theory must have a *weakest preconditions* operation that respects the subterm ordering.
 1234

1235 *Definition 3.26 (Weakest preconditions)*. The *weakest precondition* operation of the client theory is
 1236 a relation $\text{WP} \subseteq \mathcal{T}_\pi \times \mathcal{T}_\alpha \times \mathcal{P}(\mathcal{T}_{\text{pred}}^*)$, where \mathcal{T}_π are the primitive actions and \mathcal{T}_α are the primitive
 1237 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
 1238 to yield $\sum a_i \cdot \pi$ ”; the second π is redundant but written for clarity. We require that if $\pi \cdot \alpha \text{ WP}$
 1239 $\{a_1, \dots, a_k\} \cdot \pi$, then $\pi \cdot \alpha \equiv \sum_{i=1}^k a_i \cdot \pi$, and $a_i \leq \alpha$.
 1240

1241 Given the client theory’s weakest-precondition relation WP , we define a normalization procedure
 1242 for \mathcal{T}^* by extending the client’s WP relation to a more general *pushback* relation, PB (Fig. 16). The
 1243 client’s WP relation need not be a function, nor do the a_i need to be obviously related to α or π in
 1244 any way. Even when the WP relation is a function, the PB relation will generally not be a function.
 1245 While WP computes the classical weakest precondition, the PB relations do something different:
 1246 when pushing back we have the freedom to *change the program itself*—not normally an option for
 1247 weakest preconditions (see Sec. 6).
 1248

1249 We define the top-level normalization routine with the p norm x relation (Fig. 16), a syntax
 1250 directed relation that takes a term p and produces a normal form $x = \sum_i a_i \cdot m_i$. Most syntactic forms
 1251 are easy to normalize: predicates are already normal forms (PRED); primitive actions π are normal
 1252 forms where there’s just one summand and the predicate is 1 (ACT); and parallel composition of
 1253 two normal forms means just joining the sums (PAR). But sequence and Kleene star are harder: we
 1254 define judgments using PB to lift these operations to normal forms (SEQ , STAR).
 1255

1256 For sequences, we can recursively take $p \cdot q$ and normalize p into $x = \sum a_i \cdot m_i$ and q into
 1257 $y = \sum b_j \cdot n_j$. But how can we combine x and y into a new normal form? We can concatenate
 1258 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
 1259 to 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
 1260 $x \cdot y \text{ PB}^J z$ to mean that the concatenation of x and y is equivalent to the normal form z —the \cdot is
 1261 suggestive notation, as are other operators that appear on the left-hand side of the PB judgments.
 1262

1263 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
 1264 need to somehow move all of the tests out of the star and to the front. We do so with the PB^*
 1265 relation (Fig. 16), writing $x^* \text{ PB}^* y$ to mean that the Kleene star of x is equivalent to the normal
 1266 form y —the $*$ on the left is again suggestive notation. The PB^* relation is more subtle than PB^J .
 1267 There are four possible ways to treat x , based on how it splits (Lemma 3.25): if $x = 0$, then our work
 1268 is trivial since $0^* \equiv 1$ (STARZERO); if x splits into $a \cdot x'$ where a is a maximal test and there are no
 1269 other summands, then we can either use the KAT sliding lemma (Lemma 3.29) to pull the test out
 1270 when a is strictly the largest test in x (SLIDE) or by using the KAT expansion lemma (Lemma 3.32)
 1271 otherwise (EXPAND); if x splits into $a \cdot x' + z$, we use the KAT denesting lemma (Lemma 3.30) to
 1272 pull a out before recurring on what remains (DENEST).
 1273

1274 The bulk of the pushback’s work happens in the PB^\bullet relation, which pushes a test back through
 a restricted action; PB^R and PB^T use PB^\bullet to push tests back through normal forms and normal
 forms back through restricted actions, respectively. We write $m \cdot a \text{ PB}^\bullet y$ to mean that pushing
 the test a back through restricted action m yields the equivalent normal form y . The PB^\bullet relation

1275 **Normalization**

$$\begin{array}{c}
1276 \\
1277 \\
1278 \\
1279 \\
1280 \\
1281
\end{array}
\frac{\frac{}{a \text{ norm } a} \quad \text{PRED} \quad \frac{}{\pi \text{ norm } 1 \cdot \pi} \quad \text{ACT} \quad \frac{p \text{ norm } x \quad q \text{ norm } y}{p + q \text{ norm } x + y} \text{PAR}}{\frac{p \text{ norm } x \quad q \text{ norm } y \quad x \cdot y \text{ PB}^{\downarrow} z}{p \cdot q \text{ norm } z} \quad \text{SEQ} \quad \frac{p \text{ norm } x \quad x^* \text{ PB}^* y}{p^* \text{ norm } y} \quad \text{STAR}}$$

1282 **Sequential composition of normal forms**

$$\begin{array}{c}
1283 \\
1284 \\
1285 \\
1286
\end{array}
\frac{m_i \cdot b_j \text{ PB}^* x_{ij}}{(\sum_i a_i \cdot m_i) \cdot (\sum_j b_j \cdot n_j) \text{ PB}^{\downarrow} \sum_i \sum_j a_i \cdot x_{ij} \cdot n_j} \quad \text{JOIN}$$

1287 **Normalization of star**

$$\begin{array}{c}
1288 \\
1289 \\
1290 \\
1291 \\
1292 \\
1293 \\
1294
\end{array}
\frac{\frac{}{0^* \text{ PB}^* 1} \quad \text{STARZERO} \quad \frac{x < a \quad x \cdot a \text{ PB}^{\top} y \quad y^* \text{ PB}^* y' \quad y' \cdot x \text{ PB}^{\downarrow} z}{(a \cdot x)^* \text{ PB}^* 1 + a \cdot z} \quad \text{SLIDE}}{\frac{x \not< a \quad x \cdot a \text{ PB}^{\top} a \cdot t + u \quad (t + u)^* \text{ PB}^* y \quad y \cdot x \text{ PB}^{\downarrow} z}{(a \cdot x)^* \text{ PB}^* 1 + a \cdot z} \quad \text{EXPAND} \quad \frac{a \notin \text{mt}(z) \quad y \neq 0 \quad y^* \text{ PB}^* y' \quad x \cdot y' \text{ PB}^{\downarrow} x' \quad (a \cdot x')^* \text{ PB}^* z \quad y' \cdot z \text{ PB}^{\downarrow} z'}{(a \cdot x + y)^* \text{ PB}^* z'} \quad \text{DENEST}}$$

1295 **Pushback**

$$\begin{array}{c}
1296 \\
1297 \\
1298 \\
1299 \\
1300 \\
1301 \\
1302 \\
1303 \\
1304 \\
1305 \\
1306 \\
1307 \\
1308 \\
1309 \\
1310 \\
1311 \\
1312 \\
1313
\end{array}
\frac{\frac{m \cdot a \text{ PB}^* y}{m \cdot 0 \text{ PB}^* 0} \quad \text{SEQZERO} \quad \frac{m \cdot x \text{ PB}^{\text{R}} y}{m \cdot 1 \text{ PB}^* 1 \cdot m} \quad \text{SEQONE} \quad \frac{m \cdot a \text{ PB}^* y \quad y \cdot b \text{ PB}^{\top} z}{m \cdot (a \cdot b) \text{ PB}^* z} \quad \text{SEQSEQTEST} \quad \frac{n \cdot a \text{ PB}^* x \quad m \cdot x \text{ PB}^{\text{R}} y}{(m \cdot n) \cdot a \text{ PB}^* y} \quad \text{SEQSEQACTION} \quad \frac{m \cdot a \text{ PB}^* x \quad m \cdot b \text{ PB}^* y}{m \cdot (a + b) \text{ PB}^* x + y} \quad \text{SEQPARTEST} \quad \frac{m \cdot a \text{ PB}^* x \quad n \cdot a \text{ PB}^* y}{(m + n) \cdot a \text{ PB}^* x + y} \quad \text{SEQPARACTION} \quad \frac{\pi \cdot \alpha \text{ WP } \{a_1, \dots\}}{\pi \cdot \alpha \text{ PB}^* \sum_i a_i \cdot \pi} \quad \text{PRIM} \quad \frac{\pi \cdot a \text{ PB}^* \sum_i a_i \cdot \pi \quad \text{nfn}(-(\sum_i a_i)) = b}{\pi \cdot \neg a \text{ PB}^* b \cdot \pi} \quad \text{PRIMNEG} \quad \frac{m \cdot a \text{ PB}^* x \quad x < a \quad m^* \cdot x \text{ PB}^{\text{R}} y}{m^* \cdot a \text{ PB}^* a + y} \quad \text{SEQSTARSMALLER} \quad \frac{m \cdot a \text{ PB}^* a \cdot t + u \quad m^* \cdot u \text{ PB}^{\text{R}} x \quad t^* \text{ PB}^* y \quad x \cdot y \text{ PB}^{\downarrow} z}{m^* \cdot a \text{ PB}^* a \cdot y + z} \quad \text{SEQSTARINV} \quad \frac{m \cdot a_i \text{ PB}^* x_i}{m \cdot \sum_i a_i \cdot n_i \text{ PB}^{\text{R}} \sum_i x_i \cdot n_i} \quad \text{RESTRICTED} \quad \frac{m_i \cdot a \text{ PB}^* \sum_j b_{ij} \cdot m_{ij}}{(\sum_i a_i \cdot m_i) \cdot a \text{ PB}^{\top} \sum_i \sum_j a_i \cdot b_{ij} \cdot m_{ij}} \quad \text{TEST}$$

Fig. 16. Normalization for \mathcal{T}^*

works by analyzing both the action and the test. The client theory's WP relation is used in PB^* when we try to push a primitive predicate α through a primitive action π (PRIM); all other KAT predicates can be handled by rules matching on the action or predicate structure, deferring to other PB relations. To handle negation, the function nfn translates predicates to *negation normal form*, where negations only appear on primitive predicates (Fig. 15); PUSHBACK-NEG justifies the PRIMNEG case (PUSHBACK-NEG); we use nfn to respect the maximal subterm ordering.

1324 *Definition 3.27 (Negation normal form).* The negation normal form of a term p is a term p' such
 1325 that $p \equiv p'$ and negations occur only on primitive predicates in p' .

1326 LEMMA 3.28 (TERMS ARE EQUIVALENT TO THEIR NEGATION-NORMAL FORMS). $\text{nnf}(p) \equiv p$ and
 1327 $\text{nnf}(p)$ is in negation normal form.
 1328

1329 PROOF. By induction on the size of p . The only interesting case is when $p = \neg a$; we go by cases
 1330 on a .

1331 ($a = 0$) We have $\neg 0 \equiv 1$ immediately, and the latter is clearly negation free.

1332 ($a = 1$) We have $\neg 1 \equiv 0$; as above.

1333 ($a = \alpha$) We have $\neg \alpha$, which is in normal form.

1334 ($a = b + c$) We have $\neg(b + c) \equiv \neg b \cdot \neg c$ as a consequence of BA-EXCL-MID and soundness
 1335 (Theorem 3.5). By the IH on $\neg b$ and $\neg c$, we find that $\text{nnf}(\neg b) \equiv \neg b$ and $\text{nnf}(\neg c) \equiv$
 1336 $\neg c$ —where the left-hand sides are negation normal. So transitively, we have $\neg(b +$
 1337 $c) \equiv \text{nnf}(\neg b) \cdot \text{nnf}(\neg c)$, and the latter is negation normal.
 1338

1339 ($a = b \cdot c$) We have $\neg(b \cdot c) \equiv \neg b + \neg c$ as a consequence of BA-EXCL-MID and soundness
 1340 (Theorem 3.5). By the IH on $\neg b$ and $\neg c$, we find that $\text{nnf}(\neg b) \equiv \neg b$ and $\text{nnf}(\neg c) \equiv$
 1341 $\neg c$ —where the left-hand sides are negation normal. So transitively, we have $\neg(b \cdot c) \equiv$
 1342 $\text{nnf}(\neg b) + \text{nnf}(\neg c)$, and the latter is negation normal. \square
 1343

1344 To elucidate the way PB^\bullet handles structure, suppose we have the term $(\pi_1 + \pi_2) \cdot (\alpha_1 + \alpha_2)$. One of
 1345 two rules could apply: we could split up the tests and push them through individually (SEQPARTEST),
 1346 or we could split up the actions and push the tests through together (SEQPARACTION). It doesn't
 1347 particularly matter which we do first: the next step will almost certainly be the other rule, and in
 1348 any case the results will be equivalent from the perspective of our equational theory. It *could* be
 1349 the case that choosing a one rule over another could give us a smaller term, which might yield a
 1350 more efficient normalization procedure. Similarly, a given normal form may have more than one
 1351 maximal test—and therefore be splittable in more than one way (Lemma 3.25)—and it may be that
 1352 different splits produce more or less efficient terms. We haven't yet studied differing strategies for
 1353 pushback.
 1354

1355 LEMMA 3.29 (SLIDING). $p \cdot (q \cdot p)^* \equiv (p \cdot q)^* \cdot p$.

1356 PROOF. Following Kozen [35], as a corollary of a related result: if $p \cdot x \equiv x \cdot q$ then $p^* \cdot x \equiv x \cdot q^*$.
 1357 We prove this separate property by mutual inclusion.

1358 (\Rightarrow) We use KA-LFP-L with $p = p$ and $q = x$ and $r = x \cdot q^*$. We must show that $x + p \cdot x \cdot q^* \leq x \cdot q^*$
 1359 to find $p^* \cdot x \leq x \cdot q^*$.

1360 If $p \cdot q \leq x \cdot q$ then $p \cdot x \cdot q^* \leq x \cdot q \cdot q^*$ by monotonicity. We have $x + x \cdot q \cdot q^* \leq x \cdot q^*$ by
 1361 KA-UNROLL-L and KA-PLUS-IDEM. Therefore $x + p \cdot x \cdot q^* \leq x + x \cdot q \cdot q^* \leq x \cdot q^*$, as desired.

1362 (\Leftarrow) This case is symmetric to the first, using -R rules instead of -L rules. We apply KA-LFP-R
 1363 with $p = x$ and $r = q$ and $q = p^* \cdot x$. We must show $x + p^* \cdot x \cdot q \leq p^* \cdot x$ to find $x \cdot q^* \leq p^* \cdot x$.

1364 If $x \cdot q \leq p \cdot x$, then $p^* \cdot x \cdot q \leq p^* \cdot p \cdot x$ by monotonicity. We have $x + p^* \cdot p \cdot x \leq p^* \cdot x$ by
 1365 KA-UNROLL-R and KA-PLUS-IDEM. Therefore $x + p^* \cdot x \cdot q \leq x + p^* \cdot p \cdot x \leq p^* \cdot x$, as desired.

1366 We can now find sliding by letting $p = p \cdot q$ and $x = p$ and $q = q \cdot p$ in the above, i.e., we have the
 1367 premise $p \cdot q \cdot p \equiv p \cdot q \cdot p$ by reflexivity, and so $(p \cdot q)^* \cdot p \equiv p \cdot (q \cdot p)^*$. \square
 1368

1369 LEMMA 3.30 (DENESTING). $(p + q)^* \equiv p^* \cdot (q \cdot p^*)^*$.

PROOF. Following Kozen [35], we do the proof by mutual inclusion. The proof is surprisingly challenging, so we include it here.

(\Rightarrow) To show $(p + q)^* \leq a^* \cdot (b \cdot a^*)^*$, we apply induction with $q = 1$ and $r = p^* \cdot (q \cdot p^*)^*$ (to show $(a + b)^* \cdot 1 \leq r$). We must show that $1 + (p + q) \cdot p^* \cdot (q \cdot p^*)^* \leq p^* \cdot (q \cdot p^*)^*$. We do so in several parts, working our way there in five steps.

First, we observe that $1 \leq p^* \cdot (q \cdot p^*)^*$ (A) because:

$$\begin{aligned}
 & 1 + p^* \cdot (q \cdot p^*)^* \\
 \equiv & 1 + (1 + p \cdot p^* \cdot (q \cdot p^*)^*) && \text{KA-UNROLL-L} \\
 \equiv & 1 + p \cdot p^* \cdot (q \cdot p^*)^* && \text{KA-PLUS-ASSOC, KA-PLUS-IDEM} \\
 \equiv & p^* \cdot (q \cdot p^*)^* && \text{KA-UNROLL-L}
 \end{aligned}$$

Next, $p \cdot p^* \cdot (q \cdot p^*)^* \leq p^* \cdot (q \cdot p^*)^*$ (B) because:

$$\begin{aligned}
 & p \cdot p^* \cdot (q \cdot p^*)^* + p^* \cdot (q \cdot p^*)^* \\
 \equiv & p \cdot p^* \cdot (q \cdot p^*)^* + 1 + p \cdot p^* \cdot (q \cdot p^*)^* && \text{KA-UNROLL-L} \\
 \equiv & 1 + p \cdot p^* \cdot (q \cdot p^*)^* && \text{KA-PLUS-IDEM} \\
 \equiv & p^* \cdot (q \cdot p^*)^* && \text{KA-UNROLL-L}
 \end{aligned}$$

We have $q \cdot p^* \cdot (q \cdot p^*)^* \leq (q \cdot p^*)^*$ because:

$$\begin{aligned}
 & q \cdot p^* \cdot (q \cdot p^*)^* + (q \cdot p^*)^* \\
 \equiv & q \cdot p^* \cdot (q \cdot p^*)^* + 1 + q \cdot p^* \cdot (q \cdot p^*)^* && \text{KA-UNROLL-L} \\
 \equiv & 1 + q \cdot p^* \cdot (q \cdot p^*)^* && \text{KA-PLUS-IDEM} \\
 \equiv & (q \cdot p^*)^* && \text{KA-UNROLL-L}
 \end{aligned}$$

Further, $(q \cdot p^*)^* \leq p^* \cdot (q \cdot p^*)^*$ because:

$$\begin{aligned}
 & (q \cdot p^*)^* + p^* \cdot (q \cdot p^*)^* \\
 \equiv & (q \cdot p^*)^* + 1 \cdot (q \cdot p^*)^* + p \cdot p^* \cdot (q \cdot p^*)^* && \text{KA-UNROLL-L, KA-DIST-R} \\
 \equiv & 1 \cdot (q \cdot p^*)^* + p \cdot p^* \cdot (q \cdot p^*)^* && \text{KA-PLUS-IDEM} \\
 \equiv & p^* \cdot (q \cdot p^*)^* && \text{KA-UNROLL-L}
 \end{aligned}$$

Finally, $q \cdot p^* \cdot (q \cdot p^*)^* \leq a^* \cdot (q \cdot p^*)^*$ (C) by transitivity with the last two results.

Now we can find that

$$1 + (p + q)p^* \cdot (q \cdot p^*)^* \leq 1 + p \cdot p^* \cdot (q \cdot p^*)^* + q \cdot p^* \cdot (q \cdot p^*)^* \leq p^* \cdot (q \cdot p^*)^*$$

because:

$$\begin{aligned}
 & 1 + (p + q)p^* \cdot (q \cdot p^*)^* + 1 + p \cdot p^* \cdot (q \cdot p^*)^* + q \cdot p^* \cdot (q \cdot p^*)^* \\
 \equiv & 1 + p \cdot p^* \cdot (q \cdot p^*)^* + q \cdot p^* \cdot (q \cdot p^*)^* + p \cdot p^* \cdot (q \cdot p^*)^* + q \cdot p^* \cdot (q \cdot p^*)^* && \text{KA-PLUS-IDEM} \\
 \equiv & 1 + p \cdot p^* \cdot (q \cdot p^*)^* + q \cdot p^* \cdot (q \cdot p^*)^* && \text{KA-PLUS-IDEM}
 \end{aligned}$$

because, finally:

$$\begin{aligned}
 & 1 + p \cdot p^* \cdot (q \cdot p^*)^* + q \cdot p^* \cdot (q \cdot p^*)^* + p^* \cdot (q \cdot p^*)^* \\
 \equiv & p^* \cdot (q \cdot p^*)^* + p \cdot p^* \cdot (q \cdot p^*)^* + q \cdot p^* \cdot (q \cdot p^*)^* && \text{(A)} \\
 \equiv & p^* \cdot (q \cdot p^*)^* + q \cdot p^* \cdot (q \cdot p^*)^* && \text{(B)} \\
 \equiv & p^* \cdot (q \cdot p^*)^* && \text{(C)}
 \end{aligned}$$

(\Leftarrow) To show $p^* \cdot (q \cdot p^*)^* \leq (p+q)^*((p+q)(p+q)^*)^*$, we have first that $p \leq p+q$ and $q \leq p+q$, and so $p+q \leq (p+q)^*$. And so, by monotonicity $p^* \cdot (q \cdot p^*)^* \leq (p+q)^*((p+q)(p+q)^*)^*$. We can then find that $(p+q)^* \cdot ((p+q) \cdot (p+q)^*)^* \leq (p+q)^* \cdot ((p+q)^*)^*$ because:

$$\begin{aligned}
& (p+q)(p+q)^* + (p+q)^* \\
\equiv & p \cdot (p+q)^* + q \cdot (p+q)^* + (p+q)^* && \text{KA-DIST-R} \\
\equiv & p \cdot (p+q)^* + q \cdot (p+q)^* + 1 + (p+q)(p+q)^* && \text{KA-UNROLL-L} \\
\equiv & p \cdot (p+q)^* + q \cdot (p+q)^* + 1 + p \cdot (p+q)^* + q \cdot (p+q)^* && \text{KA-DIST-R} \\
\equiv & 1 + p \cdot (p+q)^* + q \cdot (p+q)^* && \text{KA-PLUS-IDEM} \\
\equiv & 1 + (p+q) \cdot (p+q)^* && \text{KA-DIST-R} \\
\equiv & (p+q)^* && \text{KA-UNROLL-L}
\end{aligned}$$

But we also have $(p+q)^* \cdot ((p+q)^*)^* \leq (p+q)^*$ because:

$$\begin{aligned}
& (p+q)^* \cdot ((p+q)^*)^* + (p+q)^* \\
\equiv & (p+q)^* \cdot (p+q)^* + (p+q)^* && \text{because } (x^*)^* = x^* \\
\equiv & (p+q)^* + (p+q)^* && \text{because } x^* x^* = x^* \\
\equiv & (p+q)^* && \text{KA-PLUS-IDEM}
\end{aligned}$$

□

LEMMA 3.31 (STAR INVARIANT). *If $p \cdot a \equiv a \cdot q + r$ then $p^* \cdot a \equiv (a + p^* \cdot r) \cdot q^*$.*

PROOF. We show two implications using \leq to derive the equality.

(\Rightarrow) We want to show $p^*; a \leq (a + p^*; y); x^*$.

We know that $q + pr \leq r \implies p^*q \leq r$ by the induction axiom KA-LFP-L, so we can instantiate it with p as p and q as a and r as $(a + p^*; y); x^*$. We find:

$$\begin{aligned}
& a + p; (a + p^*; y); x^* && \leq (a + p^*; y); x^* \\
& a + p; a; x^* + p; p^*; y; x^* && \leq (a + p^*; y); x^* \\
& a + p; a; x^* + p; p^*; y; x^* + (a + p^*; y); x^* && = (a + p^*; y); x^* \\
& a + p; a; x^* + p; p^*; y; x^* + a; x^* + p^*; y; x^* && = (a + p^*; y); x^* \\
& (a + a; x^* + p; a; x^*) + (p; p^*; y; x^* + p^*; y; x^*) && = (a + p^*; y); x^* \\
& (a; x^* + p; a; x^*) + (p; p^*; y; x^* + p^*; y; x^*) && = (a + p^*; y); x^* \\
& (1 + p); a; x^* + (1 + p); p^*; y; x^* && = (a + p^*; y); x^* \\
& a; x^* + p^*; y; x^* && = (a + p^*; y); x^* \\
& (a + p^*; y); x^* && = (a + p^*; y); x^*
\end{aligned}$$

(\Leftarrow) We can to show $(a + p^*; y); x^* \leq p^*; a$ We can apply the other induction axiom (KA-LFP-R), $q + r; p \leq r \implies q; p^* \leq r$, with $p = x$ and $q = (a + p^*; y)$ and $r = p^*; a$. We find:

$$\begin{aligned}
& (a + p^*; y) + (p^*; a); x && \leq p^*; a \\
& a + p^*; y + p^*; a; x + p^*; a && = p^*; a \\
& a + p^*; (a; x + y + a) && = p^*; a \\
& a + p^*; (p; a + a) && = p^*; a \\
& a + p^*; (a; (p + 1)) && = p^*; a \\
& a + p^*; a && = p^*; a \\
& p^*; a && = p^*; a
\end{aligned}$$

□

LEMMA 3.32 (STAR EXPANSION). *If $p \cdot a \equiv a \cdot q + r$ then $p \cdot a \cdot (p \cdot a)^* \equiv (a \cdot q + r) \cdot (q + r)^*$.*

1471 PROOF. First we observe that $p; a; (p; a)^*$ is equivalent to $(p; a)^*; p; a$ (apply KA-SLIDING twice).
 1472 We show two implications using \leq to derive the equality.

1473 (\Rightarrow) We want to show $(p; a)^*; p; a \leq (a; x + y); (x + y)^*$.
 1474

1475 We know that $q + pr \leq r \implies p^*q \leq r$ by the induction axiom KA-LFP-L, so we can
 1476 instantiate it with p and q as $p; a$ and r as $(a; x + y); (x + y)^*$. We find:

$$\begin{aligned}
 1477 & p; a + p; a; (a; x + y); (x + y)^* \leq (a; x + y); (x + y)^* \\
 1478 & p; a + (p; a; x + p; a; y); (x + y)^* \leq (a; x + y); (x + y)^* \\
 1479 & (a; x + y) + ((a; x + y); x + (a; x + y); y); (x + y)^* \leq (a; x + y); (x + y)^* \\
 1480 & (a; x + y) + (a; x + y); (x + y); (x + y)^* \leq (a; x + y); (x + y)^* \\
 1481 & (a; x + y); (1 + (x + y); (x + y)^*) \leq (a; x + y); (x + y)^* \\
 1482 & (a; x + y); (x + y)^* \leq (a; x + y); (x + y)^*
 \end{aligned}$$

1483
 1484 (\Leftarrow) We can to show $(a; x + y); (x + y)^* \leq p; a; (p; a)^*$ We can apply the other induction axiom
 1485 (KA-LFP-R), $q + r; p \leq r \implies q; p^* \leq r$, with $p = x + y$ and $q = a; x + y$ and $r = p; a; (p; a)^*$.
 1486 We find:

$$\begin{aligned}
 1488 & (a; x + y) + p; a; (p; a)^*; (x + y) \leq (p; a)^*; p; a \\
 1489 & p; a + p; a; (p; a)^*; (x + y) \leq (p; a)^*; p; a \\
 1490 & p; a + p; a; (p; a)^*; (x + y) \leq p; a + (p; a)^*; p; a; p; a \\
 1491 & p; a + p; a; (p; a)^*; (x + y) \leq p; a + (p; a)^*; p; a; (a; x + y) \\
 1492 & p; a + p; a; (p; a)^*; (x + y) \leq p; a + (p; a)^*; (a; x + y); (x + y) \\
 1493 & p; a + p; a; (p; a)^*; (x + y) \leq p; a + (p; a)^*; (p; a); (x + y)
 \end{aligned}$$

□

1494
 1495
 1496
 1497 LEMMA 3.33 (PUSHBACK THROUGH PRIMITIVE ACTIONS). *Pushing a test back through a primitive*
 1498 *action leaves the primitive action intact, i.e., if $\pi \cdot a \text{ PB}^\bullet x$ or $(\sum b_i \cdot \pi) \cdot a \text{ PB}^\top x$, then $x = \sum a_i \cdot \pi$.*
 1499

1500 PROOF. By induction on the derivation rule used. □

1501
 1502 We show that our notion of pushback is correct in two steps. First we prove that pushback is
 1503 partially correct, i.e., if we can form a derivation in the pushback relations, the right-hand sides
 1504 are equivalent to the left-hand-sides (Theorem 3.34). Once we've established that our pushback
 1505 relations' derivations mean what we want, we have to show that we can find such derivations;
 1506 here we use our maximal subterm measure to show that the recursive tangle of our PB relations
 1507 always terminates (Theorem 3.35), which makes extensive use of our subterm ordering lemma
 1508 (Lemma 3.22) and splitting (Lemma 3.25)).

1509 THEOREM 3.34 (PUSHBACK SOUNDNESS).
 1510

- 1511 (1) If $x \cdot y \text{ PB}^\downarrow z'$ then $x \cdot y \equiv z'$.
- 1512 (2) If $x^* \text{ PB}^* y$ then $x^* \equiv y$.
- 1513 (3) If $m \cdot a \text{ PB}^\bullet y$ then $m \cdot a \equiv y$.
- 1514 (4) If $m \cdot x \text{ PB}^R y$ then $m \cdot x \equiv y$.
- 1515 (5) If $x \cdot a \text{ PB}^\top y$ then $x \cdot a \equiv y$.

1516
 1517
 1518 PROOF. By simultaneous induction on the derivations. Cases are grouped by judgment.
 1519

1520 *Sequential composition of normal forms* ($x \cdot y \text{ PB}^J z$).

1522 (JOIN) We have $x = \sum_{i=1}^k a_i \cdot m_i$ and $y = \sum_{j=1}^l b_j \cdot n_j$. By the IH on (3), each $m_i \cdot b_j \text{ PB}^\bullet x_{ij}$.
 1523 We compute:

$$\begin{aligned}
 & x \cdot y \\
 \equiv & \left[\sum_{i=1}^k a_i \cdot m_i \right] \cdot \left[\sum_{j=1}^l b_j \cdot n_j \right] \\
 \equiv & \sum_{i=1}^k a_i \cdot m_i \cdot \left[\sum_{j=1}^l b_j \cdot n_j \right] && \text{(KA-DIST-R)} \\
 \equiv & \sum_{i=1}^k a_i \cdot \left[m_i \cdot \sum_{j=1}^l b_j \cdot n_j \right] && \text{(KA-SEQ-ASSOC)} \\
 \equiv & \sum_{i=1}^k a_i \cdot \left[\sum_{j=1}^l m_i \cdot b_j \cdot n_j \right] && \text{(KA-DIST-L)} \\
 \equiv & \sum_{i=1}^k a_i \cdot \left[\sum_{j=1}^l x_{ij} \cdot n_j \right] && \text{(IH (3))} \\
 \equiv & \sum_{i=1}^k \sum_{j=1}^l a_i \cdot x_{ij} \cdot n_j && \text{(KA-DIST-L)}
 \end{aligned}$$

1538 *Kleene star of normal forms* ($x^* \text{ PB}^J y$).

1540 (STARZERO) We have $0^* \text{ PB}^* 1$. We compute:

$$\begin{aligned}
 & 0^* \\
 \equiv & 1 + 0 \cdot 0^* && \text{(KA-UNROLL-L)} \\
 \equiv & 1 + 0 && \text{(KA-ZERO-SEQ)} \\
 \equiv & 1 && \text{(KA-PLUS-ZERO)}
 \end{aligned}$$

1550 (SLIDE) We are trying to pushback the minimal term a of x through a star, i.e., we have
 1551 $(a \cdot x)^*$; by the IH on (5), we know there exists some y such that $x \cdot a \equiv y$; by the
 1552 IH on (2), we know that $y^* \equiv y'$; and by the IH on (1), we know that $y' \cdot x \equiv z$. We
 1553 must show that $(a \cdot x)^* \equiv 1 + a \cdot z$. We compute:

$$\begin{aligned}
 & (a \cdot x)^* \\
 \equiv & 1 + a \cdot x \cdot (a \cdot x)^* && \text{(KA-UNROLL-L)} \\
 \equiv & 1 + a \cdot (x \cdot a)^* \cdot x && \text{(sliding with } p = x \text{ and } q = a; \text{ Lemma 3.29)} \\
 \equiv & 1 + a \cdot y^* \cdot x && \text{(IH (5))} \\
 \equiv & 1 + a \cdot y' \cdot x && \text{(IH (2))} \\
 \equiv & 1 + a \cdot z && \text{(IH (1))}
 \end{aligned}$$

1564 (EXPAND) We are trying to pushback the minimal term a of x through a star, i.e., we have
 1565 $(a \cdot x)^*$; by the IH on (5), we know that there exist t and u such that $x \cdot a \equiv a \cdot t + u$;
 1566 by the IH on (2), we know that there exists a y such that $(t + u)^* \equiv y$; and by the IH

on (1), we know that there is some z such that $y \cdot x \equiv z$. We compute:

$$\begin{aligned}
& (a \cdot x)^* \\
\equiv & 1 + a \cdot x + a \cdot x \cdot a \cdot x \cdot (a \cdot x)^* && \text{(KA-UNROLL-L)} \\
\equiv & 1 + a \cdot x + a \cdot x \cdot a \cdot (x \cdot a)^* \cdot x && \\
& \quad \text{(sliding with } p = x \text{ and } q = a; \text{ Lemma 3.29)} \\
\equiv & 1 + a \cdot x + a \cdot [x \cdot a \cdot (x \cdot a)^*] \cdot x && \text{(KA-SEQ-ASSOC)} \\
\equiv & 1 + a \cdot x + a \cdot [(a \cdot t + u) \cdot (t + u)^*] \cdot x && \\
& \quad \text{(expansion using IH (5); Lemma 3.32)} \\
\equiv & 1 + a \cdot x + a \cdot (a \cdot t + u) \cdot (t + u)^* \cdot x && \text{(KA-SEQ-ASSOC)} \\
\equiv & 1 + a \cdot x + (a \cdot a \cdot t + a \cdot u) \cdot (t + u)^* \cdot x && \text{(KA-DIST-L)} \\
\equiv & 1 + a \cdot x + (a \cdot t + a \cdot u) \cdot (t + u)^* \cdot x && \text{(BA-SEQ-IDEM)} \\
\equiv & 1 + a \cdot x + a \cdot (t + u) \cdot (t + u)^* \cdot x && \text{(BA-SEQ-IDEM)} \\
\equiv & 1 + a \cdot 1 \cdot x + a \cdot (t + u) \cdot (t + u)^* \cdot x && \text{(KA-ONE-SEQ)} \\
\equiv & 1 + (a \cdot 1 + a \cdot (t + u) \cdot (t + u)^*) \cdot x && \text{(KA-DIST-R)} \\
\equiv & 1 + a \cdot (1 + (t + u) \cdot (t + u)^*) \cdot x && \text{(KA-DIST-L)} \\
\equiv & 1 + a \cdot (t + u)^* \cdot x && \text{(KA-UNROLL-L)} \\
\equiv & 1 + a \cdot y \cdot x && \text{(IH (2))} \\
\equiv & 1 + a \cdot z && \text{(IH (1))}
\end{aligned}$$

(DENEST) We have a compound normal form $a \cdot x + y$ under a star; we will push back the maximal test a . By our first IH on (2) we know that that $y^* \equiv y'$ for some y' ; by our first IH on (1), we know that $x \cdot y' \equiv x'$ for some x' ; by our second IH on (2), we know that $(a \cdot x')^* \equiv z$ for some z ; and by our second IH on (1), we know that $y' \cdot z \equiv z'$ for some z' . We must show that $(a \cdot x + y)^* \equiv z'$. We compute:

$$\begin{aligned}
& (a \cdot x + y)^* \\
\equiv & y^* \cdot (a \cdot x \cdot y^*)^* && \text{(denesting with } p = a \cdot x \text{ and } q = y; \text{ Lemma 3.30)} \\
\equiv & y' \cdot (a \cdot x \cdot y')^* && \text{(first IH (2))} \\
\equiv & y' \cdot (a \cdot x')^* && \text{(first IH (1))} \\
\equiv & y' \cdot z && \text{(second IH (2))} \\
\equiv & z' && \text{(second IH (1))}
\end{aligned}$$

*Pushing tests through actions ($m \cdot a$ PB * y).*

(SEQZERO) We are pushing 0 back through a restricted action m . We immediately find $m \cdot 0 \equiv 0$ by KA-SEQ-ZERO.

(SEQONE) We are pushing 1 back through a restricted action m . We find:

$$\begin{aligned}
& m \cdot 1 \\
\equiv & m && \text{(KA-ONE-SEQ)} \\
\equiv & 1 \cdot m && \text{(KA-SEQ-ONE)}
\end{aligned}$$

(SEQSEQTEST) We are pushing the tests $a \cdot b$ through the restricted action m . By our first IH on (3), we have $m \cdot a \equiv y$; by our second IH on (3), we have $y \cdot b \equiv z$. We compute:

$$\begin{aligned}
& m \cdot (a \cdot b) \\
\equiv & m \cdot a \cdot b && \text{(KA-SEQ-ASSOC)} \\
\equiv & y \cdot b && \text{(first IH (3))} \\
\equiv & z && \text{(second IH (3))}
\end{aligned}$$

1618 (SEQSEQACTION) We are pushing the test a through the restricted actions $m \cdot n$. By our IH on (3),
 1619 we have $n \cdot a \equiv x$; by our IH on (4), we have $m \cdot x \equiv y$. We compute:

$$\begin{aligned}
 1620 & (m \cdot n) \cdot a \\
 1621 & \equiv m \cdot (n \cdot a) && \text{(KA-SEQ-ASSOC)} \\
 1622 & \equiv m \cdot x && \text{(IH (3))} \\
 1623 & \equiv y && \text{(IH (4))}
 \end{aligned}$$

1624
 1625
 1626 (SEQPARTTEST) We are pushing the tests $a + b$ through the restricted action m . By our first IH on
 1627 (3), we have $m \cdot a \equiv x$; by our second IH on (3), we have $m \cdot b \equiv y$. We compute:

$$\begin{aligned}
 1628 & m \cdot (a + b) \\
 1629 & m \cdot a + m \cdot b && \text{(KA-DIST-L)} \\
 1630 & \equiv x + m \cdot b && \text{(first IH (3))} \\
 1631 & \equiv x + y && \text{(second IH (3))}
 \end{aligned}$$

1632
 1633
 1634 (SEQPARACTION) We are pushing the test a through the restricted actions $m + n$. By our first IH on
 1635 (3), we have $m \cdot a \equiv x$; by our second IH on (3), we have $n \cdot a \equiv y$. We compute:

$$\begin{aligned}
 1636 & (m + n) \cdot a \\
 1637 & m \cdot a + n \cdot a && \text{(KA-DIST-R)} \\
 1638 & \equiv x + n \cdot a && \text{(first IH (3))} \\
 1639 & \equiv x + y && \text{(second IH (3))}
 \end{aligned}$$

1640
 1641
 1642 (PRIM) We are pushing a primitive predicate α through a primitive action π . We have, by
 1643 assumption, that $\pi \cdot a$ WP $\{a_1, \dots, a_k\}$. By definition of the WP relation, it must
 1644 be the case that $\pi \cdot \alpha \equiv \sum_{i=1}^k a_i \cdot \pi$

1645 (PRIMNEG) We are pushing a negated predicate $\neg a$ back through a primitive action π . We
 1646 have, by assumption, that $\pi \cdot a$ PB $^\bullet \sum_i a_i \cdot pi$ and that $\text{nnf}(\neg(\sum_i a_i)) = b$, so
 1647 $\neg(\sum_i a_i) \equiv b$ (Lemma 3.28). By the IH, we know that $\pi \cdot a \equiv \sum_i a_i \cdot \pi$; we must
 1648 show that $\pi \cdot \neg a \equiv b \cdot \pi$. By our assumptions, we know that $b \cdot \pi \equiv \neg(\sum_i a_i) \cdot \pi$,
 1649 so by pushback negation (PUSHBACK-NEG/Lemma 3.1).

1650
 1651 (SEQSTARSMALLER) We are pushing the test a through the restricted action m^* . By our IH on (3), we
 1652 have $m \cdot a \equiv x$ for some x ; by our IH on (4), we have $m^* \cdot x \equiv y$ for some y . We
 1653 compute:

$$\begin{aligned}
 1654 & m^* \cdot a \\
 1655 & \equiv (1 + m^* \cdot m) \cdot a && \text{(KA-UNROLL-R)} \\
 1656 & \equiv a + m^* \cdot m \cdot a && \text{(KA-DIST-R)} \\
 1657 & \equiv a + m^* \cdot (m \cdot a) && \text{(KA-SEQ-ASSOC)} \\
 1658 & \equiv a + m^* \cdot x && \text{(IH (3))} \\
 1659 & \equiv a + y && \text{(IH (4))}
 \end{aligned}$$

1660
 1661
 1662 (SEQSTARINV) We are pushing the test a through the restricted action m^* . By our IH on (3), there
 1663 exist t and u such that $m \cdot a \equiv a \cdot t + u$; by our IH on (4), there exists an x such
 1664 that $m^* \cdot u \equiv x$; by our IH on (2), there exists a y such that $u^* \equiv y$; and by our IH
 1665 on (1), there exists a z such that $x \cdot y \equiv z$. We compute:

1666

$$\begin{aligned}
1667 \quad & m \cdot a = a \cdot t + u \ m^* a = (a + m^* \cdot u) + t^* \\
1668 \quad & \\
1669 \quad & \quad m^* \cdot a \\
1670 \quad & \equiv (a + m^* \cdot u) \cdot t^* \quad (\text{star invariant on IH (3); Lemma 3.31}) \\
1671 \quad & \equiv a \cdot t^* + m^* \cdot u \cdot t^* \quad (\text{KA-DIST-R}) \\
1672 \quad & \equiv a \cdot t^* + x \cdot t^* \quad (\text{IH (4)}) \\
1673 \quad & \equiv a \cdot y + x \cdot y \quad (\text{IH (2)}) \\
1674 \quad & \equiv a \cdot y + z \quad (\text{IH (1)})
\end{aligned}$$

1675 *Pushing normal forms through actions ($m \cdot x \text{ PB}^R z$).*

1676 (RESTRICTED) We have $x = \sum_{i=1}^k a_i \cdot n_i$. By the IH on (3), $m \cdot a_i \text{ PB}^\bullet y_i$. We compute:

$$\begin{aligned}
1677 \quad & \\
1678 \quad & \quad m \cdot x \\
1679 \quad & \equiv m \cdot \sum_{i=1}^k a_i \cdot n_i \\
1680 \quad & \equiv \sum_{i=1}^k m \cdot a_i \cdot n_i \quad (\text{KA-DIST-L}) \\
1681 \quad & \equiv \sum_{i=1}^k y_i \cdot n_i \quad (\text{IH (3)})
\end{aligned}$$

1682 *Pushing tests through normal forms ($x \cdot a \text{ PB}^T y$).*

1683 (TEST) We have $x = \sum_{i=1}^k a_i \cdot m_i$. By the IH on (3), we have $m_i \cdot a \text{ PB}^\bullet y_i$ where $y_i = \sum_{j=1}^l b_{ij} \cdot m_{ij}$. We compute:

$$\begin{aligned}
1684 \quad & \\
1685 \quad & \quad x \cdot a \\
1686 \quad & \equiv \left[\sum_{i=1}^k a_i \cdot m_i \right] \cdot a \\
1687 \quad & \equiv \sum_{i=1}^k a_i \cdot m_i \cdot a \quad (\text{KA-DIST-R}) \\
1688 \quad & \equiv \sum_{i=1}^k a_i \cdot (m_i \cdot a) \quad (\text{KA-SEQ-ASSOC}) \\
1689 \quad & \equiv \sum_{i=1}^k a_i \cdot y_i \quad (\text{IH (3)}) \\
1690 \quad & \equiv \sum_{i=1}^k a_i \cdot \sum_{j=1}^l b_{ij} \cdot m_{ij} \\
1691 \quad & \equiv \sum_{i=1}^k \sum_{j=1}^l a_i \cdot b_{ij} \cdot m_{ij} \quad (\text{KA-DIST-L})
\end{aligned}$$

□

1692 **THEOREM 3.35 (PUSHBACK EXISTENCE).** *For all x and m and a :*

- 1693 (1) *For all y and z , if $x \leq z$ and $y \leq z$ then there exists some $z' \leq z$ such that $x \cdot y \text{ PB}^J z'$.*
- 1694 (2) *There exists a $y \leq x$ such that $x^* \text{ PB}^\bullet y$.*
- 1695 (3) *There exists some $y \leq a$ such that $m \cdot a \text{ PB}^\bullet y$.*
- 1696 (4) *There exists a $y \leq x$ such that $m \cdot x \text{ PB}^R y$.*
- 1697 (5) *If $x \leq z$ and $a \leq z$ then there exists a $y \leq z$ such that $x \cdot a \text{ PB}^T y$.*

1698 **PROOF.** By induction on the lexicographical order of: the subterm ordering ($<$); the size of x (for (1), (2), (4), and (5)); the size of m (for (3) and (4)); and the size of a (for (3)).

1699 *Sequential composition of normal forms ($x \cdot y \text{ PB}^J z$).* We have $x = \sum_{i=1}^k a_i \cdot m_i$ and $y = \sum_{j=1}^l b_j \cdot n_j$; by the IH on (3) with the size decreasing on m_i , we know that $m_i \cdot b_j \text{ PB}^\bullet x_{ij}$ for each i and j such that $x_{ij} \leq a_i$, so by JOIN, we know that $x \cdot y \text{ PB}^J \sum_{i=1}^k \sum_{j=1}^l a_i x_{ij} n_j = z'$.

1700 Given that $x, y \leq z$, it remains to be seen that $z' \leq z$. We've assumed that $a_i \leq x \leq z$. By our IH on (3) we found earlier that $x_{ij} \leq a_i \leq z$. Therefore, by unpacking x and applying test bounding (Lemma 3.22), $a_i \cdot x_{ij} \cdot n_j \leq z$. By normal form parallel congruence (Lemma 3.22), we have $z' \leq z$.

1716 *Kleene star of normal forms* ($x^* \text{PB}^\downarrow y$). If x is vacuous, we find that $0^* \text{PB}^* 1$ by STARZERO, with
 1717 $1 \leq 0$ since they have the same maximal terms (just 1).

1718 If x isn't vacuous, then we have $x \equiv a \cdot x_1 + x_2$ where $x_1, x_2 < x$ and $a \in \text{mt}(x)$ by splitting
 1719 (Lemma 3.25). We first consider whether x_2 is vacuous.

1720 (x_2 is vacuous) We have $x \equiv a \cdot x_1 + 0 \equiv a \cdot x_1$.

1721 By our IH on (5) with x_1 decreasing in size, we have $x_1 \cdot a \text{PB}^\uparrow w$ where $w \leq x$
 1722 (because $x_1 < x$ and $a \leq x$). By maximal test inequality (Lemma 3.24), we have
 1723 two cases: either $a \in \text{mt}(w)$ or $w < a \leq x$.

1725 ($a \in \text{mt}(w)$) By splitting (Lemma 3.25), we have $w \equiv a \cdot t + u$ for some normal
 1726 forms $t, u < w$.

1727 By normal-form parallel congruence (Lemma 3.22), $t + u < x$; so
 1728 by the IH on (2) with our subterm ordering decreasing on $t + u < x$,
 1729 we find that $(t + u)^* \text{PB}^* w'$ for some $w' \leq (t + u)^* < w \leq x$.
 1730 Since $w' < x$, we can apply our IH on (1) with our subterm
 1731 ordering decreasing on $w' < x$ to find that $w' \cdot x_1 \text{PB}^\downarrow z$ such that
 1732 $z \leq x_1 < x$ (since $w' \leq x$ and $x_1 < x$).

1733 Finally, we can see by EXPAND that $x = (a \cdot x_1)^* \text{PB}^* 1 + a \cdot z = y$.
 1734 Since each $1, a, z \leq x$, we have $y = 1 + a \cdot z \leq x$ as needed.

1735 ($w < a$) Since $w < a$, we can apply our IH on (2) with our subterm order
 1736 decreasing on $w < x$ to find that $w^* \text{PB}^* w'$ such that $w' \leq w <$
 1737 $a \leq x$. By our IH on (1) with our subterm order decreasing on
 1738 $w' < x$ to find that $w' \cdot x_1 \text{PB}^\downarrow z$ where $z \leq x$ (because $w' \leq x$
 1739 and $x_1 < x$).

1740 We can now see by SLIDE that $x = (a \cdot x_1)^* \text{PB}^* 1 + a \cdot z = y$. Since
 1741 each $1, a, z \leq x$, we have $y = 1 + a \cdot z \leq x$ as needed.

1742 (x_2 isn't vacuous) We have $x \equiv a \cdot x_1 + x_2$ where $x_i < x$ and $a \in \text{mt}(x)$. Since x_2 isn't vacuous, we
 1743 must have $a < x$, not just $a \leq x$.

1744 By the IH on (2) with the subterm ordering decreasing on $x_2 < x$, we find $x_2 \text{PB}^* w$
 1745 such that $w \leq x_2$. By the IH on (1) with the subterm ordering decreasing on $x_1 < x$,
 1746 we have $x_1 \cdot w \text{PB}^\downarrow v$ where $v \leq x$ (because $x_1 \leq x$ and $w \leq x$). By the IH on (2)
 1747 with the subterm ordering decreasing on $a \cdot v < x$, we find $(a \cdot v)^* \text{PB}^* z$ where
 1748 $z \leq a \cdot v < x$. By our IH on (1) with the subterm ordering decreasing on $w < x$,
 1749 we find $w \cdot z \text{PB}^\downarrow y$ where $y < x$ (because $w < x$ and $z < x$).

1750 By DENEST, we can see that $x \equiv (a \cdot x_1 + x_2)^* \text{PB}^* y$, and we've already found
 1751 that $y \leq x$ as needed.

1752 *Pushing tests through actions* ($m \cdot a \text{PB}^\bullet y$). We go by cases on a and m to find the $y \leq a$ such
 1753 that $m \cdot a \text{PB}^\bullet y$.

1754 ($m, 0$) We have $m \cdot 0 \text{PB}^\bullet 0$ by SEQZERO, and $0 \leq 0$ immediately.

1755 ($m, 1$) We have $m \cdot 1 \text{PB}^\bullet 1 \cdot m$ by SEQONE and $1 \leq 1$ immediately.

1756 ($m, a \cdot b$) By the IH on (3) decreasing in size on a , we know that $m \cdot a \text{PB}^\bullet x$ where $x \leq a \leq a \cdot b$.
 1757 By the IH on (5) decreasing in size on b , we know that $x \cdot b \text{PB}^\uparrow y$. Finally, we know
 1758 by SEQSEQTEST that $m \cdot (a \cdot b) \text{PB}^\bullet y$. Since $x \leq a \cdot b$ and $b \leq a \cdot b$, we know by the
 1759 IH on (5) earlier that $y \leq a \cdot b$.

1760

- 1765 $(m, a + b)$ By the IH on (3) decreasing in size on a , we know that $m \cdot a \text{ PB}^\bullet x$ such that $x \leq a \leq$
 1766 $a + b$. Similarly, by the IH on (3) decreasing in size on b , we know that $m \cdot b \text{ PB}^\bullet z$
 1767 such that $z \leq b \leq a + b$. By SEQPARTEST, we know that $m \cdot (a + b) \text{ PB}^\bullet x + z = y$;
 1768 by normal form parallel congruence, we know that $y = x + z \leq a + b$ as needed.
- 1769 $(m \cdot n, a)$ By the IH on (3) decreasing in size on n , we know that $n \cdot a \text{ PB}^\bullet x$ such that $x \leq a$.
 1770 By the IH on (4) decreasing in size on m , we know that $m \cdot x \text{ PB}^R y$ such that
 1771 $y \leq x \leq a$ (which are the size bounds on y we needed to show). All that remains to
 1772 be seen is that $(m \cdot n) \cdot a \text{ PB}^\bullet y$, which we have by SEQSEQACTION.
- 1773 $(m + n, a)$ By the IH on (3) decreasing in size on m , we know that $m \cdot a \text{ PB}^\bullet x$. Similarly, by
 1774 the IH on (3) decreasing in size on n , we know that $n \cdot a \text{ PB}^\bullet z$. By SEQPARACTION,
 1775 we know that $(m + n) \cdot a \text{ PB}^\bullet x + z = y$. Furthermore, both IHs let us know that
 1776 $x, z \leq a$, so by normal form parallel congruence, we know that $y = x + z \leq a$.
- 1777 $(\pi, \neg a)$ By the IH on (3) decreasing in size on a , we can find that $\pi \cdot a \text{ PB}^\bullet \sum_i a_i \cdot \pi$ where
 1778 $\sum_i a_i \leq a$, and $\text{nnf}(\neg(\sum_i a_i)) = b$ for some term b . It remains to be seen that $b \leq \neg a$,
 1779 which we have by monotonicity of nnf (Lemma 3.21).
- 1780 (π, α) In this case, we fall back on the client theory's pushback operation (Definition 3.26).
 1781 We have $\pi \cdot \alpha \text{ WP } \{a_1, \dots, a_k\}$ such that $a_i \leq \alpha$. By PRIM, we have $\pi \cdot \alpha \text{ PB}^\bullet$
 1782 $\sum_{i=1}^k a_i \cdot \pi = y$; since each $a_i \leq \alpha$, we find $y \leq \alpha$ by the monotonicity of union
 1783 (Lemma 3.18).
- 1784 (m^*, a) We've already ruled out the case where $a = b \cdot c$, so it must be the case that
 1785 $\text{seqs}(a) = \{a\}$, so $\text{mt}(a) = \{a\}$.
 1786 By the IH on (3) decreasing in size on m , we know that $m \cdot a \text{ PB}^\bullet x$ such that $x \leq a$.
 1787 There are now two possibilities: either $x < a$ or $a \in \text{mt}(x) = \{a\}$.
- 1788 $(x < a)$ By the IH on (4) with $x < a$, we know by SEQSTARSMALLER that $m^* \cdot x \text{ PB}^R y$
 1789 such that $y \leq x < a$.
- 1790 $(a \in \text{mt}(x))$ By splitting (Lemma 3.25), we have $x \equiv a \cdot t + u$, where t and u are normal
 1791 forms such that $t, u < x \leq a$.
 1792 By the IH on (4) with $t < a$, we know that $m^* \cdot t \text{ PB}^R w$ such that $w \leq$
 1793 $t < x \leq a$. By the IH on (2) with $u < x \leq a$, we know that $u^* \text{ PB}^\bullet z$ such
 1794 that $z \leq u < x \leq a$. By the IH on (1) with $w < a$ and $z < a$, we find that
 1795 $w \cdot z \text{ PB}^J v$ such that $v \leq w < a$.
 1796 Finally we have our y : by SEQSTARINV, we have $m^* \cdot a \text{ PB}^\bullet a \cdot z + v = y$.
 1797 Since $z \leq a$ and $a \leq a$, we have $a \cdot z \leq a$ (mixed sequence congruence;
 1798 Lemma 3.22) and $v < a$. By normal form parallel congruence, we have
 1799 $a \cdot z + v \leq a$ (Lemma 3.22).
- 1800 *Pushing normal forms through actions* ($m \cdot x \text{ PB}^R z$). We have $x = \sum_{i=1}^k a_i \cdot n_i$; by the IH on (3)
 1801 with the size decreasing on n_i , we know that $m \cdot a_i \text{ PB}^\bullet x_i$ for each i such that $x_i \leq a_i$, so by
 1802 RESTRICTED, we know that $m \cdot x \text{ PB}^R \sum_{i=1}^k x_i n_i = y$.
 1803 We must show that $y \leq x$. By our IH on (3) we found earlier that $x_i \leq a_i$. By normal form parallel
 1804 congruence (Lemma 3.22), we have $y \leq x$.
- 1805 *Pushing tests through normal forms* ($x \cdot a \text{ PB}^T y$). We have $x = \sum_{i=1}^k a_i \cdot m_i$; by the IH on (3) with
 1806 the size decreasing on m_i , we know that $m_i \cdot a \text{ PB}^\bullet y_i = \sum_{j=1}^l b_{ij} m_{ij}$ where $y_i \leq a$. Therefore, we
 1807 know that $x \cdot a \text{ PB}^T \sum_{i=1}^k \sum_{j=1}^l a_i \cdot b_{ij} \cdot m_{ij} = y$ by TEST.

1814 Given that $x \leq z$ and $a \leq z$, We must show that $y \leq z$. We already know that $a_i \leq x \leq z$, and we
 1815 found from the IH on (3) earlier that $b_{ij} \leq y_i \leq a \leq z$. By test bounding (Lemma 3.22), we have
 1816 $a_i \cdot b_{ij} \leq z$, and therefore $y \leq z$ by normal form parallel congruence (Lemma 3.22).
 1817 □

1818 Finally, to reiterate our discussion of PB^\bullet , Theorem 3.35 shows that every left-hand side of the
 1819 pushback relation has a corresponding right-hand side. We *haven't* proved that the pushback
 1820 relation is functional— if a term has more than one maximal test, there could be many different
 1821 choices of how we perform the pushback.
 1822 Now that we can push back tests, we can show that every term has an equivalent normal form.

1823
 1824 **COROLLARY 3.36 (NORMAL FORMS).** *For all $p \in \mathcal{T}^*$, there exists a normal form x such $p \text{ norm } x$*
 1825 *and that $p \equiv x$.*

1826 **PROOF.** By induction on p .

1827 (PRED) We have $a \equiv a$ immediately.

1828 (ACT) We have $\pi \equiv 1 \cdot \pi$ by KA-SEQ-ONE.

1829 (PAR) By the IHs and congruence.

1830 (SEQ) We have $p = q \cdot r$; by the IHs, we know that $q \text{ norm } x$ and $r \text{ norm } y$. By pushback
 1831 existence (Theorem 3.35), we know that $x \cdot y \text{ PB}^1 z$ for some z . By pushback
 1832 soundness (Theorem 3.34), we know that $x \cdot y \equiv z$. By congruence, $p \equiv z$.
 1833

1834 (STAR) We have $p = q^*$. By the IH, we know that $q \text{ norm } x$. By pushback existence
 1835 (Theorem 3.35), we know that $x^* \text{ PB}^* y$. By pushback soundness (Theorem 3.34),
 1836 we know that $x^* \equiv y$.
 1837 □

1838
 1839
 1840 The PB relations and these two proofs are one of the contributions of this paper: we believe it is
 1841 the first time that a KAT normalization procedure has been made so explicit, rather than hiding
 1842 inside of completeness proofs. Temporal NetKAT, which introduced the idea of pushback, proved a
 1843 concretization of Theorems 3.34 and 3.35 as a single theorem and without any explicit normalization
 1844 or pushback relation.
 1845

1846 3.4 Completeness

1847 We prove completeness—if $\llbracket p \rrbracket = \llbracket q \rrbracket$ then $p \equiv q$ —by normalizing p and q and comparing the
 1848 resulting terms. Our completeness proof uses the completeness of Kleene algebra (KA) as its
 1849 foundation: the set of possible traces of actions performed for a restricted (test-free) action in our
 1850 denotational semantics is a regular language, and so the KA axioms are sound and complete for it. In
 1851 order to relate our denotational semantics to regular languages, we define the regular interpretation
 1852 of restricted actions $m \in \mathcal{T}_{\text{RA}}$ in the conventional way and then relate our denotational semantics
 1853 to the regular interpretation (Fig. 17). Readers familiar with NetKAT's completeness proof may
 1854 notice that we've omitted the language model and gone straight to the regular interpretation. We're
 1855 able to shorten our proof because our tracing semantics is more directly relatable to its regular
 1856 interpretation, and because our completeness proof separately defers to the client theory's decision
 1857 procedure for the predicates at the front. Our normalization routine—the essence of our proof—only
 1858 uses the KAT axioms and doesn't rely on any property of our tracing semantics. We conjecture
 1859 that one could prove a similar completeness result and derive a similar decision procedure with
 1860 a merging, non-tracing semantics, like in NetKAT or KAT+B! [1, 30]. We haven't attempted the
 1861 proof or an implementation.
 1862

1863	$\mathcal{R} : \overline{\mathcal{T}}_{\text{RA}} \rightarrow \mathcal{P}(\Pi_{\mathcal{T}}^*)$		$\text{label} : \text{Trace} \rightarrow \Pi_{\mathcal{T}}^*$
1864	$\mathcal{R}(1) = \{\epsilon\}$		$\text{label}(\langle\sigma, \perp\rangle) = \epsilon$
1865	$\mathcal{R}(\pi) = \{\pi\}$		$\text{label}(t\langle\sigma, \pi\rangle) = \text{label}(t)\pi$
1866	$\mathcal{R}(m+n) = \mathcal{R}(m) \cup \mathcal{R}(n)$		
1867	$\mathcal{R}(m \cdot n) = \{uv \mid u \in \mathcal{R}(m), v \in \mathcal{R}(n)\}$		$\mathcal{L}^0 = \{\epsilon\}$
1868	$\mathcal{R}(m^*) = \bigcup_{0 \leq i} \mathcal{R}(m)^i$		$\mathcal{L}^{n+1} = \{uv \mid u \in \mathcal{L}, v \in \mathcal{L}^n\}$

Fig. 17. Regular interpretation of restricted actions

LEMMA 3.37 (RESTRICTED ACTIONS ARE AHISTORICAL). *If $\llbracket m \rrbracket(t_1) = t_1, t$ and $\text{last}(t_1) = \text{last}(t_2)$ then $\llbracket m \rrbracket(t_2) = t_2, t$.*

PROOF. By induction on m .

($m = 1$) Immediate, since t is empty.

($m = \pi$) We immediately have $t = \langle \text{last}(t_1), \pi \rangle$.

($m = m+n$) We have $\llbracket m+n \rrbracket(t_1) = \llbracket m \rrbracket(t_1) \cup \llbracket n \rrbracket(t_1)$ and $\llbracket m+n \rrbracket(t_2) = \llbracket m \rrbracket(t_2) \cup \llbracket n \rrbracket(t_2)$. By the IHs.

($m = m \cdot n$) We have $\llbracket m \cdot n \rrbracket(t_1) = (\llbracket m \rrbracket \bullet \llbracket n \rrbracket)(t_1)$ and $\llbracket m \cdot n \rrbracket(t_2) = (\llbracket m \rrbracket \bullet \llbracket n \rrbracket)(t_2)$. It must be that $\llbracket m \rrbracket(t_1) = \{t_1, t_{mi}\}$, so by the IH we have $\llbracket m \rrbracket(t_2) = \{t_2, t_{mi}\}$. These sets have the same last states, so we can apply the IH again for n , and we are done.

($m = m^*$) We have $\llbracket m^* \rrbracket(t_1) = \bigcup_{0 \leq i} \llbracket m \rrbracket^i(t_1)$. By induction on i .

($i = 0$) Immediate, since $\llbracket m \rrbracket^0(t_i) = t_i$ and so t is empty.

($i = i+1$) By the IH and the reasoning above for \cdot .

□

LEMMA 3.38 (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 .

($m = 1$) We have $\mathcal{R}(1) = \{\epsilon\}$. For all σ , we find $\llbracket 1 \rrbracket(\langle\sigma, \perp\rangle) = \{\langle\sigma, \perp\rangle\}$, and $\text{label}(\langle\sigma, \perp\rangle) = \epsilon$.

($m = \pi$) We $\mathcal{R}(\pi) = \{\pi\}$. For all σ , we find $\llbracket \pi \rrbracket(\langle\sigma, \perp\rangle) = \{\langle\sigma, \perp\rangle \langle \text{act}(\pi, \sigma), \pi \rangle\}$, and so $\text{label}(\langle\sigma, \perp\rangle \langle \text{act}(\pi, \sigma), \pi \rangle) = \pi$.

($m = m+n$) We have $\mathcal{R}(m+n) = \mathcal{R}(m) \cup \mathcal{R}(n)$. For all σ , we have:

$$\begin{aligned} \text{label}(\llbracket m+n \rrbracket(\langle\sigma, \perp\rangle)) &= \text{label}(\llbracket m \rrbracket(\langle\sigma, \perp\rangle) \cup \llbracket n \rrbracket(\langle\sigma, \perp\rangle)) \\ &= \text{label}(\llbracket m \rrbracket(\langle\sigma, \perp\rangle)) \cup \text{label}(\llbracket n \rrbracket(\langle\sigma, \perp\rangle)) \end{aligned}$$

and we are done by the IHs.

($m = m \cdot n$) We have $\mathcal{R}(m \cdot n) = \{uv \mid u \in \mathcal{R}(m), v \in \mathcal{R}(n)\}$. For all σ , we have:

$$\begin{aligned} \text{label}(\llbracket m \cdot n \rrbracket(\langle\sigma, \perp\rangle)) &= \text{label}(\llbracket m \rrbracket \bullet \llbracket n \rrbracket)(\langle\sigma, \perp\rangle) \\ &= \text{label}(\bigcup_{t \in \llbracket m \rrbracket(\langle\sigma, \perp\rangle)} \text{label}(\llbracket n \rrbracket(t))) \\ &= \text{label}(\bigcup_{t \in \llbracket m \rrbracket(\langle\sigma, \perp\rangle)} \text{label}(t \llbracket n \rrbracket(\langle\sigma, \perp\rangle))) \quad \text{by Lemma 3.37} \\ &= \text{label}(\llbracket m \rrbracket(\langle\sigma, \perp\rangle)) \text{label}(\llbracket n \rrbracket(\langle\sigma, \perp\rangle)) \end{aligned}$$

and we are done by the IHs.

($m = m^*$) We have $\mathcal{R}(m^*) = \bigcup_{0 \leq i} \mathcal{R}(m)^i$. For all σ , we have:

$$\begin{aligned} \text{label}(\llbracket m^* \rrbracket(\langle\sigma, \perp\rangle)) &= \text{label}(\bigcup_{0 \leq i} \llbracket m \rrbracket^i(\langle\sigma, \perp\rangle)) \\ &= \bigcup_{0 \leq i} \text{label}(\llbracket m \rrbracket^i(\langle\sigma, \perp\rangle)) \end{aligned}$$

and we are done by the IH. □

Our proof of completeness works by normalizing each side of the equation, making each side locally unambiguous, making the entire equation unambiguous, and then using word equality to ensure that normal forms with equivalent predicates have equivalent actions.

THEOREM 3.39 (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 3.36); by soundness (Theorem 3.5), 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's predicates must be unambiguous locally, and then the predicates must be pairwise comparable 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 the predicates a_i (excluding the case where we select none) and run the actions corresponding to the true predicates, i.e., m_i gets run precisely when a_i is true:

$$\left. \begin{array}{l} \hat{x} = a_1 \cdot a_2 \cdot \dots \cdot a_n \cdot (m_1 + m_2 + \dots + m_n) \\ + \neg a_1 \cdot a_2 \cdot \dots \cdot a_n \cdot (m_2 + \dots + m_n) \\ + a_1 \cdot \neg a_2 \cdot \dots \cdot a_n \cdot (m_1 + \dots + m_n) \\ + \dots \\ + \neg a_1 \cdot \neg a_2 \cdot \dots \cdot a_n \cdot m_n \\ + \neg a_1 \cdot \neg a_2 \cdot \dots \cdot \neg a_n \cdot 0 \end{array} \right\} \text{all combinations of } a_i \text{ (} 2^n \text{ summands)}$$

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

Observe that the sum of all of the predicates in \hat{x} and \hat{y} are respectively equivalent to 1, since it enumerates all possible combinations of each a_i (BA-PLUS-DIST, BA-EXCL-MID); i.e., if $\hat{x} = \sum_i c_i \cdot l_i$ and $\hat{y} = \sum_j d_j \cdot m_j$, then $\sum_i c_i \equiv 1$ and $\sum_j d_j \equiv 1$. We can take advantage of exhaustiveness of these sums to translate the locally disjoint but syntactically unequal predicates in each \hat{x} and \hat{y} to a single set of predicates on both, which allows us to do a *syntactic* comparison on each of the predicates. Let \tilde{x} and \tilde{y} be the extension of \hat{x} and \hat{y} with the tests from the other form, giving us $\tilde{x} = \sum_{i,j} c_i \cdot d_j \cdot l_i$ and $\tilde{y} = \sum_{i,j} c_i \cdot d_j \cdot m_j$. Extending the normal forms to be disjoint between the two normal forms is still provably equivalent using commutativity (BA-SEQ-COMM) and the exhaustiveness above (KA-SEQ-ONE).

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 have assumed that the predicates of \mathcal{T} can be checked for emptiness, so we can eliminate those cases where the expanded tests at the front of \tilde{x} and \tilde{y} are equivalent to zero, which is sound by the client theory's completeness and zero-cancellation (KA-ZERO-SEQ and KA-SEQ-ZERO). If one normal form is empty, the other one must be empty as well.

1961 Finally, we can defer to deductive completeness for KA to find proofs that the commands are
 1962 equivalent. To use KA's completeness to get a proof over commands, we have to show that if our
 1963 commands have equal denotations in our semantics, then they will also have equal denotations
 1964 in the KA semantics. We've done exactly this by showing that restricted actions have regular
 1965 interpretations: because the zero-canceled \hat{x} and \hat{y} are provably equivalent, soundness guarantees
 1966 that their denotations are equal. Since their tests are pairwise disjoint, if their denotations are
 1967 equal, it must be that any non-canceled commands are equal, which means that each label of these
 1968 commands must be equal—and so $\mathcal{R}(l_i) = \mathcal{R}(o_j)$ (Lemma 3.38). By the deductive completeness of
 1969 KA, we know that $\text{KA} \vdash l_i \equiv o_j$. Since we have the KA axioms in our system, then $l_i \equiv o_j$; by
 1970 reflexivity, we know that $c_i \cdot d_j \equiv c_i \cdot d_j$, and we have proved that $\hat{x} \equiv \hat{y}$. By transitivity, we can see
 1971 that $\hat{x} \equiv \hat{y}$ and so $x \equiv y$ and $p \equiv q$, as desired. \square

1972

1973 4 IMPLEMENTATION

1974 We have implemented our ideas in an OCaml library; the library's source code, tests, and our
 1975 evaluation workbench are available online.¹ Sec. 1.3 summarizes the high-level idea and gives an
 1976 example library implementation for the theory of increasing natural numbers. To use a higher-order
 1977 theory such as that of product theories, one need only instantiate the appropriate modules in the
 1978 library:

1979

```
1980 module P = Product(IncNat)(Boolean)
1981 module D = Decide(P) (* normalization-based decision procedure *)
1982 let a = P.K.parse "y<1; (a=F + a=T; inc(y)); y>0" in
1983 let b = P.K.parse "y<1; a=T; inc(y)" in
1984 assert (D.equivalent a b)
```

1985

1986 The module P instantiates Product over our theories of incrementing naturals and booleans; the
 1987 module D gives a way to normalize terms based on the completeness proof. User's of the library
 1988 can access these representations to perform any number of tasks such as compilation, verification,
 1989 inference, and so on. For example, checking language equivalence is then simply a matter of reading
 1990 in KMT terms and calling the equivalence function. Our implementation currently supports both
 1991 a decision procedure based on automata (not yet completely correct, and so omitted from this
 1992 article) and one based on the normalization term-rewriting from the completeness proof. We've
 1993 implemented a command-line tool that can be configured to work these theories; given a variety of
 1994 KMT terms as input, it partitions them into equivalence classes using the decision procedure of the
 1995 user's choice.

1996

1997 4.1 Optimizations

1998 In practice, our implementation uses several optimizations, with the two most prominent being (1)
 1999 hash-consing all KAT terms to ensure fast set operations, and (2) lazy construction and exploration
 2000 of automata during equivalence checking.

2001 Our hash-consing constructors are *smart* constructors, automatically rewriting common identities
 2002 (e.g., constructing $p \cdot 1$ will simply return p ; constructing $(p^*)^*$ will simply return p^*). Client theories
 2003 can extend our smart constructors to witness theory-specific identities. These optimizations are
 2004 partly responsible for the speed of our normalization routine (when it avoids the costly DENEST case).
 2005 When deciding equivalence using normalization, we use the Hopcroft and Karp algorithm [32] on
 2006 implicit automata using the Brzozowski derivative [9] to generate the transition relation on-the-fly.

2007

2008 ¹<https://github.com/mgree/kmt>

2009

	Benchmark	\mathcal{T}	Time to check equivalence
2010			
2011	$a^* \neq a$ (for random arithmetic predicate a)	\mathbb{N}	0.034s
2012	$\text{inc}_x^*; x > 10 \equiv \text{inc}_x^*; \text{inc}_x^*; x > 10$	\mathbb{N}	<0.001s
2013	$\text{inc}_x^*; x > 3; \text{inc}_y^*; y > 3 \equiv \text{inc}_x^*; \text{inc}_y^*; x > 3; y > 3$	\mathbb{N}	<0.001s
2014	$x = \bar{f}; (\text{flip } x; \text{flip } x)^* \equiv (\text{flip } x; \text{flip } x)^*; x = \bar{f}$	\mathcal{B}	<0.001s
2015	$w := \bar{f}; x := t; y := \bar{f}; z := \bar{f};$		
2016	$((w = t + x = t + y = t + z = t); a := t +$		
2017	$(\neg(w = t + x = t + y = t + z = t)); a := \bar{f})$		
2018	$\equiv w := \bar{f}; x := t; y := \bar{f}; z := \bar{f};$	\mathcal{B}	<0.001s
2019	$((w = t + x = t) + (y = t + z = t)); a := t +$		
2020	$(\neg((w = t + x = t) + (y = t + z = t))); a := \bar{f})$		
2021	$y < 1; a = t; \text{inc}_y;$		
2022	$(1 + b = t; \text{inc}_y);$		
2023	$(1 + c = t; \text{inc}_y); y > 2$	$\mathbb{N} \times \mathcal{B}$	0.309s
2024	$\equiv y < 1; a = t; b = t; c = t; \text{inc}_y; \text{inc}_y; \text{inc}_y$		
2025	$(\text{flip } x + \text{flip } y + \text{flip } z)^* = (\text{flip } x + \text{flip } y + \text{flip } z)^*$	\mathbb{B}	>30s (timeout)
2026			

Fig. 18. Implementation microbenchmarks

Client theories can implement custom solvers or rely on Z3 embeddings—custom solvers are typically faster. We’ve implemented a few of these domain-specific optimizations: satisfiability procedure for IncNat makes a heuristic decision between using our incomplete custom solver or Z3 [18]—our solver is much faster on its restricted domain.

5 EVALUATION

We performed a few experiments to evaluate our tool on a collection of simple microbenchmarks. Fig. 18 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. We use a timeout of thirty seconds.

Our 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, 37]) 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 fourth, parity-swapping benchmark is one such example – it flips the parity of a boolean variable an even number of times 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.

6 RELATED WORK

Kozen and Mamouras’s Kleene algebra with equations [40] is perhaps the most closely related work: they also devise a framework for proving extensions of KAT sound and complete. Our works share a similar genesis: Kleene algebra with equations generalizes the NetKAT completeness proof (and then reconstructs it); our work generalizes the Temporal NetKAT completeness proof (and then

reconstructs it—while also developing several other, novel KATs). Both their work and ours use rewriting to find normal forms and prove deductive completeness. Their rewriting systems work on mixed sequences of actions and predicates, but they can only delete these sequences wholesale or replace them with a single primitive action or predicate; our rewriting system’s pushback operation only works on predicates (since the trace semantics preserves the order of actions), but pushback isn’t restricted to producing at most a single primitive predicate. Each framework can do things the other cannot. Kozen and Mamouras can accommodate equations that combine actions, like those that eliminate redundant writes in KAT+B! and NetKAT [1, 30]; we can accommodate more complex predicates and their interaction with actions, like those found in Temporal NetKAT [8] or those produced by the compositional theories (Sec. 2). It may be possible to build a hybrid framework, with ideas from both. A trace semantics occurs in previous work on KAT as well [26, 37].

Kozen studies KATs with arbitrary equations $x := e$ [38], also called Schematic KAT, where e comes from arbitrary first-order structures over a fixed signature Σ . He has a pushback-like axiom $x := e \cdot p \equiv \phi[x/e] \cdot x := e$. Arbitrary first-order structures over Σ ’s theory are much more expressive than anything we can handle—the pushback may or may not decrease in size, depending on Σ ; KATs over such theories are generally undecidable. We, on the other hand, are able to offer pay-as-you-go results for soundness and completeness as well as an implementations for deciding equivalence—but only for first-order structures that admit a non-increasing weakest precondition. Other extensions of KAT often give up on decidability, too. Larsen et al. [42] allow comparison of variables, but this of course leads to an incomplete theory. They are, able, however, to decide emptiness of an entire expression.

Coalgebra provides a general framework for reasoning about state-based systems [39, 54, 59], which has proven useful in the development of automata theory for KAT extensions. Although we do not explicitly develop the connection in this paper, KMT uses tools similar to those used in coalgebraic approaches, and one could perhaps adapt our theory and implementation to that setting. In principle, we ought to be able to combine ideas from the two schemes into a single, even more general framework that supports complex actions *and* predicates.

Smolka et al. [61] find an almost linear algorithm for checking equivalence of *guarded* KAT terms ($O(n \cdot \alpha(n))$, where α is the inverse Ackermann function), i.e., terms which use if and while instead of $+$ and $*$, respectively. Their guarded KAT is completely abstract, while our KMTs are completely concrete.

Our work is loosely related to Satisfiability Modulo Theories (SMT) [19]. The high-level motivation is the same—to create an extensible framework where custom theories can be combined [47] and used to increase the expressiveness and power [62] of the underlying technique (SAT vs. KA). However, the specifics vary greatly—while SMT is used to reason about the formula satisfiability, KMT is used to reason about how program structure interacts with tests. Some of our KMT theories implement satisfiability checking by calling out to Z3 [18].

The pushback requirement detailed in this paper is closely related to the classical notion of weakest precondition [6, 20, 55]. The pushback operation isn’t quite a generalization of weakest preconditions because the various PB relations can change the program itself. Automatic weakest precondition generation is generally limited in the presence of loops in while-programs. While there has been much work on loop invariant inference [24, 25, 27, 34, 49, 57], the problem remains undecidable in most cases; however, the pushback restrictions of “growth” of terms makes it possible for us to automatically lift the weakest precondition generation to loops in KAT. In fact, this is exactly what the normalization proof does when lifting tests out of the Kleene star operator.

7 CONCLUSION

Kleene algebra modulo theories (KMT) is a new framework for extending Kleene algebra with tests with the addition of actions and predicates in a custom domain. KMT uses an operation that pushes tests back through actions to go from a decidable client theory to a domain-specific KMT. Derived KMTs are sound and complete with respect to a trace semantics; we derive decision procedures for the KMT in an implementation that mirrors our formalism. The KMT framework captures common use cases and can reproduce *by simple composition* several results from the literature, some of which were challenging results in their own right, as well as several new results: we offer theories for bitvectors [30], natural numbers, unbounded sets and maps, networks [1], and temporal logic [8]. Our ability to reason about unbounded state is novel. Our work, however, is limited to tracing semantics; we conjecture that it is possible to merge actions (as in KAT+B!, NetKAT, and Kleene algebra with equations [1, 30, 40]), but leave it to future work.

ACKNOWLEDGMENTS

Dave Walker and Aarti Gupta provided valuable advice. Ryan Beckett was supported by NSF CNS award 1703493. Justin Hsu provided advice and encouragement.

REFERENCES

- [1] Carolyn Jane Anderson, Nate Foster, Arjun Guha, Jean-Baptiste Jeannin, Dexter Kozen, Cole Schlesinger, and David Walker. 2014. NetKAT: Semantic Foundations for Networks. In *Proceedings of the 41st ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages* (San Diego, California, USA) (POPL '14). ACM, New York, NY, USA, 113–126.
- [2] M Anderson. 2014. Time Warner Cable Says Outages Largely Resolved. <http://www.seattletimes.com/business/time-warner-cable-says-outages-largely-resolved>.
- [3] Allegra Angus and Dexter Kozen. 2001. *Kleene Algebra with Tests and Program Schematology*. Technical Report. Cornell University, Ithaca, NY, USA.
- [4] Mina Tahmasbi Arashloo, Yaron Koral, Michael Greenberg, Jennifer Rexford, and David Walker. 2016. SNAP: Stateful Network-Wide Abstractions for Packet Processing. In *Proceedings of the 2016 ACM SIGCOMM Conference* (Florianopolis, Brazil) (SIGCOMM '16). ACM, New York, NY, USA, 29–43.
- [5] Jorge A. Baier and Sheila A. McIlraith. 2006. Planning with First-order Temporally Extended Goals Using Heuristic Search. In *National Conference on Artificial Intelligence* (Boston, Massachusetts) (AAAI'06). AAAI Press, 788–795. <http://dl.acm.org/citation.cfm?id=1597538.1597664>
- [6] Mike Barnett and K. Rustan M. Leino. 2005. Weakest-precondition of Unstructured Programs. In *Proceedings of the 6th ACM SIGPLAN-SIGSOFT Workshop on Program Analysis for Software Tools and Engineering* (Lisbon, Portugal) (PASTE '05). ACM, New York, NY, USA, 82–87.
- [7] Adam Barth and Dexter Kozen. 2002. *Equational verification of cache blocking in lu decomposition using kleene algebra with tests*. Technical Report. Cornell University.
- [8] Ryan Beckett, Michael Greenberg, and David Walker. 2016. Temporal NetKAT. In *Proceedings of the 37th ACM SIGPLAN Conference on Programming Language Design and Implementation* (Santa Barbara, CA, USA) (PLDI '16). ACM, New York, NY, USA, 386–401.
- [9] Janusz A. Brzozowski. 1964. Derivatives of Regular Expressions. *J. ACM* 11, 4 (Oct. 1964), 481–494. <https://doi.org/10.1145/321239.321249>
- [10] Eric Hayden Campbell. 2017. *Infiniteness and Linear Temporal Logic: Soundness, Completeness, and Decidability*. Undergraduate thesis. Pomona College.
- [11] Eric Hayden Campbell and Michael Greenberg. 2018. Injecting finiteness to prove completeness for finite linear temporal logic. In submission.
- [12] Ernie Cohen. 1994. Hypotheses in Kleene Algebra. <http://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.56.6067>
- [13] Ernie Cohen. 1994. Lazy Caching in Kleene Algebra. <http://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.57.5074>
- [14] Ernie Cohen. 1994. *Using Kleene algebra to reason about concurrency control*. Technical Report. Telcordia.
- [15] Anupam Das and Damien Pous. 2017. A Cut-Free Cyclic Proof System for Kleene Algebra. In *Automated Reasoning with Analytic Tableaux and Related Methods*, Renate A. Schmidt and Cláudia Nalon (Eds.). Springer International Publishing, Cham, 261–277.
- [16] Giuseppe De Giacomo, Riccardo De Masellis, and Marco Montali. 2014. Reasoning on LTL on Finite Traces: Insensitivity to Infiniteness.. In *AAAI Citeseer*, 1027–1033.

- 2157 [17] Giuseppe De Giacomo and Moshe Y Vardi. 2013. Linear temporal logic and linear dynamic logic on finite traces.
2158 In *IJCAI'13 Proceedings of the Twenty-Third international joint conference on Artificial Intelligence*. Association for
2159 Computing Machinery, 854–860.
- 2160 [18] Leonardo De Moura and Nikolaj Bjørner. 2008. Z3: An Efficient SMT Solver. In *Proceedings of the Theory and Practice of*
2161 *Software, 14th International Conference on Tools and Algorithms for the Construction and Analysis of Systems* (Budapest,
2162 Hungary) (*TACAS'08/ETAPS'08*). Springer-Verlag, Berlin, Heidelberg, 337–340.
- 2163 [19] Leonardo De Moura and Nikolaj Bjørner. 2011. Satisfiability Modulo Theories: Introduction and Applications. *Commun.*
2164 *ACM* 54, 9 (Sept. 2011), 69–77.
- 2165 [20] Edsger W. Dijkstra. 1975. Guarded Commands, Nondeterminacy and Formal Derivation of Programs. *Commun. ACM*
2166 18, 8 (Aug. 1975), 453–457.
- 2167 [21] Nate Foster, Rob Harrison, Michael J. Freedman, Christopher Monsanto, Jennifer Rexford, Alec Story, and David Walker.
2168 2011. Frenetic: a network programming language. In *Proceeding of the 16th ACM SIGPLAN international conference on*
2169 *Functional Programming, ICFP 2011, Tokyo, Japan, September 19-21, 2011*. 279–291. <https://doi.org/10.1145/2034773.2034812>
- 2170 [22] Nate Foster, Dexter Kozen, Konstantinos Mamouras, Mark Reitblatt, and Alexandra Silva. 2016. Probabilistic NetKAT.
2171 In *Programming Languages and Systems: 25th European Symposium on Programming, ESOP 2016, Held as Part of the*
2172 *European Joint Conferences on Theory and Practice of Software, ETAPS 2016, Eindhoven, The Netherlands, April 2–8, 2016,*
2173 *Proceedings*, Peter Thiemann (Ed.). Springer Berlin Heidelberg, Berlin, Heidelberg, 282–309.
- 2174 [23] Nate Foster, Dexter Kozen, Matthew Milano, Alexandra Silva, and Laure Thompson. 2015. A Coalgebraic Decision Pro-
2175 cedure for NetKAT. In *Proceedings of the 42Nd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming*
2176 *Languages* (Mumbai, India) (*POPL '15*). ACM, New York, NY, USA, 343–355.
- 2177 [24] Carlo A. Furia and Bertrand Meyer. 2009. Inferring Loop Invariants using Postconditions. *CoRR* abs/0909.0884 (2009).
- 2178 [25] Carlo Alberto Furia and Bertrand Meyer. 2010. *Inferring Loop Invariants Using Postconditions*. Springer Berlin Heidelberg,
2179 Berlin, Heidelberg, 277–300.
- 2180 [26] Murdoch J. Gabbay and Vincenzo Ciancia. 2011. Freshness and Name-restriction in Sets of Traces with Names. In
2181 *Proceedings of the 14th International Conference on Foundations of Software Science and Computational Structures: Part of*
2182 *the Joint European Conferences on Theory and Practice of Software* (Saarbrücken, Germany) (*FOSSACS'11/ETAPS'11*).
2183 Berlin, Heidelberg, 365–380.
- 2184 [27] Juan P. Galeotti, Carlo A. Furia, Eva May, Gordon Fraser, and Andreas Zeller. 2014. Automating Full Functional
2185 Verification of Programs with Loops. *CoRR* abs/1407.5286 (2014). <http://arxiv.org/abs/1407.5286>
- 2186 [28] Phillipa Gill, Navendu Jain, and Nachiappan Nagappan. 2011. Understanding Network Failures in Data Centers:
2187 Measurement, Analysis, and Implications. In *SIGCOMM*.
- 2188 [29] Joanne Godfrey. 2016. The Summer of Network Misconfigurations. <https://goo.gl/ViU9uS>.
- 2189 [30] Niels Bjørn Bugge Grathwohl, Dexter Kozen, and Konstantinos Mamouras. 2014. KAT + B!. In *Proceedings of the Joint*
2190 *Meeting of the Twenty-Third EACSL Annual Conference on Computer Science Logic (CSL) and the Twenty-Ninth Annual*
2191 *ACM/IEEE Symposium on Logic in Computer Science (LICS)* (Vienna, Austria) (*CSL-LICS '14*). ACM, New York, NY, USA,
2192 Article 44, 44:1–44:10 pages.
- 2193 [31] Arjun Guha, Mark Reitblatt, and Nate Foster. 2013. Machine-verified network controllers. In *ACM SIGPLAN Conference*
2194 *on Programming Language Design and Implementation, PLDI '13, Seattle, WA, USA, June 16-19, 2013*. 483–494. <https://doi.org/10.1145/2462156.2462178>
- 2195 [32] John E. Hopcroft and R. M. Karp. 1971. *A Linear Algorithm for Testing Equivalence of Finite Automata*. Technical Report
2196 71-114. Cornell University.
- 2197 [33] Zeus Kerravala. 2004. What is Behind Network Downtime? Proactive Steps to Reduce Human Error and Improve
2198 Availability of Networks. <https://www.cs.princeton.edu/courses/archive/fall10/cos561/papers/Yankee04.pdf>.
- 2199 [34] Soonho Kong, Yungbum Jung, Cristina David, Bow-Yaw Wang, and Kwangkeun Yi. 2010. Automatically Inferring
2200 Quantified Loop Invariants by Algorithmic Learning from Simple Templates. In *Proceedings of the 8th Asian Conference*
2201 *on Programming Languages and Systems* (Shanghai, China) (*APLAS'10*). 328–343.
- 2202 [35] Dexter Kozen. 1994. A Completeness Theorem for Kleene Algebras and the Algebra of Regular Events. *Inf. Comput.*
2203 110, 2 (1994), 366–390. <https://doi.org/10.1006/inco.1994.1037>
- 2204 [36] Dexter Kozen. 1997. Kleene Algebra with Tests. *ACM Trans. Program. Lang. Syst.* 19, 3 (May 1997), 427–443. <https://doi.org/10.1145/256167.256195>
- 2205 [37] Dexter Kozen. 2003. *Kleene algebra with tests and the static analysis of programs*. Technical Report. Cornell University.
- [38] Dexter Kozen. 2004. Some results in dynamic model theory. *Science of Computer Programming* 51, 1 (2004), 3 – 22. <https://doi.org/10.1016/j.scico.2003.09.004> Mathematics of Program Construction (MPC 2002).
- [39] Dexter Kozen. 2017. On the Coalgebraic Theory of Kleene Algebra with Tests. In *Rohit Parikh on Logic, Language and Society*. Springer, 279–298.

- 2206 [40] Dexter Kozen and Konstantinos Mamouras. 2014. Kleene Algebra with Equations. In *Automata, Languages, and*
 2207 *Programming: 41st International Colloquium, ICALP 2014, Copenhagen, Denmark, July 8-11, 2014, Proceedings, Part II*,
 2208 Javier Esparza, Pierre Fraigniaud, Thore Husfeldt, and Elias Koutsoupias (Eds.). Springer Berlin Heidelberg, Berlin,
 2209 Heidelberg, 280–292.
- 2210 [41] Dexter Kozen and Maria-Christina Patron. 2000. Certification of Compiler Optimizations Using Kleene Algebra with
 2211 Tests. In *Proceedings of the First International Conference on Computational Logic (CL '00)*. Springer-Verlag, London, UK,
 2212 UK, 568–582.
- 2213 [42] Kim G Larsen, Stefan Schmid, and Bingtian Xue. 2016. WNetKAT: Programming and Verifying Weighted Software-
 2214 Defined Networks. In *OPODIS*.
- 2215 [43] Ratul Mahajan, David Wetherall, and Tom Anderson. 2002. Understanding BGP Misconfiguration. In *SIGCOMM*.
- 2216 [44] Jedidiah McClurg, Hossein Hojjat, Nate Foster, and Pavol Černý. 2016. Event-driven Network Programming. In
 2217 *Proceedings of the 37th ACM SIGPLAN Conference on Programming Language Design and Implementation* (Santa Barbara,
 2218 CA, USA) (*PLDI '16*). ACM, New York, NY, USA, 369–385.
- 2219 [45] Christopher Monsanto, Nate Foster, Rob Harrison, and David Walker. 2012. A compiler and run-time system for network
 2220 programming languages. In *Proceedings of the 39th ACM SIGPLAN-SIGACT Symposium on Principles of Programming*
 2221 *Languages, POPL 2012, Philadelphia, Pennsylvania, USA, January 22-28, 2012*. 217–230. <https://doi.org/10.1145/2103656.2103685>
- 2222 [46] Yoshiki Nakamura. 2015. Decision Methods for Concurrent Kleene Algebra with Tests: Based on Derivative. *RAMiCS*
 2223 *2015* (2015), 1.
- 2224 [47] Greg Nelson and Derek C. Oppen. 1979. Simplification by Cooperating Decision Procedures. *ACM Trans. Program.*
 2225 *Lang. Syst.* 1, 2 (Oct. 1979), 245–257.
- 2226 [48] Juniper Networks. 2008. As the Value of Enterprise Networks Escalates, So Does the Need for Configuration Manage-
 2227 ment. <https://www-935.ibm.com/services/au/gts/pdf/200249.pdf>.
- 2228 [49] Saswat Padhi, Rahul Sharma, and Todd Millstein. 2016. Data-driven Precondition Inference with Learned Features.
 2229 In *Proceedings of the 37th ACM SIGPLAN Conference on Programming Language Design and Implementation* (Santa
 2230 Barbara, CA, USA) (*PLDI '16*). New York, NY, USA, 42–56.
- 2231 [50] Damien Pous. 2015. Symbolic Algorithms for Language Equivalence and Kleene Algebra with Tests. In *Proceedings of*
 2232 *the 42Nd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages* (Mumbai, India) (*POPL*
 2233 *'15*). New York, NY, USA, 357–368.
- 2234 [51] Mark Reitblatt, Marco Canini, Arjun Guha, and Nate Foster. 2013. FatTire: declarative fault tolerance for software-
 2235 defined networks. In *Proceedings of the Second ACM SIGCOMM Workshop on Hot Topics in Software Defined Networking,*
 2236 *HotSDN 2013, The Chinese University of Hong Kong, Hong Kong, China, Friday, August 16, 2013*. 109–114. <https://doi.org/10.1145/2491185.2491187>
- 2237 [52] Yakov Rekhter and Tony Li. 1995. *A Border Gateway Protocol 4 (BGP-4)*. RFC 1654. RFC Editor. 1–56 pages. <http://www.rfc-editor.org/rfc/rfc1654.txt>
- 2238 [53] Grigore Roşu. 2016. Finite-Trace Linear Temporal Logic: Coinductive Completeness. In *International Conference on*
 2239 *Runtime Verification*. Springer, 333–350.
- 2240 [54] J. J.M.M. Rutten. 1996. *Universal Coalgebra: A Theory of Systems*. Technical Report. CWI (Centre for Mathematics and
 2241 Computer Science), Amsterdam, The Netherlands, The Netherlands.
- 2242 [55] Andrew E. Santosa. 2015. Comparing Weakest Precondition and Weakest Liberal Precondition. *CoRR* abs/1512.04013
 2243 (2015).
- 2244 [56] Cole Schlesinger, Michael Greenberg, and David Walker. 2014. Concurrent NetCore: From Policies to Pipelines. In
 2245 *Proceedings of the 19th ACM SIGPLAN International Conference on Functional Programming* (Gothenburg, Sweden)
 2246 (*ICFP '14*). ACM, New York, NY, USA, 11–24.
- 2247 [57] Rahul Sharma and Alex Aiken. 2014. From Invariant Checking to Invariant Inference Using Randomized Search. In
 2248 *Proceedings of the 16th International Conference on Computer Aided Verification - Volume 8559*. New York, NY, USA,
 2249 88–105.
- 2250 [58] Simon Sharwood. 2016. Google cloud wobbles as workers patch wrong routers. http://www.theregister.co.uk/2016/03/01/google_cloud_wobbles_as_workers_patch_wrong_routers/.
- 2251 [59] Alexandra Silva. 2010. *Kleene Coalgebra*. PhD Thesis. University of Minho, Braga, Portugal.
- 2252 [60] Steffen Smolka, Spiridon Eliopoulos, Nate Foster, and Arjun Guha. 2015. A Fast Compiler for NetKAT. In *Proceedings*
 2253 *of the 20th ACM SIGPLAN International Conference on Functional Programming* (Vancouver, BC, Canada) (*ICFP 2015*).
 2254 ACM, New York, NY, USA, 328–341.
- 2255 [61] Steffen Smolka, Nate Foster, Justin Hsu, Tobias Kappé, Dexter Kozen, and Alexandra Silva. 2019. Guarded Kleene
 Algebra with Tests: Verification of Uninterpreted Programs in Nearly Linear Time. *Proc. ACM Program. Lang.* 4, POPL,
 Article 61 (Dec. 2019), 28 pages. <https://doi.org/10.1145/3371129>

- 2255 [62] Aaron Stump, Clark W. Barrett, David L. Dill, and Jeremy R. Levitt. 2001. A Decision Procedure for an Extensional
2256 Theory of Arrays. In *LICS*.
- 2257 [63] Yevgeny Sverdlik. 2012. Microsoft: misconfigured network device led to Azure outage. <https://goo.gl/Se7DzD>
- 2258
- 2259
- 2260
- 2261
- 2262
- 2263
- 2264
- 2265
- 2266
- 2267
- 2268
- 2269
- 2270
- 2271
- 2272
- 2273
- 2274
- 2275
- 2276
- 2277
- 2278
- 2279
- 2280
- 2281
- 2282
- 2283
- 2284
- 2285
- 2286
- 2287
- 2288
- 2289
- 2290
- 2291
- 2292
- 2293
- 2294
- 2295
- 2296
- 2297
- 2298
- 2299
- 2300
- 2301
- 2302
- 2303