Proof Theory for Kleene Algebra

Chris Hardin
Department of Mathematics
Cornell University
Ithaca, NY 14853-4201, USA
hardin@math.cornell.edu

Abstract

*The universal Horn theory of relational Kleene algebra with tests (*RKAT*) is of practical interest, particularly for program semantics. We develop an (infinitary) proof system, based on well-founded trees of finite automata, which is sound and complete for this theory. A small modification of this system yields a proof system which is sound and complete for the universal Horn theory of $*$-continuous Kleene algebras with tests (*KAT^**). This sheds light on the relationship between* RKAT *and* KAT^*.

1 Introduction

Kleene algebra (KA) arises in many areas of computer science, such as automata theory, the design and analysis of algorithms, dynamic logic, and program semantics. Many of these applications are enhanced by using Kleene algebra with tests (KAT), which combines KA with Boolean algebra. The use of KAT in program verification largely motivates our work here.

We can use KAT to reason propositionally about programs (see [1, 12] for examples). The equivalence of an optimized and unoptimized program, the equivalence of an annotated and unannotated program, and partial correctness assertions can all be expressed as equations. The equational theory of KAT is well understood and has many useful properties; in particular, it is decidable (in *PSPACE*) and the theory remains unchanged when we restrict to relational interpretations [3, 13]. (Relational interpretations are of the greatest interest because the intended semantics are generally relational.)

However, we frequently wish to reason about programs under certain assumptions about the interaction of atomic programs and tests. For example, if p is the program "x := 3" and b is the assertion "x = 3", then we want to be able to make use of the facts $pb = p$ and $bp = b$ when reasoning about programs in which p and b appear; for instance, the equation $p^2 = p$ is not valid in KAT, but the formula $(pb = p \ \wedge \ bp = b) \rightarrow p^2 = p$ is. Thus, the *universal Horn theory* is of interest. A *univeral Horn formula* is an implication $E \rightarrow s = t$, where E is a finite set of equations. The word "universal" refers to the fact that the atomic symbols of E, s, and t are implicitly universally quantified. The *universal Horn theory* of a class of structures C, denoted $\mathcal{H}\mathsf{C}$, is the set of universal Horn formulas valid under all interpretations over structures in C.

The increased generality of the universal Horn theory is accompanied by greater complexity, and the theory does not remain the same when we restrict to important classes of Kleene algebras such as $*$-continuous Kleene algebras with tests (KAT^*) and relational Kleene algebras with tests (RKAT). $\mathcal{H}\mathsf{KAT}$ is Σ^0_1-complete (undecidable), $\mathcal{H}\mathsf{KAT}^*$ and $\mathcal{H}\mathsf{RKAT}$ are Π^1_1-complete (highly undecidable), and we have proper inclusions $\mathcal{H}\mathsf{KAT} \subsetneq \mathcal{H}\mathsf{KAT}^* \subsetneq \mathcal{H}\mathsf{RKAT}$ [11, 7].

Although the complexity makes it impossible to thoroughly understand these theories, their study yields practical results. For example, if we restrict premises to the form $s = 0$, this portion of $\mathcal{H}\mathsf{KAT}$ is *PSPACE*-complete, subsumes Hoare logic, and is complete for relational interpretations [10]. Further work has shown that we can also admit premises of the form $bp = b$ while remaining in *PSPACE* (and preserving the completeness of KAT over relational interpretations) [6]; premises of this form are useful for eliminating redundant code.

This paper contributes to our understanding of these universal Horn theories by developing sound and complete proof systems for $\mathcal{H}\mathsf{KAT}^*$ and $\mathcal{H}\mathsf{RKAT}$. Built around well-founded trees of finite automata, these systems are closely related to each other, and their difference sheds light on the relationship between $\mathcal{H}\mathsf{KAT}^*$ and $\mathcal{H}\mathsf{RKAT}$. Although the systems are infinitary (a necessity given the Π^1_1-completeness of $\mathcal{H}\mathsf{KAT}^*$ and $\mathcal{H}\mathsf{RKAT}$), many formulas do have finite proofs, or can be seen by inspection to not be provable; this has provided many new examples of formulas that are in $\mathcal{H}\mathsf{RKAT}$ but not in $\mathcal{H}\mathsf{KAT}^*$. However, this is somewhat beside the point: the systems' foremost purpose is not to prove individual formulas, but rather to provide a platform for proof-theoretic arguments yielding general theorems about the theories $\mathcal{H}\mathsf{KAT}$, $\mathcal{H}\mathsf{KAT}^*$, and $\mathcal{H}\mathsf{RKAT}$. For example, we can give syntactical criteria for Horn formulas that guarantee that when such formulas are provable at all (in either system), they are provable via a finite proof; this immediately reduces the complexity of $\mathcal{H}\mathsf{KAT}^*$ and $\mathcal{H}\mathsf{RKAT}$, restricted to such formulas, from Π^1_1 to Σ^0_1. The power of such proof-theoretic techniques becomes clear when one observes that they yield short proofs of many existing theorems whose known proofs are much longer.

The greatest significance of these proof systems is their utility in studying $\mathcal{H}\mathsf{RKAT}$. Despite its complexity, $\mathcal{H}\mathsf{KAT}^*$ is at least reasonably transparent, in the sense that we have an axiomatization for KAT^* (that makes use of an infinitary Horn formula to capture $*$-continuity). $\mathcal{H}\mathsf{RKAT}$ is very murky by comparison; our current knowledge of it has been gained slowly, most often through ad hoc constructions. The proof systems presented here provide a significant tool for its study.

2 Preliminaries

For a more complete introduction to Kleene algebra and Kleene algebra with tests, see [9].

2.1 Kleene Algebra

Definition 1 *An* idempotent semiring *is a structure* $(S, +, \cdot, 0, 1)$ *satisfying*

$$\begin{aligned} x + x &= x \quad \textit{(idempotence)} \\ x + 0 &= x \\ x + y &= y + x \\ x + (y + z) &= (x + y) + z \\ 0 \cdot x &= x \cdot 0 = 0 \\ 1 \cdot x &= x \cdot 1 = x \\ x \cdot (y \cdot z) &= (x \cdot y) \cdot z \\ x \cdot (y + z) &= x \cdot y + x \cdot z \\ (y + z) \cdot x &= y \cdot x + z \cdot x \end{aligned}$$

(In other words, $(S, +, 0)$ *is an upper semilattice with bottom element* 0*,* $(S, \cdot, 1)$ *is a monoid,* 0 *is an annihilator for* $\cdot$*, and* $\cdot$ *distributes over* $+$ *on the right and left.)*

We often drop $\cdot$, writing xy for $x \cdot y$. The upper semilattice structure induces a natural partial order on any idempotent semiring: $x \le y \Leftrightarrow x + y = y$. (Note also that this means inequalities $s \le t$ may appear anywhere equations appear, such as in universal Horn formulas.) $+$ and $\cdot$ enjoy the following form of monotonicity: if $x \le x'$ and $y \le y'$, then $x + y \le x' + y'$, and $xy \le x'y'$. (For $+$, this is trivial. For $\cdot$, suppose $x \le x'$ and $y \le y'$. Then $x + x' = x'$, so we have $xy + x'y = (x + x')y = x'y$, so $xy \le x'y$. We similarly have $x'y \le x'y'$, so $xy \le x'y'$.)

Definition 2 *A* Kleene algebra *is a structure* $(K, +, \cdot, ^*, 0, 1)$ *such that* $(K, +, \cdot, 0, 1)$ *forms an idempotent semiring, and which satisfies*

$$1 + xx^* \le x^* \quad (1)$$
$$1 + x^*x \le x^* \quad (2)$$
$$p + qx \le x \;\rightarrow\; q^*p \le x \quad (3)$$
$$p + xq \le x \;\rightarrow\; pq^* \le x \quad (4)$$

(The order of precedence among the operators is $^* > \cdot > +$*, so that* $p + qr^* = p + (q \cdot (r^*))$*.) We let* KA *denote the category of all Kleene algebras and their homomorphisms.*

(The names of the categories we consider will also serve as convenient abbreviations for the type of algebra they contain. For example, "the KA *axioms" will mean "the axioms of Kleene algebra".)*

Equation (1) implies that q^*p *is a solution to the inequality* $p + qx \le x$*, and (3) implies that it is the least solution; (2) and (4) say that* pq^* *is the least solution to* $p + xq \le x$*.*

Given a set Σ *of constant symbols, let* RExp_Σ *be the set of Kleene algebra terms over* Σ*. We call the elements of* RExp_Σ regular expressions*, and the elements of* Σ atomic program symbols*. An* interpretation *is a homomorphism* $I : \mathsf{RExp}_\Sigma \to K$*, where* K *is a Kleene algebra.* I *is determined uniquely by its values on* Σ*.*

We use $\models$ to denote ordinary Tarskian satisfaction. However, since we have constant symbols from Σ not in the signatures of the underlying algebras, we will pair each algebra with an interpretation when speaking about satisfaction. For example, given a Kleene algebra K, interpretation $I : \mathsf{RExp}_\Sigma \to K$, and formula φ whose atomic program symbols are among Σ, we will write $K, I \models \varphi$ to indicate that K satisfies φ when the symbols in Σ are evaluated according to I. $K \models \varphi$ means that $K, I \models \varphi$ for every interpretation $I : \mathsf{RExp}_\Sigma \to K$. We also use $\models$ in two other standard ways: for a class C of algebras, $\mathsf{C} \models \varphi$ means that $K \models \varphi$ for each $K \in \mathsf{C}$; for a set Φ of formulas, $\Phi \models \varphi$ means that $K \models \varphi$ for each algebra K satisfying every formula in Φ.

Definition 3 *For an arbitrary monoid* M*, its powerset* 2^M

forms a Kleene algebra as follows.

$$\begin{array}{rcl} 0 & = & \emptyset \\ 1 & = & \{1^M\} \quad \textit{(where } 1^M \textit{ is the identity of } M\textit{)} \\ A+B & = & A \cup B \\ A \cdot B & = & \{xy \mid x \in A,\ y \in B\} \\ A^* & = & \bigcup_{k \in \omega} A^k \end{array}$$

We let REG M *denote the smallest subalgebra of* 2^M *containing the singletons* $\{x\}$, $x \in M$. *(The elements of* REG M *are the* regular subsets *of* M.*)* 2^M *and its subalgebras are known as* language algebras.

Of particular interest is the case $M = \Sigma^*$*, the monoid of all strings over alphabet* Σ *under concatenation. The empty string* ε *is the identity of this monoid. We define the canonical interpretation* $R : \mathsf{RExp}_\Sigma \to \mathsf{REG}\ \Sigma^*$ *by letting* $R(p) = \{p\}$ *(and extending* R *homomorphically to the rest of* RExp_Σ*). Note that we can interpret elements of* Σ^* *as elements of* RExp_Σ *in the obvious fashion.*

Relational Kleene algebras are also of interest.

Definition 4 *For an arbitrary set* X*, the set* $2^{X \times X}$ *of all binary relations on* X *forms a Kleene algebra* $\mathcal{R}(X)$ *as follows.*

$$\begin{array}{rcl} 0 & = & \emptyset \\ 1 & = & \iota_X = \{(x,x) \mid x \in X\} \\ S+T & = & S \cup T \\ S \cdot T & = & S \circ T \quad \textit{(the composition of } S \textit{ with } T\textit{)} \\ S^* & = & \bigcup_{k \in \omega} S^k \quad \textit{(the reflexive transitive closure of } S\textit{)} \end{array}$$

A Kleene algebra K *is* relational *if it is a subalgebra of* $\mathcal{R}(X)$ *for some* X*;* X *is called the* base *of* K*. We let* RKA *denote the category of all relational Kleene algebras and their homomorphisms.*

The definitions of $*$ in 2^M and $\mathcal{R}(X)$ exemplify the most common intuition about the meaning of $*$, which is that $y^* = \sup_{n \in \omega} y^n$, or informally, $y^* = 1 + y + y^2 + \cdots$. (More generally, if we require that multiplication distributes over this supremum, we have $xy^*z = x1z + xyz + xy^2z + \cdots = \sup_{n \in \omega} xy^nz$.) However, this property of $*$ does not follow from the KA $*$-axioms, and must be postulated separately.

Definition 5 *A Kleene algebra* K *is* $*$-continuous *if it satisfies*

$$xy^*z = \sup_{k \in \omega} xy^kz$$

for all $x, y, z \in K$*. We let* KA^* *denote the category of all* $*$*-continuous Kleene algebras and their homomorphisms.*

Since relational composition distributes over arbitrary union, it is immediate from the definition of * in $\mathcal{R}(X)$ that relational Kleene algebras are $*$-continuous, so $\mathsf{RKA} \subseteq \mathsf{KA}^*$.

The following ubiquitous lemma is a useful generalization of $*$-continuity.

Lemma 6 *Suppose* $K \in \mathsf{KA}^*$*,* $I : \mathsf{RExp}_\Sigma \to K$ *is an interpretation, and* $t \in \mathsf{RExp}_\Sigma$*. Then*

$$I(t) = \sup_{\sigma \in R(t)} I(\sigma) \ .$$

Proof By induction on structure of t. For details, see [8, Lemma 7.1, pp. 246–248]. □

Corollary 7 *Suppose* $K \in \mathsf{RKA}$*,* $I : \mathsf{RExp}_\Sigma \to K$ *is an interpretation, and* $t \in \mathsf{RExp}_\Sigma$*. Then*

$$I(t) = \bigcup_{\sigma \in R(t)} I(\sigma) \ .$$

In particular, if $(x,y) \in I(t)$*, then there exists* $\sigma \in R(t)$ *such that* $(x,y) \in I(\sigma)$*.*

Proof Let X be the base of K. Union coincides with supremum in $\mathcal{R}(X)$, and we can treat I as an interpretation into $\mathcal{R}(X)$; applying Lemma 6 to $\mathcal{R}(X)$ and I gives the desired result. (This trick was necessary to get around the fact that in an arbitrary RKA such as K, suprema do not always coincide with union, although, as this lemma shows, suprema over regular sets must coincide with union.) □

Definition 8 *A* universal Horn formula *is a formula of the form*

$$s_1 = t_1 \wedge \cdots \wedge s_t = t_k \to s = t \ ,$$

where s_i, t_i, s, t *are terms in the appropriate language. The set of universal Horn formulas valid over a class* C *of algebras is the* universal Horn theory *of* C*, which we denote by* $\mathcal{H}\mathsf{C}$*. We will often drop the word "universal".*

Lemma 9 *Let* Γ *be any class of* $*$*-continuous Kleene algebras with interpretations. (That is,* Γ *consists of pairs* (K, I) *where* $K \in \mathsf{KA}^*$ *and* $I : \mathsf{RExp}_\Sigma \to K$ *is an interpretation.) Then for any Horn formula of the form* $E \to s \le t$*,*

$$\Gamma \models E \to s \le t \iff (\forall \sigma \in R(s))\ \Gamma \models E \to \sigma \le t \ .$$

Proof For any $K \in \mathsf{KA}^*$ with interpretation $I : \mathsf{RExp}_\Sigma \to K$, the equivalence

$$K, I \models E \to s \le t \iff (\forall \sigma \in R(s))\ K, I \models E \to \sigma \le t$$

is a straightforward consequence of Lemma 6. The lemma then follows by exchanging the universal quantifiers $(\forall \sigma \in R(s))$ and $(\forall (K, I) \in \Gamma)$. (This latter quantifier comes from $\Gamma \models E \to s \le t \Leftrightarrow (\forall (K, I) \in \Gamma)\ K, I \models E \to s \le t$.) □

2.2 Kleene Algebra with Tests

Definition 10 *A* Kleene algebra with tests *is a two-sorted structure* $(K, B, +, \cdot, ^*, \overline{\ }, 0, 1)$, *where* $(K, +, \cdot, ^*, 0, 1)$ *is a Kleene algebra, and* $(B, +, \cdot, \overline{\ }, 0, 1)$ *is a Boolean subalgebra. The elements of* B *are called* tests. *We let* KAT *denote the category of all Kleene algebras with tests and their homomorphisms; we let* KAT^* *denote the subcategory of all* $*$*-continuous Kleene algebras with tests.*

We now have two types of atomic symbols: programs and tests. For a finite set P *of atomic program symbols and a finite set* B *of atomic test symbols,* $\mathsf{RExp}_{\mathsf{P,B}}$ *is the set of* KAT *terms over* P *and* B*; negation can only be applied to Boolean terms, which are terms built from* $0, 1, +, \cdot, \overline{\ }$*, and atomic test symbols. An interpretation* $I : \mathsf{RExp}_{\mathsf{P,B}} \to K$ *must map each atomic test to a test in* K *(and it follows by induction that it will map all Boolean terms to tests).*

$\mathcal{R}(X)$ *forms a Kleene algebra with tests by keeping the previously defined Kleene algebra structure, and letting* $B = \{r \in \mathcal{R}(X) \mid r \leq 1\}$, $\overline{b} = \iota_X - b$. *A Kleene algebra with tests* K *is relational if it is a subalgebra of* $\mathcal{R}(X)$ *for some* X. *We let* RKAT *denote the category of all relational Kleene algebras with tests and their homomorphisms.*

Every Kleene algebra induces a Kleene algebra with tests by letting $B = \{0, 1\}$, the two-element Boolean algebra; conversely, every Kleene algebra with tests induces a Kleene algebra by taking its reduct to the signature of Kleene algebra (*i.e.*, taking its image under the map $(K, B, +, \cdot, ^*, \overline{\ }, 0, 1) \mapsto (K, +, \cdot, ^*, 0, 1)$). With this in mind, it is easy to see that for any formula φ in the language of Kleene algebra, $\mathsf{KAT} \models \varphi \Leftrightarrow \mathsf{KA} \models \varphi$, $\mathsf{KAT}^* \models \varphi \Leftrightarrow \mathsf{KA}^* \models \varphi$, and $\mathsf{RKAT} \models \varphi \Leftrightarrow \mathsf{RKA} \models \varphi$.

3 A Proof System for $\mathcal{H}\mathsf{RKA}$

3.1 Preview

Notation: Given binary relations S, T over some set X, xSy will mean $(x, y) \in S$. $xSzTy$ will be shorthand for xSz & zTy. ST will denote the composition

$$S \circ T = \{(x, y) \mid \exists z\, xSzTy\}.$$

The most direct way to prove that a Horn formula $E \to s \leq t$ is relationally valid is to suppose $K \in \mathsf{RKA}$, with interpretation I, such that $K, I \models E$, and show that for any $(x, y) \in I(s)$, we must have $(x, y) \in I(t)$. Let us give a concrete example of such a proof.

Let φ be the Horn formula

$$r_0 \leq r_1 + r_2 \wedge s_0 \leq s_1 + s_2 \to r_0 s_0 \leq r_0 s_1 + r_1 s_2 + r_2 s_0.$$

We wish to show $\mathsf{RKA} \models \varphi$. Suppose K is a relational model, with interpretation I, such that $K, I \models r_0 \leq r_1 + r_2 \wedge s_0 \leq s_1 + s_2$. Let $R_i = I(r_i)$, $S_i = I(s_i)$, for $0 \leq i \leq 2$.

Suppose xR_0S_0y. Then xR_0zS_0y for some z.

Applying the hypothesis $r_0 \leq r_1 + r_2$ gives us xR_1z or xR_2z. We break into cases 1 and 2.

Case 1: xR_1z. Applying the hypothesis $s_0 \leq s_1 + s_2$ gives us zS_1y or zS_2y. We break into subcases 1.1 and 1.2.

Case 1.1: zS_1y. We have xR_0zS_1y, so $(x, y) \in I(r_0s_1 + r_1s_2 + r_2s_0)$.

Case 1.2: zS_2y. We have xR_1zS_2y, so $(x, y) \in I(r_0s_1 + r_1s_2 + r_2s_0)$.

Case 2: xR_2z. We have xR_2zS_0y, so $(x, y) \in I(r_0s_1 + r_1s_2 + r_2s_0)$.

Therefore, $K, I \models r_0s_0 \leq r_0s_1 + r_1s_2 + r_2s_0$, so φ is relationally valid.

The structure of the cases forms a tree as in Figure 1.

To represent the information about our relations more compactly, we can let each point be a vertex in a graph, with an edge from v to w labeled p indicating that $(v, w) \in I(p)$:

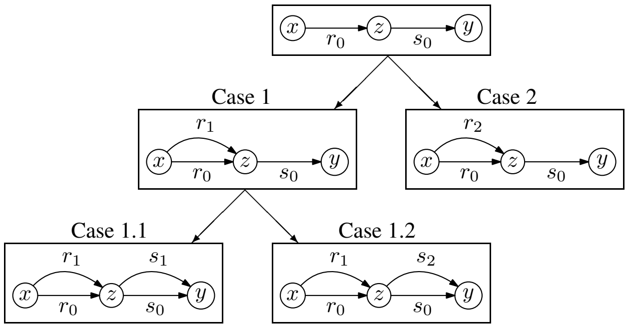

The actual variable names x, y, z are irrelevant, except that we must distinguish x and y in some way, as (x, y) is the pair that must be shown to be in $I(r_0s_1 + r_1s_2 + r_2s_0)$. So we will associate the leftmost vertex with x and the rightmost vertex with y:

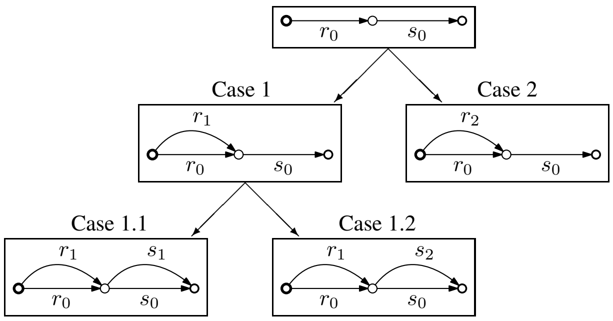

What we now have is a tree of finite automata. In a moment, we will generalize this idea into a proof system for the Horn theory of RKA in which proofs are well-founded trees of automata. As to why this is natural, consider: a relational model is basically a Kripke frame, a countable Kripke frame can be constructed as a sequence of finite Kripke frames

Suppose xR_0zS_0y.

Case 1: xR_1z		Case 2: xR_2z
Case 1.1: yS_1z	Case 1.2: yS_2z	$\therefore (x,y) \in I(r_0s_1 + r_1s_2 + r_2s_0)$
$\therefore (x,y) \in I(r_0s_1 + r_1s_2 + r_2s_0)$	$\therefore (x,y) \in I(r_0s_1 + r_1s_2 + r_2s_0)$	

Figure 1. A tree of cases.

(each one extending the previous one), and a finite Kripke frame is basically a finite automaton. So, the process of constructing a counterexample to a given Horn formula φ can be viewed as a tree of finite automata (partial constructions), in which each path is a particular way of proceeding with the construction—*i.e.*, branching represents different ways of proceeding—such that an infinite path through this tree would give us a model in which φ fails, while the well-foundedness of this tree would imply the relational validity of φ.

3.2 Finite Automata and Trees

Our proof systems for $\mathcal{H}\mathsf{RKA}$ and $\mathcal{H}\mathsf{KA}^*$ will be based on trees of finite automata, and we must define a number of notions related to trees and automata before continuing.

Assume we have a fixed finite alphabet Σ. We let NFA denote the set of all nondeterministic finite automata over Σ, requiring that all states are natural numbers; ε-moves (also called ε-edges) are allowed.

We will also use NFA as shorthand for *nondeterministic finite automaton*. For any NFA A, $L(A)$ denotes the language of A, and $|A|$ denotes the states of A. For states $v, w \in |A|$, let $A^{v,w}$ denote the NFA which is identical to A except that it has v and w as its unique start and accept states, respectively. We fix distinct natural numbers a and b (say, $a = 0$ and $b = 1$), and let $\mathrm{NFA}^{a,b}$ be the set of all $A \in \mathrm{NFA}$ which have unique start state a and unique accept state b. For the rest of Section 3, we will focus on automata in $\mathrm{NFA}^{a,b}$, and in all diagrams of automata, the leftmost state will be the start state, and the rightmost state will be the accept state.

Given $A, B \in \mathrm{NFA}^{a,b}$, we define their *wedge product* $A \wedge B \in \mathrm{NFA}^{a,b}$ as follows. $A \wedge B$ consists of disjoint copies of A and B, except that the start states of A and B are combined into a single start state, and the accept states of A and B are combined into a single accept state. (So $A \wedge B$ is essentially the result of gluing A and B together, in parallel, at their start and end states.) We assume that any renumbering of states needed for the disjoint union occurs only in B; apart from this bias in numbering states, the wedge product is commutative. It is also associative.

We now define some special automata in $\mathrm{NFA}^{a,b}$.

F_0 has states $\{a, b\}$ and no edges.

F_1 has states $\{a, b\}$ and an ε-edge from a to b.

F_2 has states $\{a, b\}$ and ε-edges from a to b and from b to a.

For $\sigma = p_1 \cdots p_k \in \Sigma^*$, we define F_σ so that $L(F_\sigma) = \{\sigma\}$ as follows. For the case $\sigma = \varepsilon$, we let $F_\sigma = F_\varepsilon = F_2$; otherwise, we let F_σ have states $\{a, x_1, x_2, \ldots, x_{k-1}, b\}$ and edges

$$a \xrightarrow{p_1} x_1 \xrightarrow{p_2} \cdots \xrightarrow{p_{k-1}} x_{k-1} \xrightarrow{p_k} b \ .$$

We will sometimes want to "insert" strings into automata. For $A \in \mathrm{NFA}^{a,b}$, $v, w \in |A|$, and $\tau \in \Sigma^*$, we define

$$\mathrm{insert}(A, v, w, \tau) = (A^{v,w} \wedge F_\tau)^{a,b} \ .$$

This has the effect of inserting the string τ into A from v to w, with the exception of $\mathrm{insert}(A, v, w, \varepsilon)$, which also inserts an ε-edge from w to v. We refer to this extra ε-edge as a *reverse ε-edge*. Where it is used, $\mathrm{insert}(A, v, w, \varepsilon)$ will correspond to identifying v and w with each other. (In Section 4, we will have a different notion of insertion, insert_F.)

We now move on to trees. ω is the set of naturals $\{0, 1, \ldots\}$. $\omega^{<\omega}$ is the set of all finite strings of naturals (including the empty string). A set $T \subseteq \omega^{<\omega}$ is a *tree* if it is closed under taking initial segments. A function $f : \omega \to \omega$ can be treated as an infinite sequence of naturals, and for $n \in \omega$, we let $f \restriction n$ denote the initial segment of f of length n. Such an f is a *path* through a tree T if $(f \restriction n) \in T$ for all $n \in \omega$.

3.3 Relational Proofs

We first define our system in a special case, and then generalize the definition.

Definition 11 *Let $E \to \sigma \leq t$ be a Horn formula in the language of* KA *with $\sigma \in \Sigma^*$ and $t \in \mathsf{RExp}_\Sigma$. We assume that all hypotheses in E are inequalities $x \leq y$, by breaking any equations $x = y$ into $x \leq y \wedge y \leq x$ as necessary. We fix a special symbol* CON*, which will signify contradiction.*

A relational tree *for $E \to \sigma \le t$ is a pair (T, A) where $T \subseteq \omega^{<\omega}$ is a tree and $A : T \to \mathrm{NFA}^{a,b} \cup \{\mathrm{CON}\}$ such that the following conditions hold. (A_f will denote $A(f)$.)*

1. *At the root, we have $A_{\langle\rangle} = F_\sigma$*
2. *$f \in T$ is a leaf node if and only if $A_f = \mathrm{CON}$ or $R(t) \cap L(A_f) \neq \emptyset$.*
3. *If f is not a leaf node, then there exist $v, w \in |A_f|$ (possibly equal), $\rho \in L(A_f^{v,w})$, and an inequality $r \le r'$ in E such that $\rho \in R(r)$ and*
   - *(a) if $R(r') = \emptyset$ (typically because $r' = 0$), then f has one child g, with $A_g = \mathrm{CON}$;*
   - *(b) if $R(r') \neq \emptyset$, then f has one child g_τ for each $\tau \in R(r')$, with $A_{g_\tau} = \mathrm{insert}(A_f, v, w, \tau)$.*

   (We say that the hypothesis $r \le r'$ is applied at *f.)*

A relational proof *of $E \to \sigma \le t$ is a well-founded relational tree for $E \to \sigma \le t$. We say $E \to \sigma \le t$ is* relationally provable *if such a proof exists.*

Although our proof of soundness and completeness will stand on its own, Definition 11 is best understood in terms of the reasoning done in Section 3.1. The root node corresponds to the supposition that $(x, y) \in I(\sigma)$. The CON nodes indicate that we have followed a subcase that has $(x, y) \in I(0)$, a contradiction; the other type of leaf node indicates that, within the current subcase, we have $(x, y) \in I(t)$ as desired. The intermediate steps are essentially just a matter of bookkeeping. The well-foundedness of the tree expresses the fact that the cases are exhaustive.

We can reinterpret the tree of automata in Section 3.1 as a relational proof. We now give two further examples that illustrate the purpose of CON and the ε-edges.

Consider the formula $p \le 0 \to p \le q$. If we follow the same reasoning as in the motivating example, the assumption that $(x, y) \in I(p)$ will lead to a contradiction if we apply the hypothesis $p \le 0$. As a relational proof, we would have a tree with F_p at the root, and one child, CON.

Now consider the formula $p \le 1 \to p \le p^2$. The assumption that $(x, y) \in I(p)$ lets us conclude that $(x, y) \in I(1)$ by applying the hypothesis $p \le 1$. Since $I(1)$ is the identity relation, we can conclude that $x = y$. Letting $P = I(p)$ and $E = I(1)$, we now have $xPyExPy$, so $(x, y) \in PEP = I(p)I(1)I(p) = I(p^2)$. Note that this argument required the reflexivity of E: from xEy, we needed to conclude yEx to finish the argument. This is essentially why, in relational proofs, ε-edges are added in both directions. In terms of relational proofs, we start with F_p at the root. We then apply the hypothesis $p \le 1$ by adding ε-edges in both directions between a and b, yielding one child, $F_p \wedge F_2$. We have $p^2 \in L(F_p \wedge F_2)$ (we can follow the reverse ε-edge from b to a, just as we used yEx above), so this is a leaf and we are done.

Note that the above two examples had no branching. This is because the right hand side t of each hypothesis had no more than one element in $R(t)$. Branching can occur only when $|R(t)| > 1$, and infinite branching can occur only when $R(t)$ is infinite; note that $R(t)$ is finite whenever t is $*$-free.

Lemma 12 *For any Horn formula of the form $E \to \sigma \le t$, the following are equivalent.*

(i) $\mathsf{RKA} \models E \to \sigma \le t$

(ii) *$E \to \sigma \le t$ is relationally provable.*

Proof A detailed proof can be found in Section A.1 of the appendix, but we give a rough sketch here.

The completeness is proven by defining a canonical relational tree for the formula, with the property that hypotheses are applied in a fair manner, so that along any path through the tree, any applicable hypothesis is eventually applied. If $E \to \sigma \le t$ is not relationally provable, then this tree is not well-founded, and a path through it can be used to produce a model witnessing $\mathsf{RKA} \not\models E \to \sigma \le t$.

For soundness, we fix a relational proof of $E \to \sigma \le t$. Suppose $K \in \mathsf{RKA}$ with interpretation $I : \mathsf{RExp}_\Sigma \to K$ such that $K, I \models E$, and suppose that $(x, y) \in I(\sigma)$. Thinking of the relational proof as a table of subcases organized into a tree, we start at the root, and iterate the process of choosing the subcase appropriate for K, I, x, and y; because the tree is well founded, we must eventually hit a leaf. This leaf cannot be a CON (since that would yield a contradiction), so instead we reach the other type of leaf, which will witness $(x, y) \in I(t)$. This establishes $I(\sigma) \subseteq I(t)$. The details are essentially a verification that Definition 11 accurately captures the kind of reasoning seen in Section 3.1. □

We now extend the notion of relational provability to arbitrary Horn formulas.

Definition 13 *We say that $E \to s \le t$ is* relationally provable *if $E \to \sigma \le t$ is* relationally provable *for all $\sigma \in R(s)$. We say that $E \to s = t$ is* relationally provable *if both $E \to s \le t$ and $E \to t \le s$ are relationally provable.*

Theorem 14 *The following are equivalent for any Horn formula $E \to s = t$.*

(i) $\mathsf{RKA} \models E \to s = t$

(ii) $E \to s = t$ is relationally provable.

Proof It suffices to show that equivalence holds for the formulas $E \to s \le t$ and $E \to t \le s$. The former equivalence is as follows.

$$\begin{aligned} & \mathsf{RKA} \models E \to s \le t \\ \Leftrightarrow\quad & (\forall \sigma \in R(s))\ \ \mathsf{RKA} \models E \to \sigma \le t \quad \text{(Lemma 9)} \\ \Leftrightarrow\quad & (\forall \sigma \in R(s))\ \ E \to \sigma \le t \text{ is relationally provable} \\ \Leftrightarrow\quad & E \to s \le t \text{ is relationally provable} \end{aligned}$$

The equivalence for $E \to t \le s$ is similar. □

4 A Proof System (or Two) for $\mathcal{H}\mathsf{KA}^*$

4.1 Star Proofs

We can get a proof system for $\mathcal{H}\mathsf{KA}^*$ by taking our proof system for $\mathcal{H}\mathsf{RKA}$ and replacing insert with an appropriately modified notion of insertion, insert_F. The critical difference will be that, while insert allows automata to develop internal structure (such as cycles and intersecting paths), insert_F will prevent any such structure from developing.

Definition 15 *Given $A, B \in \mathrm{NFA}$, we define $A \vee B$ to be the disjoint union of A and B (without the gluing that occurred with $\wedge$).*

We let NFA^F denote the set of all automata of the form

$$F_{\sigma_1} \vee \cdots \vee F_{\sigma_k}$$

where $\sigma_1, \ldots, \sigma_k \in \Sigma^*$. (This includes the case $k = 0$, which yields the empty automaton.) Automata in NFA^F are called free automata *(for their lack of structure).*

Given $A \in \mathrm{NFA}^F$ and $y \in |A|$, there is a unique start state x and a unique accept state z such that y lies on a path from x to z. $L(A^{x,y})$ and $L(A^{y,z})$ are both singletons. We let the prefix *of y in A be the unique $\rho \in L(A^{x,y})$, and let the* suffix *of y in A be the unique $\rho' \in L(A^{y,z})$.*

Given $A \in \mathrm{NFA}^F$, $v, w \in |A|$, and $\tau \in \Sigma^*$, we define

$$\mathrm{insert}_F(A, v, w, \tau) = A \vee F_{\rho\tau\rho'} \ ,$$

where ρ is the prefix of v and ρ' is the suffix of w in A.

The intuitive difference between insert and insert_F is that while $\mathrm{insert}(A, v, w, \tau)$ inserts the new string "in place", $\mathrm{insert}_F(A, v, w, \tau)$ copies a small portion of A (namely, F_ρ and $F_{\rho'}$) off to the side first and inserts τ there.

The following definition of star trees and star proofs is the result of replacing insert with insert_F in Definition 11. A side effect of this change is that the automata in the tree now happen to be NFA^F instead of $\mathrm{NFA}^{a,b}$.

Definition 16 *Let $E \to \sigma \le t$ be a Horn formula in the language of KA with $\sigma \in \Sigma^*$ and $t \in \mathsf{RExp}_\Sigma$. We assume that all hypotheses in E are inequalities $x \le y$, by breaking any equations $x = y$ into $x \le y \wedge y \le x$ as necessary. We fix a special symbol CON, which will signify contradiction.*

A star tree *for $E \to \sigma \le t$ is a pair (T, A) where $T \subseteq \omega^{<\omega}$ is a tree and $A : T \to \mathrm{NFA}^F \cup \{\mathrm{CON}\}$ such that the following conditions hold. (A_f will denote $A(f)$.)*

1. *At the root, we have $A_{\langle\rangle} = F_\sigma$.*
2. *$f \in T$ is a leaf node if and only if $A_f = \mathrm{CON}$ or $R(t) \cap L(A_f) \neq \emptyset$.*
3. *If f is not a leaf node, then there exist $v, w \in |A_f|$ (possibly equal), $\rho \in L(A_f^{v,w})$, and an inequality $r \le r'$ in E such that $\rho \in R(r)$ and*
   (a) *if $R(r') = \emptyset$, then f has one child g, with $A_g = \mathrm{CON}$;*
   (b) *if $R(r') \neq \emptyset$, then f has one child g_τ for each $\tau \in R(r')$, with $A_{g_\tau} = \mathrm{insert}_F(A_f, v, w, \tau)$.*

A star proof *of $E \to \sigma \le t$ is a well-founded star tree for $E \to \sigma \le t$. We say $E \to \sigma \le t$ is* star provable *if such a proof exists.*

A free automaton is uniquely determined by its language (putting aside the minor issue of strings occurring multiple times), so Definition 16 could be reformulated without reference to automata, using instead sets of finite strings. However, there are advantages to preserving as much structural similarity as possible between the definitions of relational and star proofs: it better isolates the differences between $\mathcal{H}\mathsf{RKA}$ and $\mathcal{H}\mathsf{KA}^*$, and it sometimes allows the same proof-theoretic argument to be made with both systems with minimal modification (see Theorem 27, for example).

Lemma 17 *For any Horn formula of the form $E \to \sigma \le t$, the following are equivalent.*

(i) $\mathsf{KA}^* \models E \to \sigma \le t$

(ii) *$E \to \sigma \le t$ is star provable.*

The definition of star provability is extended to arbitrary Horn formulas exactly as it was for relational provability, and once the above lemma is established, the soundness and completeness will follow analogously. However, the proof of the above soundness and completeness lemma will have little resemblance to the proof of the corresponding lemma for relational proofs, and the intuitive meaning of a star tree will be unrelated to the reasoning in Section 3.1. Instead, we will develop a language-theoretic proof system for $\mathcal{H}\mathsf{KA}^*$, and show that star provability coincides with provability in this language-theoretic system. We must develop this extra theory before proving Lemma 17.

4.2 A Language-Theoretic Look at $\mathcal{H}\mathsf{KA}^*$

In this section, we develop a characterization of $\mathcal{H}\mathsf{KA}^*$ in terms of a closure operator. This is used establish soundness of completeness of star provability, and can also be treated as another proof system.

Fix a finite set of hypotheses E. As before, assume all hypotheses are inequalities by breaking equations $x = y$ into $x \leq y$ and $y \leq x$.

Definition 18 *For $t \in \mathsf{RExp}_\Sigma$, we define*

$$\mathrm{below}_E(t) = \{\sigma \in \Sigma^* \mid \mathsf{KA}^* \models E \to \sigma \leq t\} \ .$$

We will now see how to construct $\mathrm{below}_E(t)$ using a closure operator.

Definition 19 *For any $S \subseteq \Sigma^*$, we define $\mathrm{new}_E(S)$ to be the set of all strings of the form $\sigma\rho\tau$ such that $\sigma, \rho, \tau \in \Sigma^*$ and there is a hypothesis $r \leq r'$ in E such that $\rho \in R(r)$, and $\sigma R(r')\tau \subseteq S$ (that is, $\sigma\rho'\tau \in S$ for all $\rho' \in R(r')$). We say $S \subseteq \Sigma^*$ is E-*closed *if* $\mathrm{new}_E(S) \subseteq S$.

We define the E-closure of $S \subseteq \Sigma^*$ by

$$\mathrm{cl}_E(S) = \bigcap\{T \subseteq \Sigma^* \mid S \subseteq T \text{ and } T \text{ is } E\text{-closed}\} \ .$$

For $t \in \mathsf{RExp}_\Sigma$, we define

$$\mathrm{sp}_E(t) = \{\sigma \in \Sigma^* \mid E \to \sigma \leq t \text{ is star provable}\} \ .$$

Theorem 20 *For any $t \in \mathsf{RExp}_\Sigma$, we have*

$$\mathrm{below}_E(t) = \mathrm{cl}_E(R(t)) = \mathrm{sp}_E(t) \ .$$

Proof See Section A.2 of the appendix. ☐

Corollary 21 *The following are equivalent.*

(i) $\mathsf{KA}^* \models E \to \sigma \leq t$

(ii) $\sigma \in \mathrm{cl}_E(R(t))$

(iii) *$E \to \sigma \leq t$ is star provable.*

Proof Observe that (i)–(iii) each state that σ is in one of the three sets shown to be equal in the above theorem. ☐

Note that Corollary 21 trivially implies Lemma 17.

Theorem 22 *The following are equivalent.*

(i) $\mathsf{KA}^* \models E \to s = t$

(ii) *$R(s) \subseteq \mathrm{cl}_E(R(t))$ and $R(t) \subseteq \mathrm{cl}_E(R(s))$.*

(iii) $\mathrm{cl}_E(R(s)) = \mathrm{cl}_E(R(t))$

(iv) *$E \to s = t$ is star provable.*

Proof (i)⇔(ii) is immediate from Corollary 21 and Lemma 9.

(ii)⇔(iii) follows from the easy observation that for any $S \subseteq \Sigma^*$, $S \subseteq \mathrm{cl}_E(S) = \mathrm{cl}_E(\mathrm{cl}_E(S))$.

(i)⇔(iv) is analogous to the proof of Theorem 14. ☐

In addition to giving us the desired soundness and completeness of star provability, Theorem 22 (and Corollary 21) also show us how to use E-closure as a sort of proof system for $\mathcal{H}\mathsf{KA}^*$. For example, recall the Horn formula

$$r_0 \leq r_1 + r_2 \wedge s_0 \leq s_1 + s_2 \to r_0 s_0 \leq r_0 s_1 + r_1 s_2 + r_2 s_0$$

from Section 3.1. This has the form $E \to \sigma \leq t$ where

$$\begin{aligned} E &= \{r_0 \leq r_1 + r_2,\ s_0 \leq s_1 + s_2\} \\ \sigma &= r_0 s_0 \\ t &= r_0 s_1 + r_1 s_2 + r_2 s_0 \ . \end{aligned}$$

We can confirm by inspection that $\mathrm{new}_E(R(t)) = \mathrm{new}_E(\{r_0 s_1, r_1 s_2, r_2 s_0\}) = \emptyset$. Thus $R(t)$ is E-closed, so $\mathrm{cl}_E(R(t)) = R(t) = \{r_0 s_1, r_1 s_2, r_2 s_0\}$. Thus, $\sigma \notin \mathrm{cl}_E(R(t))$, so $\mathsf{KA}^* \not\models E \to \sigma \leq t$.

By comparison, it is a lot of work to directly construct a KA^* in which this Horn formula fails.

4.3 Example Application 1: Simple Formulas

This example uses a proof-theoretic argument to give a short proof of a known result.

Definition 23 *The Horn formula*

$$(s_1 \leq t_1 \wedge \cdots \wedge s_n \leq t_n) \to s \leq t$$

is simple *if $*$ does not occur in s or any of the t_i.*

Theorem 24 *$\mathcal{H}\mathsf{KA}^*$ and $\mathcal{H}\mathsf{RKA}$, restricted to simple Horn formulas, are Σ^0_1.*

(In fact, these theories are Σ^0_1-complete, but we do not argue that here.)

Proof Clearly, if $r \in \mathsf{RExp}_\Sigma$ is $*$-free, then $R(r)$ is finite.

Let φ be a simple Horn formula of the form $(s_1 \leq t_1 \wedge \cdots \wedge s_n \leq t_n) \to s \leq t$. For each $\sigma \in R(s)$, observe that a relational proof of $(s_1 \leq t_1 \wedge \cdots \wedge s_n \leq t_n) \to \sigma \leq t$ would be finitely branching, since $R(t_i)$ is finite for each t_i. Well-founded finitely-branching trees must be finite, so if $(s_1 \leq t_1 \wedge \cdots \wedge s_n \leq t_n) \to \sigma \leq t$ is provable at all, it has a finite relational proof. This proof could be coded with a single natural number. So $\mathsf{RKA} \models \varphi$ is equivalent to the Σ^0_1 statement, "For each $\sigma \in R(s)$, there exists an $n \in \omega$ encoding a finite relational proof of $(s_1 \leq t_1 \wedge \cdots \wedge s_n \leq t_n) \to \sigma \leq t$." (The finiteness of $R(s)$ makes this Σ^0_1, as opposed to Π^0_2.)

Therefore, $\mathcal{H}$RKA, restricted to simple formulas, is Σ_1^0. The same argument (with star proofs in place of relational proofs) shows that $\mathcal{H}\mathsf{KA}^*$, restricted to simple formulas, is Σ_1^0. □

4.4 Example Application 2: Eliminating $r = 0$

We can use proof-theoretic methods to improve the known results involving the elimination of $r = 0$.

Definition 25 *The* universal regular expression *is* $u = (p_1 + \cdots + p_n)^*$*, where* $\Sigma = \{p_1, \ldots, p_n\}$.

The following well-known theorem [2, 10, 13] shows that we can replace any Horn formula $r = 0 \to s = t$ with the equation $s + uru = t + uru$. This process is called *eliminating* $r = 0$.

Theorem 26 *For any* $r, s, t \in \mathsf{RExp}_\Sigma$*, the following are equivalent.*

(i) $\mathsf{KA} \models r = 0 \to s = t$

(ii) $\mathsf{KA}^* \models r = 0 \to s = t$

(iii) $\mathsf{RKA} \models r = 0 \to s = t$

(iv) $\mathsf{KA} \models s + uru = t + uru$

We can improve Theorem 26 to the following.

Theorem 27 *Let* E *be any finite set of hypotheses, and* $r, s, t \in \mathsf{RExp}_\Sigma$*. Let* φ *be the formula*

$$E \wedge r = 0 \to s = t \ ,$$

and let φ' *be the formula*

$$E \to s + uru = t + uru \ .$$

Then the following equivalences hold.

$$\mathsf{KA} \models \varphi \iff \mathsf{KA} \models \varphi' \qquad (5)$$
$$\mathsf{KA}^* \models \varphi \iff \mathsf{KA}^* \models \varphi' \qquad (6)$$
$$\mathsf{RKA} \models \varphi \iff \mathsf{RKA} \models \varphi' \qquad (7)$$

Note that the special case $E = \emptyset$ gives us Theorem 26: when $E = \emptyset$, the right hand sides of (5)–(7) are equivalent, since the equational theories of KA, KA^*, and RKA coincide. The significance of Theorem 27 is that, beyond simply eliminating $r = 0$, it shows that we can extend any method of eliminating hypotheses to also include hypotheses of the form $r = 0$. For, supposing we have a technique to eliminate $q_1 = q_2$, we can eliminate $q_1 = q_2 \wedge r = 0$ by eliminating $r = 0$ first with Theorem 27, leaving only the hypothesis $q_1 = q_2$ which we eliminate with its existing technique.

Proof The proof of (5) is not of interest here, and can be found in [5].

For (7), the right-to-left implication is trivial: assuming φ' and $E \wedge r = 0$, we have $s = s + 0 = s + uru = t + uru = t + 0 = t$. For the left-to-right implication, suppose $\mathsf{RKA} \models E \wedge r = 0 \to \sigma \leq t$, where $\sigma \in R(s)$. $r = 0$ is equivalent to $r \leq 0$, and $\mathsf{KA} \models t \leq t + uru$, so $\mathsf{RKA} \models E \wedge r \leq 0 \to \sigma \leq t + uru$. Let (T, A) be a relational proof of $E \wedge r \leq 0 \to \sigma \leq t + uru$.

We claim that the hypothesis $r \leq 0$ is never even applied in the proof! Suppose $r \leq 0$ is applied at node $f \in T$ (so f has one child g with $A_g = \mathrm{CON}$). For $r \leq 0$ to be applied at f, there must be states $v, w \in |A_f|$ and $\rho \in R(r)$ with $\rho \in L(A_f^{v,w})$. A property that is preserved in the automata of relational trees is that every state is accessible from the start state a, and the accept state b is accessible from every state. So there exist $\pi \in L(A_f^{a,v})$ and $\pi' \in L(A_f^{w,b})$. Thus, we have $\pi\rho\pi' \in L(A_f)$; we also have $\pi\rho\pi' \in R(uru) \subseteq R(t + uru)$. Therefore, $R(t + uru) \cap L(A_f) \neq \emptyset$, so f is in fact a leaf node, contradicting the assumption that we are applying $r \leq 0$ at f. (In other words, at any point in a relational tree for $E \wedge r \leq 0 \to \sigma \leq t + uru$ where we could apply $r \leq 0$, we would already be done along that branch.)

So, because $r \leq 0$ is never applied, (T, A) is also a relational proof of $E \to \sigma \leq t + uru$. Therefore, $\mathsf{RKA} \models E \to \sigma \leq t + uru$ for all $\sigma \in R(s)$. By Lemma 9, $\mathsf{RKA} \models E \to s \leq t + uru$, so $\mathsf{RKA} \models E \to s + uru \leq t + uru$. $\mathsf{RKA} \models E \to t + uru \leq s + uru$ is similar, and we now have $\mathsf{RKA} \models E \to s + uru = t + uru$.

For (6), the above argument also works for star proofs, with the small modification that we take π to be the prefix of v in A_f and π' to be the suffix of w in A_f.

□

5 Incorporating Tests

There is no particular difficulty in incorporating tests into relational and star provability, to yield systems that are sound and complete for $\mathcal{H}$RKAT and $\mathcal{H}\mathsf{KAT}^*$. The reason for this is that Boolean algebra is equationally axiomatized, so we can modify our proof systems to act as though appropriate instances of Boolean algebra axioms are among the hypotheses of the Horn formula in question.

Fix finite sets of atomic program symbols P and tests B. For each $b \in \mathsf{B}$, we introduce two fresh atomic program symbols $\tilde{b}$ and $\tilde{\bar{b}}$, and we let $\tilde{\mathsf{B}} = \left\{\tilde{b}, \tilde{\bar{b}} \mid b \in \mathsf{B}\right\}$. We let $\Sigma = \mathsf{P} \cup \tilde{\mathsf{B}}$.

Applying DeMorgan's laws as necessary, we can assume that Boolean negation is only applied to atomic tests in the Horn formulas we consider. In this way, any $s \in \mathsf{RExp}_{\mathsf{P},\mathsf{B}}$ can be treated as if it is in RExp_Σ.

So at least language-wise, our proof systems can handle Horn formulas in the language of KAT. We must make some minor modifications for them to prove the correct Horn formulas, though. The justification of these modifications is straightforward but tedious and is omitted here.

In our definition of a star tree for $E \to \sigma \le t$, we now allow each of the following to be used as if they were in E:

$$\begin{aligned} \tilde{b} + \tilde{\tilde{b}} &= 1 \quad (b \in \mathsf{B}) && (8)\\ \tilde{b}\tilde{\tilde{b}} &= 0 \quad (b \in \mathsf{B}) && (9)\\ bc &= cb \quad (b, c \in \tilde{\mathsf{B}}) && (10)\\ b &= bb \quad (b \in \tilde{\mathsf{B}}) && (11) \end{aligned}$$

In our definition of a relational tree for $E \to \sigma \le t$, we now allow (8) and (9) to be used as if they were in E. (We do not need (10) and (11) here, since they are made redundant by (8) in the context of relational trees.)

6 Acknowledgments

I would like to thank Dexter Kozen for his suggestions.

This work was supported in part by NSF grant CCR-0105586 and by ONR Grant N00014-01-1-0968. The views and conclusions contained herein are those of the author and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of these organizations or the US Government.

References

[1] A. Barth and D. Kozen. Equational Verification of Cache Blocking in LU Decomposition using Kleene Algebra with Tests. Technical Report 2002-1865, Computer Science Department, Cornell University, June 2002.

[2] E. Cohen. Hypotheses in Kleene Algebra. Unpublished, 1994.

[3] E. Cohen, D. Kozen, and F. Smith. The complexity of Kleene algebra with tests. Technical Report 96-1598, Computer Science Department, Cornell University, July 1996.

[4] J. H. Conway. *Regular Algebra and Finite Machines*. Chapman and Hall, London, 1971.

[5] C. Hardin. *The Horn Theory of Relational Kleene Algebra*. PhD thesis, Cornell University, 2005. In preparation.

[6] C. Hardin and D. Kozen. On the Elimination of Hypotheses in Kleene Algebra with Tests. Technical Report 2002-1879, Computer Science Department, Cornell University, Oct. 2002.

[7] C. Hardin and D. Kozen. On the Complexity of the Horn Theory of REL. Technical Report 2003-1896, Computer Science Department, Cornell University, May 2003.

[8] D. Kozen. *The Design and Analysis of Algorithms*. Springer-Verlag, New York, 1991.

[9] D. Kozen. Kleene algebra with tests. *Transactions on Programming Languages and Systems*, pages 427–443, 1997.

[10] D. Kozen. On Hoare logic and Kleene algebra with tests. *Trans. Computational Logic*, 1(1):60–76, July 2000.

[11] D. Kozen. On the Complexity of Reasoning in Kleene Algebra. *Information and Computation*, 179:159–162, 2002.

[12] D. Kozen and M.-C. Patron. Certification of compiler optimizations using Kleene algebra with tests. In J. Lloyd, V. Dahl, U. Furbach, M. Kerber, K.-K. Lau, C. Palamidessi, L. M. Pereira, Y. Sagiv, and P. J. Stuckey, editors, *Proc. 1st Int. Conf. Computational Logic (CL2000)*, volume 1861 of *Lecture Notes in Artificial Intelligence*, pages 568–582, London, July 2000. Springer-Verlag.

[13] D. Kozen and F. Smith. Kleene algebra with tests: completeness and decidability. In D. van Dalen and M. Bezem, editors, *Proc. 10th Int. Workshop on Computer Science Logic (CSL'96)*, volume 1258 of *Springer-Verlag Lecture Notes in Computer Science*, pages 244–259, Utrecht, The Netherlands, Sept. 1996.

A Selected Proofs

A.1 Soundness and Completeness of Relational Provability

We now show that relational provability coincides with relational validity. We start with soundness.

Lemma 28 *If $E \to \sigma \le t$ is relationally provable, then* $\mathsf{RKA} \models E \to \sigma \le t$.

Proof Suppose (T, A) is a relational proof of $E \to \sigma \le t$. Take any $K \in \mathsf{RKA}$ and interpretation $I : \mathsf{RExp}_\Sigma \to K$ such that $K, I \models E$. We wish to show $K, I \models \sigma \le t$, *i.e.*, $I(\sigma) \subseteq I(t)$. Let X be the base of K.

Suppose $(x, y) \in I(\sigma)$. We will use K, I to pick out a path $f_0, f_1, \ldots$ through the proof tree until we hit a leaf at some f_m, and that leaf will witness that $(x, y) \in I(t)$. We will interpret the automata along that path as approximations of K; to do this, we will need to specify how the states of the automata correspond to states in X: for each state c of an automaton on the path, ν_c will be the corresponding state in X. We will say that an edge $v \xrightarrow{r} w$ of an automaton is *I-sound* if $(\nu_v, \nu_w) \in I(r)$. At each f_i which is not a leaf, f_{i+1} will be a child of f_i, and we will inductively preserve the following property.

$$\text{Every edge in } A_{f_i} \text{ is } I\text{-sound.} \qquad (12)$$

When (12) holds, we clearly also have the following.

$$\begin{array}{c}\text{For each } v, w \in |A_{f_i}| \text{ and } \pi \in L(A_{f_i}^{v,w}),\\ \text{we have } (\nu_v, \nu_w) \in I(\pi).\end{array} \qquad (13)$$

We define $f_0 = \langle\rangle$. Recall from Definition 11 that $A_{\langle\rangle}$ has starting state a and accept state b; we define $\nu_a = x$ and $\nu_b = y$. We can choose states that witness $(x, y) \in I(\sigma)$

to define ν_c for the other states $c \in |A_{\langle\rangle}|$, and it is easy to verify (12).

Now suppose we have defined f_i such that (12) holds, and that f_i is not a leaf node. By the definition of a relational tree for $E \to p \le q$, there must exist $v, w \in |A_{f_i}|$, $\rho \in L(A^{v,w}_{f_i})$, and a hypothesis $r \le r'$ in E such that $\rho \in R(r)$ and conditions 3a and 3b from Definition 11 hold.

By (13), $(\nu_v, \nu_w) \in I(\rho)$. By Corollary 7, $I(\rho) \subseteq I(r)$, since $\rho \in R(r)$. Because $K, I \models r \le r'$, we have $I(r) \subseteq I(r')$. So $(\nu_v, \nu_w) \in I(\rho) \subseteq I(r) \subseteq I(r')$. Therefore, by Corollary 7, there must be $\rho' \in R(r')$ such that $(\nu_v, \nu_w) \in I(\rho')$. (This also gives us $R(r') \neq \emptyset$.)

By condition 3b from Definition 11, f_i must have a child g_τ with $A_{g_\tau} = \mathrm{insert}(A_{f_i}, v, w, \tau)$. Let $f_{i+1} = g_\tau$. We must define ν_u for any new states appearing in $A_{f_{i+1}}$, and show that every edge in $A_{f_{i+1}}$ is I-sound. Because we are taking (12) as our inductive hypothesis, we only need to verify I-soundness for the new edges—those edges appearing in $A_{f_{i+1}}$ that do not already appear in A_{f_i}.

Case 1: $\rho' = \varepsilon$. Then the above insert operation does not add any new states, so there are no new ν_u to define. We have two new edges, $v \xrightarrow{\varepsilon} w$ and $w \xrightarrow{\varepsilon} v$. For $v \xrightarrow{\varepsilon} w$, we already have $(\nu_v, \nu_w) \in I(\rho')$ from above. For $w \xrightarrow{\varepsilon} v$, we simply observe that $I(\rho') = I(1)$ is the identity relation, whose symmetry gives us $(\nu_w, \nu_v) \in I(\rho')$.

Case 2: $\rho' = p_1 \cdots p_k$, $k > 0$. Since $(\nu_v, \nu_w) \in I(\rho') = I(p_1) \cdots I(p_k)$, there must be states $y_1, \ldots, y_{k-1} \in X$ such that $(\nu_v, y_1) \in I(p_1)$, $(y_i, y_{i+1}) \in I(p_{i+1})$ for $1 \le i \le k-2$, and $(y_{k-1}, \nu_w) \in I(p_k)$. For $1 \le i \le k-1$, let $\nu_{x_i} = y_i$ (where the x_i are the new states created by the insert operation). The new edges in $A_{f_{i+1}}$ are $v \xrightarrow{p_1} x_1$, $x_i \xrightarrow{p_{i+1}} x_{i+1}$ for $1 \le i \le k-2$, and $x_{k-1} \xrightarrow{p_k} w$; our definition of the ν_{x_i} gives us I-soundness immediately for these new edges.[1]

These cases are exhaustive, so ν_u is now defined for all states $u \in |A_{f_{i+1}}|$, and every edge in $A_{f_{i+1}}$ is I-sound.

That completes the induction, leaving us with $f_0, f_1, \ldots, f_m$, where (12) holds at each f_i, and f_m is a leaf node. (We must eventually hit a leaf node because T is well founded.) By condition 2 of Definition 11, $A_{f_m} = \mathrm{CON}$ or $R(t) \cap L(A_{f_m}) \neq \emptyset$; the former is impossible because our construction never defined any f_i to be CON, so we must have $R(t) \cap L(A_{f_m}) \neq \emptyset$. Note that $A_{f_m} = A^{a,b}_{f_m}$, since each automaton in the tree has unique start and end states a and b, so $R(t) \cap L(A^{a,b}_{f_m}) \neq \emptyset$. Letting $\pi \in R(t) \cap L(A^{a,b}_{f_m})$, (13) gives us $(x, y) = (\nu_a, \nu_b) \in I(\pi) \subseteq I(t)$.

Therefore, $I(\sigma) \subseteq I(t)$, so $K, I \models \sigma \le t$.

Therefore, $\mathsf{RKA} \models E \to \sigma \le t$. □

We now turn to completeness. Fix a Horn formula $E \to \sigma \le t$. We will construct a canonical relational tree for $E \to \sigma \le t$ with the property that any infinite path through it will give us a relational model in which $E \to \sigma \le t$ fails. We will generate a model from a path by taking the union of the automata along that path (and modding out by an appropriate equivalence relation). As we will see, a path through an arbitrary relational tree for $E \to \sigma \le t$ would give us a model in which $\sigma \le t$ fails, so our work in producing the canonical tree will consist of making sure any resulting models satisfy E. Once we have done this, the relational validity of $E \to \sigma \le t$ will imply that the canonical tree can have no path (without getting a contradiction), so it must be well founded, leaving us with a relational proof of $E \to \sigma \le t$.

We first look at how to construct relational models from paths through relational trees. Suppose (T, A) is a relational tree for $E \to \sigma \le t$, and that $f : \omega \to \omega$ is a path through T, so that $A_{\langle\rangle}, A_{f\restriction 1}, A_{f\restriction 2}, \ldots$ is the sequence of automata appearing along f, each one extending the previous one. (In particular, we have the following monotonicity property: if $m \le n$, and $x, y \in |A_{f\restriction m}|$, then $L(A^{x,y}_{f\restriction m}) \subseteq L(A^{x,y}_{f\restriction n})$.) Let $Y_0 = \bigcup_{n\in\omega} |A_{f\restriction n}|$. We define an equivalence relation $\approx$ on Y_0 by

$$x \approx y \iff (\exists n \in \omega)\ \varepsilon \in L(A^{x,y}_{f\restriction n}) \ .$$

(We say n *witnesses* $x \approx y$ if $\varepsilon \in L(A^{x,y}_{f\restriction n})$.) The reflexivity of $\approx$ is trivial: if $x \in |A|$, then $\epsilon \in L(A^{x,x})$. Symmetry follows from the fact that, in relational trees, whenever some ε-edge $x \xrightarrow{\varepsilon} y$ is added (by an "insert"), the edge $y \xrightarrow{\varepsilon} x$ is also added. For transitivity, suppose $x \approx y$ (witnessed by n_0) and $y \approx z$ (witnessed by n_1); then, letting $n = \max(n_0, n_1)$, we have

$$\varepsilon = \varepsilon\varepsilon \in L(A^{x,y}_{f\restriction n}) \cdot L(A^{y,z}_{f\restriction n}) \subseteq L(A^{x,z}_{f\restriction n}) \ ,$$

so $x \approx z$. Therefore, $\approx$ is an equivalence relation on Y_0.

Let $[x]$ denote the $\approx$-equivalence class of x, and let $Y = Y^{T,A,f} = \{[x] \mid x \in Y_0\}$. We define $J = J^{T,A,f}$ by $J : \Sigma \to \mathcal{R}(Y)$ by

$$J(r) = \{([x], [y]) \mid (\exists n \in \omega)\ r \in L(A^{x,y}_{f\restriction n})\} \ ,$$

and extend homomorphically to an interpretation $J : \mathsf{RExp}_\Sigma \to \mathcal{R}(Y)$.

Lemma 29 *For any* $\rho \in \Sigma^*$,

$$J(\rho) = \{([x], [y]) \mid (\exists n \in \omega)\ \rho \in L(A^{x,y}_{f\restriction n})\} \ .$$

Proof This is a straightforward induction on the length of ρ. □

Lemma 30 $\mathcal{R}(Y), J \not\models \sigma \le t$

Proof $([a], [b]) \in J(\sigma)$ is immediate. Suppose, though, that $([a], [b]) \in J(t)$. Then there exist $x, y \in Y_0$ with $x \approx a$ (witnessed by some n_0) and $y \approx b$ (witnessed by some n_1),

[1] The notation here is slightly broken in the case that $k = 1$, since we will have $(\nu_v, \nu_w) \in I(p_1)$, but the argument is the same.

such that for some n_2, $R(t) \cap L(A^{x,y}_{f\restriction n_2}) \neq \emptyset$, witnessed by some π. Letting $n = \max(n_0, n_1, n_2)$, we have

$$\begin{aligned} \pi = \varepsilon\pi\varepsilon &\in L(A^{a,x}_{f\restriction n}) \cdot L(A^{x,y}_{f\restriction n}) \cdot L(A^{y,b}_{f\restriction n}) \\ &\subseteq L(A^{a,b}_{f\restriction n}) = L(A_{f\restriction n}) \ . \end{aligned}$$

By condition 2 of Definition 11, this would make $f \restriction n$ a leaf node, contradicting the fact that f is a path through T. Therefore, $([a],[b]) \notin J(t)$, so $\mathcal{R}(Y), J \not\models \sigma \leq t$. □

By the previous lemma, if we can construct a relational tree (T, A) for $E \to \sigma \leq t$ such that $\mathcal{R}(Y^{T,A,f}), J^{T,A,f} \models E$ for any path f through T, then the relational validity of $E \to \sigma \leq t$ will imply that T must be well founded, yielding a relational proof for $E \to \sigma \leq t$.

Lemma 31 *If* $\mathsf{RKA} \models E \to \sigma \leq t$*, then* $E \to \sigma \leq t$ *is relationally provable.*

Proof We define the *canonical relational tree* (T, A) for $E \to \sigma \leq t$ as follows. Note that the only degree of freedom that Definition 11 gives us is the choice of what hypothesis to apply at a node g, and where within A_g to apply it.

Our goal is to guarantee that $\mathcal{R}(Y^{T,A,f}), J^{T,A,f} \models E$ for any path f through T. To accomplish this, we make sure that for any hypothesis $r \leq r'$ in E, if a pair $([x],[y])$ enters $J(r)$, we make sure that it also enters $J(r')$ by having appropriate children as in 3b of Definition 11. Of course, we may have infinitely many such requirements to take care of. Fortunately, there are only countably many such requirements that could *ever* occur, so we fix a function $C : \omega \to \{\text{all possible requirements}\}$ such that each requirement is hit infinitely often. (Formally, a "posible requirement" is just a tuple specifying a hypothesis $r \leq r'$ and a pair of states v, w, but the details of the coding are irrelevant.) To construct the canonical proof tree, if we are at a node on level n and need to choose what hypothesis to apply, and where, we apply the hypothesis specified by $C(n)$ at the states specified by $C(n)$ if possible; otherwise we apply any hypothesis we like. (It is safe to assume that some hypothesis can be applied, because we can pretend that we have a hypothesis $1 \leq 1$, and simply add ε-edges from a to itself; allowing ourselves to do so clearly would not let us prove any Horn formulas that we could not prove under the strict interpretation of Definition 11.)

Supposing $E \to \sigma \leq t$ is not relationally provable, the canonical proof tree defined above must not be well founded, so it has a path f. Along this path, if a pair $([x],[y])$ enters $J(r)$ for some hypothesis $r \leq r'$ in E, our construction above guarantees that we will eventually apply the hypothesis $r \leq r'$ at states x, y along this path, and this will ensure $([x],[y]) \in J(r')$. It follows that $J(r) \subseteq J(r')$ for every hypothesis $r \leq r'$ in E, so $\mathcal{R}(Y), J \models E$. However, by Lemma 30, $\mathcal{R}(Y), J \not\models \sigma \leq t$. Therefore, $\mathsf{RKA} \not\models E \to \sigma \leq t$. □

Proof of Lemma 12 Immediate from Lemmas 28 and 31. □

A.2 Soundness and Completeness of Star Provability

Our goal is to prove Theorem 20. Let $t \in \mathsf{RExp}_\Sigma$ and E be a finite set of hypotheses.

Lemma 32 $\mathrm{below}_E(t) \subseteq \mathrm{cl}_E(R(t))$

Proof We construct a $K \in \mathsf{KA}^*$ as follows. Let $K = \{\mathrm{cl}_E(S) \mid S \subseteq \Sigma^*\}$. We define

$$\begin{aligned} 0 &= \mathrm{cl}_E(\emptyset) \\ 1 &= \mathrm{cl}_E(\{\varepsilon\}) \\ A \oplus B &= \mathrm{cl}_E(A \cup B) \\ A \odot B &= \mathrm{cl}_E(AB) \\ A^{\circledast} &= \mathrm{cl}_E\left(\bigcup_{n\in\omega} A^n\right) \ . \end{aligned}$$

It is straightforward to verify that K forms a $*$-continuous Kleene algebra under these operations, using the usual properties of closure, as well as the fact that for any $S, T \subseteq \Sigma^*$, we have $\mathrm{cl}_E(S)\mathrm{cl}_E(T) \subseteq \mathrm{cl}_E(ST)$.

We define the interpretation $I : \mathsf{RExp}_\Sigma \to K$ by $I(s) = \mathrm{cl}_E(R(s))$. For any hypothesis $r \leq r'$ in E, $R(r) \subseteq \mathrm{cl}_E(R(r'))$, so $K, I \models E$.

Suppose $\sigma \in \mathrm{below}_E(t)$. Then $K, I \models \sigma \leq t$, so

$$\sigma \in R(\sigma) \subseteq I(\sigma) \subseteq I(t) = \mathrm{cl}_E(R(t)) \ .$$

Therefore, $\mathrm{below}_E(t) \subseteq \mathrm{cl}_E(R(t))$. □

Lemma 33 $\mathrm{cl}_E(R(t)) \subseteq \mathrm{sp}_E(t)$

Proof $R(t) \subseteq \mathrm{sp}_E(t)$ is trivial from the definition of star proof, so it suffices to show that $\mathrm{sp}_E(t)$ is E-closed.

Take any $\sigma\rho\tau \in \mathrm{new}_E(\mathrm{sp}_E(t))$, with $r \leq r'$ a hypothesis in E such that $\rho \in R(r)$ and $\sigma R(r')\tau \subseteq \mathrm{sp}_E(t)$. We can see that $E \to \sigma\rho\tau \leq t$ must be star provable: at the root node, we apply the hypothesis $r \leq r'$ so that the children of the root have a string $\sigma\rho'\tau$ with $\rho' \in R(r')$, and we can build proof trees below each of these children because these $\sigma\rho'\tau$ are in $\mathrm{sp}_E(t)$. So $\sigma\rho\tau \in \mathrm{sp}_E(t)$. Therefore $\mathrm{new}_E(\mathrm{sp}_E(t)) \subseteq \mathrm{sp}_E(t)$, so $\mathrm{sp}_E(t)$ is E-closed. □

Lemma 34 $\mathrm{sp}_E(t) \subseteq \mathrm{below}_E(t)$

Proof Suppose $\sigma \in \mathrm{sp}_E(t)$. Fix a star proof (T, A) of $E \to \sigma \leq t$. One can show by induction in well-founded trees (that is, starting at the leaves and working up to the root) that at every node $f \in T$, either $A_f = \mathrm{CON}$ or $L(A_f) \cap \mathrm{below}_E(t) \neq \emptyset$. Then, at the root node, we have $L(A_{\langle\rangle}) = \{\sigma\}$, and this must intersect $\mathrm{below}_E(t)$. Therefore, $\sigma \in \mathrm{below}_E(t)$. $\square$

Proof of Theorem 20 Immediate from Lemmas 32, 33, 34. $\square$