How the Location of $*$ Influences Complexity in Kleene Algebra with Tests

Chris Hardin

Department of Mathematics
Cornell University
Ithaca, New York 14853-4201, USA
hardin@math.cornell.edu

Abstract. The universal Horn theory of relational Kleene algebra with tests is of practical interest, particularly for program semantics, where Horn formulas can be used to verify correctness of programs or compiler optimizations. Unfortunately, this theory is known to be Π_1^1-complete. However, many formulas arising in practice fall into fragments of the theory that are of lower complexity. In this paper, we see that the location of occurrences of the Kleene asterate operator * within a formula has a great impact on complexity. Using syntactic criteria based on the location of *, we give a fragment of the theory that is Σ_1^0-complete, and a slightly larger fragment that is Π_2^0-complete. We show that the same results hold over *-continuous Kleene algebras with tests. The techniques exhibit a relationship between first-order logic and the Horn theories of relational and *-continuous Kleene algebra, even though the theories are not first-order axiomatizable.

1 Introduction

The universal Horn theories of $*$-continuous and relational Kleene algebras (with tests) are of great interest, particularly for program semantics. Unfortunately, these Horn theories are are Π_1^1-complete [8, 5], making them difficult to work with in full generality. However, under various restrictions on the formulas, the complexity is often much lower. For example, when the hypotheses are restricted to the form $s = 0$, the theory (in both cases) is *PSPACE*-complete [2, 3, 10, 6]. See [8, 6] for further examples. In this paper, we investigate how the location of $*$ in a Horn formula affects complexity.

For the rest of the introduction, we find it convenient to be non-rigorous and use many terms without defining them, in the hopes of quickly and abstractly sketching the material that will be presented and some intuition behind it. The rest of the paper is more self-contained—barely even requiring the introduction, although it does assume a familiarity with first-order logic—and more responsible about proof.

The $*$-continuity axiom, stated succinctly as $xy^*z = \sup_{n\in\omega} xy^nz$, can equivalently be expressed by

$$xy^nz \leq xy^*z \quad \text{(for each } n \in \omega\text{)} \tag{1}$$

and the infinitary Horn formula

$$\bigwedge_{n\in\omega} xy^n z \leq w \to xy^* z \leq w \ . \tag{2}$$

Informally, (1), by bounding xy^*z from below, describes the bigness of y^*, while (2), by bounding xy^*z from above, describes the smallness of y^*. The bigness condition (1) is first-order, while the smallness condition (2) is not. By appropriately restricting where $*$ may appear in a Horn formula, we can make the validity of the Horn formula (over $*$-continuous Kleene algebras) depend only on the bigness of $*$, while the smallness is irrelevant; first-order logic will then be adequate for determining the validity of such formulas, which will pull the complexity down to Σ^0_1 ("There exists a proof…"). Specifically, the restriction will be that in the hypotheses, $*$ may only appear on the left-hand side of inequalities $s_i \leq t_i$, and in the conclusion, $*$ may only appear on the right hand side of an inequality $s \leq t$; such formulas are called *simple*.

Since everything but (2) in the definition of $*$-continuous Kleene algebra is first-order, while the Horn theory is Π^1_1-complete, it must be the smallness condition (2) that admits Π^1_1-hardness. How does this happen? Here are two intuitions, closely related to each other:

1. First-order logical consequence comes down to the well-foundedness of finitely branching trees, which (with suitable effectiveness conditions on the trees) is Σ^0_1. (One can think of these trees as proof trees; one can also think of these trees as systematic attempts to construct a counterexample, in which infinite paths yield counterexamples, while well-foundedness constitutes a proof. The distinction is only superficial.)
   We could extend first-order logic to incorporate (2) by adding the infinitary inference rule
   $$\frac{xy^n z \leq w \quad (\text{for each } n \in \omega)}{xy^* z \leq w} \ .$$
   However, our proof trees will no longer be finitely branching, and the problem of well-foundedness of (recursive) infinitely branching trees is Π^1_1-complete. Loosely, what has happened here is that well-foundedness can no longer be expressed as "There exists n such that there are no nodes of depth n," which was adequate for finitely branching trees; instead, well-foundedness must now be expressed as "All paths eventually hit a leaf node," and quantifying over paths is second-order.
2. Given a set A of first-order formulas (in the language of Kleene algebra), let $\mathsf{Th}(A) = \{\varphi \mid A \models \varphi\}$. This is a *closure operator* in the standard sense: $A \subseteq \mathsf{Th}(A)$, $\mathsf{Th}(\mathsf{Th}(A)) = \mathsf{Th}(A)$, $A \subseteq B \Rightarrow \mathsf{Th}(A) \subseteq \mathsf{Th}(B)$; A is *closed* if $A = \mathsf{Th}(A)$. There are two ways to build $\mathsf{Th}(A)$ from A: from below, and from above. To build $\mathsf{Th}(A)$ from below, we start with A, and iterate the process of throwing in axioms and applying inference rules; after countably many iterations, we will have $\mathsf{Th}(A)$, and this lets us express $\varphi \in \mathsf{Th}(A)$ with the Σ^0_1 formula "There exists a stage in this iterative process at which φ appears." To build $\mathsf{Th}(A)$ from above, we take the intersection of all closed

sets containing A;[1] this lets us express $\varphi \in \mathsf{Th}(A)$ by "For all closed sets C containing A, $\varphi \in C$," which is Π^1_1.

Suppose that we extend our notion of logical consequence to incorporate (2), and let Th' denote closure under this notion. We can build $\mathsf{Th}'(A)$ from below and above as before, but when building from below, we must iterate transfinitely (since we will have an infinitary inference rule for (2)); in particular, we cannot express $\varphi \in \mathsf{Th}'(A)$ with a Σ^0_1 formula. We must resort to the Π^1_1 definition involving intersections of closed sets.

In both instances, incorporating (2) results in a loss of compactness, breaking whatever Σ^0_1 definition of logical consequence we had.

If we only incorporate (2) in a restricted way, we can end up with a complexity between Σ^0_1 and Π^1_1. A *semisimple* Horn formula will be like a simple Horn formula, except that $*$ may appear anywhere in the conclusion. The validity of such formulas will rely on (2), but only slightly, in that (2) is only used to convert a semisimple Horn formula into an infinite conjuction of simple Horn formulas; the question of validity of such conjunctions is Π^0_2.

These will be our main results: when we restrict to simple Horn formulas, the Horn theories of $*$-continuous and relational Kleene algebras are Σ^0_1-complete; when we restrict to semisimple Horn formulas, the Horn theories are Π^0_2-complete.

2 Preliminaries

2.1 Kleene Algebra

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

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

(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 let* IS *denote the class of all idempotent semirings.*

We often drop $\cdot$, writing xy for $x \cdot y$. The upper semilattice structure induces a natural partial order on any idempotent semiring: $x \le y \Leftrightarrow x + y = y$.

$+$ 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 $xy \le x'y'$.)

[1] This might seem circular because our definition of closed was in terms of Th, but it is easy to show that a set is closed iff it contains all axioms and is closed under applications of inference rules.

The names of the several classes of algebras we consider will serve as convenient abbreviations for the type of algebra they contain. For example, "Every IS extends..." would mean "Every idempotent semiring extends...". We use the notation $\mathsf{Ax}(\mathsf{IS})$ to denote the idempotent semiring axioms.

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*

$$1 + xx^* \le x^* \tag{3}$$

$$1 + x^*x \le x^* \tag{4}$$

$$p + qx \le x \to q^*p \le x \tag{5}$$

$$p + xq \le x \to pq^* \le x \tag{6}$$

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

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.*

Equation (3) implies that q^*p is a solution to the inequality $p + qx \le x$, and (5) implies that it is the least solution; (4) and (6) say that pq^* is the least solution to $p + xq \le x$.

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 $\mathbf{C}$ of algebras, $\mathbf{C} \models \varphi$ means that $K \models \varphi$ for each $K \in \mathbf{C}$; 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.*

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

We let $\mathsf{REG}\ M$ *denote the smallest subalgebra of* 2^M *containing the singletons* $\{x\}$, $x \in M$. *(The elements of* $\mathsf{REG}\ M$ *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 is the identity). 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*).*

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.*

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

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 class of all relational Kleene algebras.*

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 class of all* $*$*-continuous Kleene algebras.*

As in the introduction, this is equivalent to the bigness condition (1) and the smallness condition (2). Because first-order logic cannot be extended to accommodate formulas such as (2) without breaking compactness, it is not surprising that compactness is well suited for violating $*$-continuity, as shown by the proof of the following proposition.

Proposition 6. *There is a Kleene algebra that is not* $*$*-continuous.*

Proof. Let $\Phi = \mathsf{Ax}(\mathsf{KA}) \cup \{1 < x,\ 1 + a < x,\ 1 + a + a^2 < x, \ldots\} \cup \{x < a^*\}$. Φ is finitely satisfiable, so it is satisfiable. Any model of Φ is a Kleene algebra, and x will witness that a^* is not the least upper bound for $\{a^n \mid n \in \omega\}$. $\square$

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

Lemma 7. *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 [9, Lemma 7.1, pp. 246–248]. □

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 $\mathsf{RKA} \subseteq \mathsf{KA}^*$.

Lemma 8. *For any monoid* M, 2^M *and its subalgebras are isomorphic to relational Kleene algebras. In particular,* $\mathsf{REG}\ M$ *is isomorphic to a relational Kleene algebra.*

Proof. Define $\varphi : 2^M \to \mathcal{R}(M)$ by

$$\varphi(A) = \{(x, xy) \ \mid \ x \in M,\ y \in A\} \ .$$

It is straightforward to show that φ is an injective homomorphism. So, 2^M (or any subalgebra of 2^M) is isomorphic to its image under φ. □

Definition 9. *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* $\mathbf{C}$ *of algebras is the* universal Horn theory *of* $\mathbf{C}$*, which we denote by* $\mathcal{H}\mathbf{C}$.

Note that, because any inequality $x \leq y$ is in fact an equation $x + y = y$, inequalities are allowed in Horn formulas.

Despite the lack of quantifiers in our presentation, universal Horn formulas are in fact universal statements, at least when speaking of their validity. (For example, $K \models p \leq 1 \to p^2 = p \iff K \models \forall x(x \leq 1 \to x^2 = x)$.) The missing quantifier is hiding in our definition of $K \models \varphi$. We will often drop the word "universal".

Horn formulas are very important in universal algebra—second only to equations, as formulas go—but take on particular importance in Kleene algebra: the hypotheses of a Horn formula are used to capture (or partially capture) the intended semantics of the atomic program symbols when reasoning about programs. (For example, if p is intended to mean "`let x := 1`" and q is intended to mean "`let y := 1`", we might wish to reason under hypotheses such as $pq = qp$, $p^2 = p$, and $q^2 = q$.)

Proposition 10. $\mathcal{H}\mathsf{KA} \subsetneq \mathcal{H}\mathsf{KA}^* \subsetneq \mathcal{H}\mathsf{RKA}$

Proof. See Sect. A.1 of appendix. □

Despite the above proposition, there are many special cases in which these Horn theories coincide. The following lemma, which will be useful for other reasons, is one such example.

Lemma 11. *Suppose $M = \Sigma^*/E$ is a finitely presented monoid (where Σ is the set of generators, and $E = \{\sigma_1 = \tau_1, \ldots, \sigma_n = \tau_n\}$ is a set of equations, with $\sigma_i, \tau_i \in \Sigma^*$). Let $J : \Sigma^* \to M$ be the interpretation mapping each element of Σ^* to its equivalence class in M. The following are equivalent for any $\sigma, \tau \in \Sigma^*$.*

(i) $M, J \models \sigma = \tau$
(ii) $E \to \sigma = \tau$ *is valid in all monoids.*
(iii) $\mathsf{KA} \models E \to \sigma = \tau$
(iv) $\mathsf{KA}^* \models E \to \sigma = \tau$
(v) $\mathsf{RKA} \models E \to \sigma = \tau$
(vi) $\mathsf{KA} \models E \to \sigma \leq \tau$
(vii) $\mathsf{KA}^* \models E \to \sigma \leq \tau$
(viii) $\mathsf{RKA} \models E \to \sigma \leq \tau$

Proof. See Sect. A.1 of appendix. □

Lemma 12. $\mathcal{H}\mathsf{RKA}$*,* $\mathcal{H}\mathsf{KA}^*$*, and* $\mathcal{H}\mathsf{KA}$*, restricted to formulas containing only monoid equations (that is, equations whose terms are built from atomic program symbols, 1, and $\cdot$), are each Σ^0_1-complete.*

Proof. The word problem for finitely presented monoids, known to be Σ^0_1-complete, is exactly the same as determining the validity (over all monoids) of Horn formulas consisting of monoid equations. By Lemma 11, this is equivalent to determining the validity of such Horn formulas in KA, KA^*, or RKA. □

2.2 Kleene Algebra with Tests

Definition 13. *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 class of all Kleene algebras with tests; we let* KAT^* *denote the subclass of all $*$-continuous Kleene algebras with tests.*

Now, instead of just having atomic program symbols, we must also have symbols to use for 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 subalgebraof $\mathcal{R}(X)$ for some X. We let RKAT *denote the class of all relational Kleene algebras with tests.*

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 $\mathsf{RKAT} \models \varphi \Leftrightarrow \mathsf{RKA} \models \varphi$.

Definition 14. *Let* $\mathsf{Atoms}_\mathsf{B}$ *denote the atoms of the free Boolean algebra on generators* B*, which we can treat as elements of* $\mathsf{RExp}_{\mathsf{P},\mathsf{B}}$ *in a canonical way by ordering* B *and writing elements of* $\mathsf{Atoms}_\mathsf{B}$ *as conjunctions in which exactly one of* b *and* $\overline{b}$ *appears for each* $b \in \mathsf{B}$*, in order. For example, if* $\mathsf{B} = \{a, b\}$*, then* $\mathsf{Atoms}_\mathsf{B} = \{ab, a\overline{b}, \overline{a}b, \overline{a}\overline{b}\}$*. We use letters* $\alpha, \beta, \gamma, \delta$ *to denote elements of* $\mathsf{Atoms}_\mathsf{B}$*.*

A guarded string *over* P *and* B *is an element of* $\mathsf{RExp}_{\mathsf{P},\mathsf{B}}$ *of the form*

$$\alpha_0 p_1 \alpha_1 \cdots p_k \alpha_k \ ,$$

where $k \geq 0$*,* $\alpha_i \in \mathsf{Atoms}_\mathsf{B}$*,* $p_i \in \mathsf{P}$*. Let* $\mathrm{GS}_{\mathsf{P},\mathsf{B}}$ *(or simply* GS*) denote the set of all guarded strings on* P *and* B*.*

We define a partial binary operation $\diamond$ *on* GS *by*

$$\alpha\sigma\beta \diamond \gamma\tau\delta = \begin{cases} \alpha\sigma\beta\tau\delta & \text{if } \beta = \gamma; \\ \text{undefined} & \text{otherwise.} \end{cases}$$

The powerset 2^{GS} *of* GS *forms a Kleene algebra with tests as follows. The tests are the subsets of* $\mathsf{Atoms}_\mathsf{B}$*, and*

$$\begin{aligned} 0 &= \varnothing \\ 1 &= \mathsf{Atoms}_\mathsf{B} \\ A + B &= A \cup B \\ A \cdot B &= \{\sigma \diamond \tau \ \mid \ \sigma \in A,\ \tau \in B, \text{ and } \sigma \diamond \tau \text{ is defined}\} \\ A^* &= \bigcup_{k \in \omega} A^k \\ \overline{A} &= \mathsf{Atoms}_\mathsf{B} - A \end{aligned}$$

The canonical interpretation $G : \mathsf{RExp}_{\mathsf{P},\mathsf{B}} \to 2^{\mathrm{GS}}$ *is defined by*

$$\begin{aligned} G(p) &= \{\alpha p \beta \ \mid \ \alpha, \beta \in \mathsf{Atoms}_\mathsf{B}\} \quad \text{for } p \in \mathsf{P} \ , \\ G(b) &= \{\alpha \ \mid \ \alpha \leq b\}^2 \quad \text{for } b \in \mathsf{B} \ , \end{aligned}$$

extended homomorphically. Let $\mathsf{REG}\ \mathrm{GS}$ *denote the elements of* 2^{GS} *which are* $G(s)$ *for some* $s \in \mathsf{RExp}_{\mathsf{P},\mathsf{B}}$*.*

[2] The partial order here is the natural partial order on the free Boolean algebra on generators B. In this case, $\alpha \leq b$ iff b appears positively in α.

REG GS is also called the *guarded string model*, and is the free KAT on generators P, B in the sense that $G(s) = G(t)$ iff $\mathsf{KAT} \models s = t$ [10]. The guarded string model can be treated in more generality as a special case of a *trace model* [6], which we do not define here.

Lemma 15. *Let $K \in \mathsf{KAT}^*$, $I : \mathsf{RExp}_{\mathsf{P},\mathsf{B}} \to K$ an interpretation, and $p, q, r \in \mathsf{RExp}_{\mathsf{P},\mathsf{B}}$. Then*

$$I(pqr) = \sup_{\sigma \in G(q)} I(p\sigma r) \ .$$

In particular, $I(q) = \sup_{\sigma \in G(q)} I(\sigma)$.

Proof. See [10]. □

2.3 Complete Idempotent Semirings

A $*$-continuous Kleene algebra can be thought of as an idempotent semiring where certain suprema are guaranteed to exist ($\sup_n q^n = q^*$), with multiplication distributing over these suprema ($pq^*r = \sup_n pq^n r$). If we strengthen this to require arbitrary suprema to exist, with multiplication distributing over these arbitrary suprema, we get the notion of *complete idempotent semiring*.[3]

The algebras we consider in this section will all have a Boolean subalgebra (that is, they will be "with tests"), but all the results still hold without tests.

Definition 16. *An* idempotent semiring with tests *is a two-sorted structure $(S, B, +, \cdot, \overline{}, 0, 1)$, where $(S, +, \cdot, 0, 1)$ is an idempotent semiring, and $(B, +, \cdot, \overline{}, 0, 1)$ is a Boolean subalgebra. The elements of B are called* tests*. We let IST denote the class of all idempotent semirings with tests.*

For a finite set P of atomic programs and a finite set B of atomic tests, $\mathsf{RExp}^0_{\mathsf{P},\mathsf{B}}$ is the set of IST terms over P and B, i.e., the $*$-free terms in $\mathsf{RExp}_{\mathsf{P},\mathsf{B}}$.

An $S \in \mathsf{IST}$ is complete *if the partial order on S induced by $+$ is complete*[4] *(We do not require the supremum of an arbitrary set of tests to be a test.)*

$\mathsf{RExp}^0_{\mathsf{P},\mathsf{B}}$ is the set of IST terms over P and B; this coincides with the $*$-free terms of $\mathsf{RExp}_{\mathsf{P},\mathsf{B}}$, and we have $\mathsf{RExp}^0_{\mathsf{P},\mathsf{B}} \subseteq \mathsf{RExp}_{\mathsf{P},\mathsf{B}}$.

A complete IST forms a $*$-continuous Kleene algebra with tests by defining $x^* = \sup_{n \in \omega} x^n$.

Lemma 17. *Given a monoid $(S, \cdot, 1)$ with a complete partial order $\leq$ such that $\cdot$ distributes over arbitrary suprema, S forms a complete idempotent semiring by defining $0 = \sup \varnothing$ and $x + y = \sup\{x, y\}$. If S has a Boolean subalgebra (with $0, 1, +, \cdot$ coinciding with the operations on S), then S forms a complete IST.*

[3] Complete idempotent semirings are often referred to as S-algebras [4].

[4] Here, a partial order $(P, \leq)$ is complete if every subset of P has a supremum, as in [1]. In particular, we require $\varnothing$ and P to have suprema; this is in contrast to other existing notions of complete partial order.

Proof. Trivial, except perhaps for the requirement that 0 is an annihilator for $\cdot$, which follows from distributivity and the definition of 0:

$$0 \cdot x = (\sup \varnothing) \cdot x = \sup(\varnothing \cdot x) = \sup \varnothing = 0 \ .$$

$x \cdot 0 = 0$ is similar. □

Theorem 18. *Every* IST *extends to a complete* IST*.*

Proof. We use ideal completion, where an ideal I is a nonempty subset of the semiring which is closed downward and closed under $+$. See Sect. A.2 of appendix for details.

(Note that if a given $S \in \mathsf{IST}$ happens to be a Kleene algebra, this construction will not typically respect *. However, if S is a $*$-continuous Kleene algebra, one can strengthen the notion of ideal to require $pq^*r \in I$ whenever $pq^nr \in I$ for all n, and using ideal completion with this notion of ideal will preserve *. For details of this see [7].) □

Corollary 19. $\mathcal{H}\mathsf{KAT}$*,* $\mathcal{H}\mathsf{KAT}^*$*, and* $\mathcal{H}\mathsf{IST}$*, restricted to* $*$*-free formulas, coincide.*

Proof. $\mathsf{KAT}^* \subseteq \mathsf{KAT} \subseteq \mathsf{IST}$, so any Horn formula valid over IST is valid over KAT and KAT^*. If a Horn formula is valid over KAT^*, then it must be valid over IST, since every IST extends to a complete IST, which is a $*$-continuous Kleene algebra, and validity of Horn formulas is preserved in subalgebras. □

3 Simple and Semisimple Horn Formulas

Definition 20. *A Horn formula* φ *of* KAT *is* semisimple *if it is of the form*

$$s_1 \le t_1 \wedge \cdots \wedge s_n \le t_n \to s \le t$$

where $t_1, \ldots, t_n$ *are* $*$*-free (i.e., have no occurrence of* $*$*). We say that* φ *is* simple *if, in addition,* s *is* $*$*-free.*

The coercion *of a Horn formula* $E \to s \le t$ *is the formula* $E \wedge p \le s \wedge t \le p' \to p \le p'$ *where* p *and* p' *are fresh atomic program symbols.* $E \to s \le t$ *is* coerced *if* s *and* t *are atomic program symbols (in particular, the coercion of a formula is coerced).*

Lemma 21. *A Horn formula* $E \to s \le t$ *is valid over any particular* $K \in \mathsf{KAT}$ *iff its coercion is. (If the formula is* $*$*-free, then validity is also preserved in any* IST*.)*

Proof. See Sect. A.3 of appendix. □

Note that the coercion of any simple Horn formula is simple, so when working with simple Horn formulas, we will often assume without loss of generality that they are coerced.

Note that there are many Horn formulas which are not simple, but are equivalent to simple formulas or conjunctions thereof. For example, $p \leq 1 \to p^2 = p$ is not simple, but is equivalent to

$$(p \leq 1 \to p^2 \leq p) \wedge (p \leq 1 \to p \leq p^2) \ .$$

A more subtle example is $q^* = 1 \to q \leq 1$. If we expand $q^* = 1$ to $q^* \leq 1 \wedge 1 \leq q^*$, we still do not have a simple formula; however, the latter hypothesis is a KA tautology, so it can be dropped, leaving us with the simple formula

$$q^* \leq 1 \to q \leq 1 \ .$$

Our goal is the following four theorems.

Theorem 22. $\mathcal{H}\mathsf{KAT}^*$*, restricted to simple Horn formulas, is* Σ^0_1*-complete.*

Theorem 23. $\mathcal{H}\mathsf{RKAT}$*, restricted to simple Horn formulas, is* Σ^0_1*-complete.*

Theorem 24. $\mathcal{H}\mathsf{KAT}^*$*, restricted to semisimple Horn formulas, is* Π^0_2*-complete.*

Theorem 25. $\mathcal{H}\mathsf{RKAT}$*, restricted to semisimple Horn formulas, is* Π^0_2*-complete.*

(In each case, the lower bound will not require tests, so the results also apply to $\mathcal{H}\mathsf{KA}^*$ and $\mathcal{H}\mathsf{RKA}$.)

In [8] and [5], the reduction used to show that $\mathcal{H}\mathsf{KA}^*$ and $\mathcal{H}\mathsf{RKA}$ are Π^1_1-complete uses formulas that are equivalent to simple formulas except for a single occurrence of $*$ on the right-hand side of one hypothesis. (Furthermore, they have no occurrence of 0, 1, or $+$, and only the one occurrence of $*$.) Semisimple formulas are as general as possible without allowing $*$ on the right-hand side of a hypothesis, which in turn allows for Π^1_1-completeness, so if we wish to find any larger fragments of $\mathcal{H}\mathsf{KA}^*$ or $\mathcal{H}\mathsf{RKA}$ that are not Π^1_1-complete, the criteria will have to be more discriminating than simply where $*$ occurs.

3.1 Simple and Semisimple Formulas in KAT^*

In KAT, first-order logic can handle $*$ well, because its axiomatization is first order. The $*$-continuity axiom ($xy^*z = \sup_{n\in\omega} xy^nz$) is not first order, though, and the fact that $\mathcal{H}\mathsf{KAT}^*$ is Π^1_1-complete shows that there is no hope of finding a first-order substitute. As we will see in this section, the notion of simple formula captures the portion of $\mathcal{H}\mathsf{KAT}^*$ that first-order logic can (indirectly) handle anyway.

The $*$-continuity axiom, when the definition of supremum is unravelled, becomes

$$xy^*z \leq w \Leftrightarrow \bigwedge_{n\in\omega} xy^nz \leq w \ ,$$

or, using Lemma 15 (and abusing notation),

$$s \leq t \Leftrightarrow \bigwedge_{\sigma\in G(s)} \sigma \leq t \ .$$

Our basic tactic will be to replace any hypothesis $s \leq t$, where t is $*$-free, with the infinite set of $*$-free hypotheses $\sigma \leq t$, $\sigma \in G(s)$, which first-order logic can better digest.

Fix a simple Horn formula φ of the form

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

and assume without loss of generality that φ is coerced, so that $s_1, \ldots, s_n$ are the only terms that may contain $*$. Let

$$\Gamma_\varphi = \{\sigma \leq t_i \ \mid \ \sigma \in G(s_i),\ 1 \leq i \leq n\} \ .$$

Lemma 26. *For φ as above, the following are equivalent.*

(i) $\mathsf{KAT} \models \varphi$
(ii) $\mathsf{KAT}^* \models \varphi$
(iii) $\mathsf{Ax}(\mathsf{IST}) \cup \Gamma_\varphi \models s \leq t$ *($\mathsf{Ax}(\mathsf{IST})$ denotes the IST axioms.)*

Proof. (i)⇒(ii) is immediate, since $\mathsf{KAT}^* \subseteq \mathsf{KAT}$.

Suppose (ii) holds. Let $S \in \mathsf{IST}$, and let $I : \mathsf{RExp}^0_{\mathsf{P,B}} \to S$ be any interpretation such that $S, I \models \Gamma_\varphi$; we must show $S, I \models s \leq t$. By Theorem 18, S extends to a complete semiring S'. I extends uniquely to an interpretation $I' : \mathsf{RExp}_{\mathsf{P,B}} \to S'$. For any $\sigma \in G(s_i)$, $1 \leq i \leq n$, we have $I'(\sigma) = I(\sigma) \leq I(t_i) = I'(t_i)$, so by Lemma 15, $I'(s_i) = \sup_{\sigma \in G(s_i)} I'(\sigma) \leq I'(t_i)$. So, since $S', I' \models \varphi$ by assumption, and we have just shown that S', I' satisfies each hypothesis of φ, we must have $S', I' \models s \leq t$. Then $I(s) = I'(s) \leq I'(t) = I(t)$, so $S, I \models s \leq t$, giving us (iii).

Now suppose (iii) holds. Take any $K \in \mathsf{KAT}$ and interpretation $I : \mathsf{RExp}_{\mathsf{P,B}} \to K$. Suppose $I(s_i) \leq I(t_i)$ for $1 \leq i \leq n$. Then for any $\sigma \in G(s_i)$, $1 \leq i \leq n$, we have $I(\sigma) \leq I(s_i) \leq I(t_i)$, so $K, I \models \Gamma_\varphi$. We also have $K, I \models \mathsf{Ax}(\mathsf{IST})$, so by (iii), $K, I \models s \leq t$. Therefore, $K, I \models \varphi$, giving us (i). □

Proof (of Theorem 22). The upper bound follows from Lemma 26, since the Horn theory of KAT is Σ^0_1. (The entire theory of KAT is Σ^0_1 since it is finitely axiomatized.)

The lower bound is by Lemma 12. (Although Horn formulas consisting of monoid equations are not technically simple, each hypothesis can be replaced by a pair of inequalities, and Lemma 11 shows that we can, in this case, replace the conclusion with an inequality.) □

Proof (of Theorem 24). Essentially, $*$-continuity lets us treat a semisimple formula $E \to s \leq t$ as the infinite conjunction $\bigwedge_{\sigma \in G(s)} E \to \sigma \leq t$. See Sect. A.3 of the appendix for details. □

Theorem 27. *Let $s_1 \leq t_1 \wedge \cdots \wedge s_n \leq t_n \to s \leq t$ be any simple Horn formula. The following are equivalent.*

(i) $\mathsf{KAT}^* \models s_1 \leq t_1 \wedge \cdots \wedge s_n \leq t_n \to s \leq t$

(ii) *There exist finite sets $T \subseteq G(t)$ and $S_i \subseteq G(s_i)$, $1 \leq i \leq n$, such that*

$$\mathsf{KAT}^* \models (\Sigma S_1) \leq t_1 \wedge \cdots \wedge (\Sigma S_n) \leq t_n \rightarrow s \leq (\Sigma T) \ .$$

(iii) *There exist finite sets $T \subseteq G(t)$ and $S_i \subseteq G(s_i)$, $1 \leq i \leq n$, such that*

$$\mathsf{IST} \models (\Sigma S_1) \leq t_1 \wedge \cdots \wedge (\Sigma S_n) \leq t_n \rightarrow s \leq (\Sigma T) \ .$$

(Here, ΣS denotes the sum of the elements of S.)

Proof. This is just an application of first-order compactness to Lemma 26. See Sect. A.3 of the appendix for details. □

3.2 Simple and Semisimple Formulas in RKAT

The ideas of Section 3.1 also work for relational algebras, but we must use first-order logic differently.

Fix finite $\mathsf{P} = \{p_1, \ldots, p_m\}$, $\mathsf{B} = \{b_1, \ldots, b_\ell\}$. Let $\mathcal{L}_{\mathsf{P},\mathsf{B}}$ be the first-order language with binary predicate symbols $P_1, \ldots, P_m, B_1, \ldots, B_\ell$ (in addition to equality) and no function or constant symbols. Let β be the formula

$$\bigwedge_{i=1}^{\ell} \forall x, y[B_i(x, y) \rightarrow x = y] \ ,$$

which will ensure that interpretations of the B_i make suitable Boolean elements in a relational algebra, by making them subsets of the identity relation.

For any $\mathcal{L}_{\mathsf{P},\mathsf{B}}$-structure $\mathcal{A}$ modeling β (or simply, "for any $\mathcal{A} \models \beta$"), $|\mathcal{A}|$ will denote the universe of $\mathcal{A}$, while $P_1^{\mathcal{A}}, \ldots, P_m^{\mathcal{A}}, B_1^{\mathcal{A}}, \ldots, B_\ell^{\mathcal{A}}$ will denote the interpretations of $P_1, \ldots, P_m, B_1, \ldots, B_\ell$ in $\mathcal{A}$, and we define the interpretation $I^{\mathcal{A}} : \mathsf{RExp}_{\mathsf{P},\mathsf{B}} \rightarrow \mathcal{R}(|\mathcal{A}|)$ by $I^{\mathcal{A}}(p_i) = P_i^{\mathcal{A}}$, $I^{\mathcal{A}}(b_i) = B_i^{\mathcal{A}}$.

For each $t \in \mathsf{RExp}^0_{\mathsf{P},\mathsf{B}}$, we define the formula $\theta_t(x, y)$ of $\mathcal{L}_{\mathsf{P},\mathsf{B}}$ by induction on t as follows.

$$\begin{aligned} \theta_0(x, y) &\Leftrightarrow \textbf{false} & \theta_{\overline{t}}(x, y) &\Leftrightarrow x = y \wedge \neg\theta_t(x, y) \\ \theta_1(x, y) &\Leftrightarrow x = y & \theta_{s+t}(x, y) &\Leftrightarrow \theta_s(x, y) \vee \theta_t(x, y) \\ \theta_{p_i}(x, y) &\Leftrightarrow P_i(x, y) & \theta_{st}(x, y) &\Leftrightarrow \exists z[\theta_s(x, z) \wedge \theta_t(z, y)] \\ \theta_{b_i}(x, y) &\Leftrightarrow B_i(x, y) & & \end{aligned}$$

(Given a formula $\varphi(x, y)$, when we write $\varphi(x, z)$, the usual convention applies: if the variable z already occurs in $\varphi(x, y)$, it is renamed as necessary before substituting z for y to get $\varphi(x, z)$, to avoid variable capture.)

Lemma 28. *Take any $\mathcal{A} \models \beta$. Then for all $a, a' \in |\mathcal{A}|$ and $t \in \mathsf{RExp}^0_{\mathsf{P},\mathsf{B}}$,*

$$(a, a') \in I^{\mathcal{A}}(t) \Leftrightarrow \mathcal{A} \models \theta_t(a, a') \ .$$

Proof. Straightforward induction on t. □

Now, for any inequality $s \le t$ for $s, t \in \mathsf{RExp}^0_{\mathsf{P},\mathsf{B}}$, we define the sentence $\theta_{s\le t}$ by

$$\theta_{s\le t} \Leftrightarrow \forall x, y[\theta_s(x,y) \to \theta_t(x,y)] \ .$$

Lemma 29. *Take any $\mathcal{A} \models \beta$. Then for all $s, t \in \mathsf{RExp}^0_{\mathsf{P},\mathsf{B}}$,*

$$\mathcal{R}(|\mathcal{A}|), I^{\mathcal{A}} \models s \le t \Leftrightarrow \mathcal{A} \models \theta_{s\le t} \ .$$

Proof. Apply previous lemma. See Sect. A.4 of appendix for details. □

As we did in Section 3.1, fix a simple Horn formula φ of the form

$$s_1 \le t_1 \wedge \cdots \wedge s_n \le t_n \to s \le t \ ,$$

assuming without loss of generality that s and t are $*$-free, and let $\Gamma_\varphi = \{\sigma \le t_i \ \mid \ \sigma \in G(s_i),\ 1 \le i \le n\}$. Let

$$\widehat{\Gamma}_\varphi = \{\theta_{u\le v} \ \mid \ u \le v \text{ appears in } \Gamma_\varphi\} \ .$$

Lemma 30. *The following are equivalent.*

(i) $\mathsf{RKAT} \models \varphi$
(ii) $\{\beta\} \cup \widehat{\Gamma}_\varphi \models \theta_{s\le t}$

Proof. Suppose (i) holds. Suppose $\mathcal{A} \models \{\beta\} \cup \widehat{\Gamma}_\varphi$. Then by Lemma 29, $\mathcal{R}(|\mathcal{A}|), I^{\mathcal{A}} \models u \le v$ for all $u \le v$ appearing in Γ_φ; that is, $\mathcal{R}(|\mathcal{A}|), I^{\mathcal{A}} \models \sigma \le t_i$ for all $\sigma \in G(s_i)$, $1 \le i \le n$. $\mathcal{R}(|\mathcal{A}|)$ is $*$-continuous, so by Lemma 15, $\mathcal{R}(|\mathcal{A}|), I^{\mathcal{A}} \models s_i \le t_i$, $1 \le i \le n$. So, applying (i), $\mathcal{R}(|\mathcal{A}|), I^{\mathcal{A}} \models s \le t$. Applying Lemma 29 again, we have $\mathcal{A} \models \theta_{s\le t}$, giving us (ii).

Now suppose (ii) holds. Take any $K \in \mathsf{RKAT}$. Let X be the base of K. Suppose $I : \mathsf{RExp}_{\mathsf{P},\mathsf{B}} \to K$ is an interpretation such that $K, I \models s_i \le t_n$, $1 \le i \le n$. Define the $\mathcal{L}_{\mathsf{P},\mathsf{B}}$-structure $\mathcal{A}$ by $|\mathcal{A}| = X$, $P_i^{\mathcal{A}} = I(p_i)$, $B_i^{\mathcal{A}} = I(b_i)$. Then $I = I^{\mathcal{A}}$, so from $K, I \models s_i \le t_i$, $1 \le i \le n$, we get $\mathcal{R}(|\mathcal{A}|), I^{\mathcal{A}} \models s_i \le t_i$, $1 \le i \le n$. It follows that $\mathcal{R}(|\mathcal{A}|), I^{\mathcal{A}} \models \sigma \le t_i$ for all $\sigma \in G(s_i)$, $1 \le i \le n$; that is, $\mathcal{R}(|\mathcal{A}|), I^{\mathcal{A}} \models \Gamma_\varphi$. Thus, by Lemma 29, $\mathcal{A} \models \widehat{\Gamma}_\varphi$. Since each $I(b_i)$ is a subset of the identity relation, we have $\mathcal{A} \models \beta$. So, applying (ii), $\mathcal{A} \models \theta_{s\le t}$. Applying Lemma 29 again, we have $\mathcal{R}(|\mathcal{A}|), I^{\mathcal{A}} \models s \le t$. This gives us $K, I \models s \le t$ (since $I = I^{\mathcal{A}}$). We now have (i), as desired. □

Proof (of Theorem 23). The upper bound comes from Lemma 30, since the predicate $\{\beta\} \cup \widehat{\Gamma}_\varphi \models \theta_{s\le t}$ is Σ^1_0 in φ. (The set $\{\beta\} \cup \widehat{\Gamma}_\varphi$ is uniformly computable in φ, so we can effectively enumerate its logical consequences.) The lower bound comes from Lemma 12. □

Lemma 31. *Let E be a finite set of monoid equations over Σ. Then for any $s \in \mathsf{RExp}_\Sigma$ and $\tau \in \Sigma^*$, the following are equivalent.*

(i) $\mathsf{KA}^* \models E \to s \le \tau$
(ii) $\mathsf{RKA} \models E \to s \le \tau$

Proof. See Sect. A.4 of appendix. □

Proof (of Theorem 25). The upper bound is by the same argument as in the proof of Theorem 24. The lower bound, also as in the proof of Theorem 24, is from the reduction in [8]; we can use the same reduction because the formulas in the reduction are of the form that Lemma 31 applies to. □

Theorem 32. *Let $s_1 \le t_1 \wedge \cdots \wedge s_n \le t_n \to s \le t$ be any simple Horn formula. The following are equivalent.*

(i) $\mathsf{RKAT} \models s_1 \le t_1 \wedge \cdots \wedge s_n \le t_n \to s \le t$
(ii) *There exist finite sets $T \subseteq G(t)$ and $S_i \subseteq G(s_i)$, $1 \le i \le n$, such that*

$$\mathsf{RKAT} \models (\Sigma S_1) \le t_1 \wedge \cdots \wedge (\Sigma S_n) \le t_n \to s \le (\Sigma T) \ .$$

(iii) *There exist finite sets $T \subseteq G(t)$ and $S_i \subseteq G(s_i)$, $1 \le i \le n$, such that*

$$\beta \models \theta_{(\Sigma S_1) \le t_1} \wedge \cdots \wedge \theta_{(\Sigma S_n) \le t_n} \to \theta_{s \le (\Sigma T)} \ .$$

Proof. Similar to Theorem 27. □

4 Bigness and Smallness of $*$

In the introduction, we stated that the validity of simple Horn formulas depends only on the bigness of $*$. We now make that precise.

Definition 33. *A* $*$-algebra *is an algebra over the signature of* KAT *satisfying the* IST *axioms (in other words, the result of dropping the $*$ axioms from* KAT*). A* big-$*$-algebra *is a $*$-algebra satisfying the bigness condition (1) from the introduction. A* small-$*$-algebra *is a $*$-algebra satisfying the smallness condition (2). We let* BIG^* *and* SMALL^* *denote the classes of big-$*$- and small-$*$-algebras, respectively.*

Clearly, $\mathsf{BIG}^* \cap \mathsf{SMALL}^* = \mathsf{KAT}^*$. We also have $\mathsf{KAT} \subseteq \mathsf{BIG}^*$, $\mathsf{KAT} \not\subseteq \mathsf{SMALL}^*$.

Theorem 34. *For any simple Horn formula φ, the following are equivalent.*

(i) $\mathsf{KAT}^* \models \varphi$
(ii) $\mathsf{BIG}^* \models \varphi$
(iii) $\mathsf{KAT} \models \varphi$

Proof. See Sect. A.5 of appendix. □

If we reverse the inequalities in the definition of simple formula, we get formulas whose validity (in KAT^*) depends only on the smallness of $*$.

Definition 35. *A Horn formula φ of* KAT *is* cosimple *if it is of the form*

$$s_1 \le t_1 \wedge \cdots \wedge s_n \le t_n \to s \le t$$

where $s_1, \ldots, s_n$ and t are $*$-free.

Theorem 36. *For any cosimple Horn formula φ, the following are equivalent.*

(i) $\mathsf{KAT}^* \models \varphi$
(ii) $\mathsf{SMALL}^* \models \varphi$

(Note the absence of the case $\mathsf{KAT} \models \varphi$*.)*

Proof. See Sect. A.5 of appendix. □

Acknowledgments

I am grateful to Dexter Kozen for his comments on various drafts of this article.

References

1. Stanley Burris and H. P. Sankappanavar. *A Course in Universal Algebra*. Springer-Verlag, New York, 1981. (Also available online: `http://www.thoralf.uwaterloo.ca/htdocs/ualg.html`)
2. Ernie Cohen. Hypotheses in Kleene Algebra. Unpublished.
3. Ernie Cohen, Dexter Kozen, and Frederick 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. Chris Hardin and Dexter Kozen. On the Complexity of the Horn Theory of REL. Technical Report 2003-1896, Computer Science Department, Cornell University, May 2003.
6. Chris Hardin and Dexter Kozen. On the Elimination of Hypotheses in Kleene Algebra with Tests. Technical Report 2002-1879, Computer Science Department, Cornell University, October 2002.
7. Dexter Kozen. On Kleene algebras and closed semirings. In Rovan, editor, *Proc. Math. Found. Comput. Sci.*, volume 452 of *Lect. Notes in Comput. Sci.*, pages 26–47. Springer, 1990.
8. Dexter Kozen. On the Complexity of Reasoning in Kleene Algebra. *Information and Computation* 179, 152-162, 2002.
9. Dexter Kozen. *The Design and Analysis of Algorithms*. Springer-Verlag, New York, 1991.
10. Dexter Kozen and Frederick Smith. Kleene algebra with tests: completeness and decidability. Proc. 10th Int. Workshop on Computer Science Logic (CSL'96), ed. D. van Dalen and M. Bezem, Utrecht, The Netherlands, Springer-Verlag Lecture Notes in Computer Science Volume 1258, September 1996, 244–259.

A Selected Proofs

A.1 Kleene Algebra

Proof (of Proposition 10). The inclusions are immediate, since $\mathsf{RKA} \subseteq \mathsf{KA}^* \subseteq \mathsf{KA}$.

$\mathcal{H}\mathsf{KA} \neq \mathcal{H}\mathsf{KA}^*$ because the former is Σ_1^0-complete, while the latter is Π_1^1 complete [8].

For $\mathcal{H}\mathsf{KA}^* \neq \mathcal{H}\mathsf{RKA}$, let φ be the formula

$$p \leq 1 \rightarrow p^2 = p \ .$$

In relational algebras, $p \leq 1$ holds iff p is a subset of the identity; among subsets of the identity, relational composition coincides with intersection, and it follows that $p^2 = p$. So $\mathsf{RKA} \models \varphi$. However, φ fails in the three element Kleene algebra $\{0, a, 1\}$ in which $0 < a < 1$ and $a^2 = 0$ (specifically, it fails under the interpretation $I(p) = a$). □

Proof (of Lemma 11). (i)⇔(ii) is standard. (ii)⇒(iii) is trivial, since every Kleene algebra forms a monoid under multiplication. (iii)⇒(iv)⇒(v) and (vi)⇒(vii)⇒(viii) are trivial, since $\mathsf{RKA} \subseteq \mathsf{KA}^* \subseteq \mathsf{KA}$.

The respective implications (iii)–(v)⇒(vi)–(viii) are trivial, since $\sigma = \tau$ is stronger than $\sigma \leq \tau$.

Suppose (viii). Define $h : M \to \mathcal{R}(M)$ by $h(y) = \{(x, xy) \mid x \in M\}$, which embeds M into the multiplicative monoid of $\mathcal{R}(M)$. $\mathcal{R}(M), h \circ J \models E$, so by (viii), $\mathcal{R}(M), h \circ J \models \sigma \leq \tau$. (Technically, $h \circ J$ is not an interpretation, since it is only defined on Σ^*, but it could be extended homomorphically to one.) So, we have $(1^M, J(\sigma)) \in h(J(\sigma)) \subseteq h(J(\tau))$, where 1^M denotes the identity of M. However, $(1^M, J(\tau))$ is the only pair of the form $(1^M, x)$ in $h(J(\tau))$, so $J(\sigma) = J(\tau)$, giving us (i). □

A.2 Ideal Completion for Idempotent Semirings

Fix an $S \in \mathsf{IST}$ with tests $B \subseteq S$.

Definition 37. *An* ideal over S *is a set* $I \subseteq S$ *such that*

(i) $0 \in I$;
(ii) *if* $x \in I$ *and* $y \leq x$*, then* $y \in I$;
(iii) *if* $x, y \in I$*, then* $x + y \in I$.

For any $x \in S$*, we let* $(x) = \{y \in S \mid y \leq x\}$*. We call* (x) *the* principal ideal generated by x*.*

We will show that the ideals of S form a complete semiring with tests, and that the map $x \mapsto (x)$ embeds S into this structure. Informally, this amounts to showing that our algebraic operations (particularly $\cdot$) interact nicely with our notion of ideal.

Let $K_S = \{I \subseteq S \mid I \text{ is an ideal}\}$, which we order by inclusion; let $B_S = \{(b) \mid b \in B\} \subseteq K_S$. $(0) = \{0\}$ is the least element of K_S.

Observe that S forms an ideal over itself, and that an arbitrary intersection of ideals is again an ideal, so given an arbitrary $X \subseteq S$, we can define

$$\mathrm{cl}(X) = \bigcap \{I \in K_S \mid X \subseteq I\} \in K_S \ ,$$

the least ideal containing X. We also define several other operations as follows.

$$\begin{aligned} \mathrm{cl}^+(X) &= \{y \in S \mid y \text{ is a finite sum}^5 \text{ of elements of } X\} \\ \mathrm{cl}{\downarrow}(X) &= \{y \in S \mid \exists x \in X\ (y \le x)\} \\ X \cdot Y &= \{xy \mid x \in X,\ y \in Y\} \\ X \odot Y &= \mathrm{cl}(X \cdot Y) \end{aligned}$$

Given $\mathcal{I} \subseteq K_S$, we define

$$\bigoplus \mathcal{I} = \mathrm{cl}(\cup \mathcal{I}) \ .$$

$\bigoplus \mathcal{I}$ gives the least upper bound of $\mathcal{I}$, witnessing that K_S is a complete partial order. For $I, J \in K_S$, we define

$$I \oplus J = \bigoplus \{I, J\} \ .$$

For $(b) \in B_S$, we define $\overline{(b)} = (\overline{b})$.

Lemma 38. *For $X, Y \subseteq S$, the following hold.*

(i) $X \subseteq \mathrm{cl}^+(X) \subseteq \mathrm{cl}(X)$
(ii) $X \subseteq \mathrm{cl}{\downarrow}(X) \subseteq \mathrm{cl}(X)$
(iii) $\mathrm{cl}(X) = \mathrm{cl}(\mathrm{cl}(X))$
(iv) *If* $X \subseteq Y$, *then* $\mathrm{cl}(X) \subseteq \mathrm{cl}(Y)$.
(v) $\mathrm{cl}{\downarrow}(\mathrm{cl}^+(X))$ *is an ideal.*
(vi) $\mathrm{cl}(X) = \mathrm{cl}{\downarrow}(\mathrm{cl}^+(X))$
(vii) $\mathrm{cl}(X \cdot Y) = \mathrm{cl}(\mathrm{cl}(X) \cdot Y) = \mathrm{cl}(X \cdot \mathrm{cl}(Y))$
(viii) *For $\mathcal{X} \subseteq \mathcal{P}(S)$ (the powerset of S),*

$$\bigcup_{X \in \mathcal{X}} \mathrm{cl}(X) \subseteq \mathrm{cl}(\cup \mathcal{X})$$

and

$$\mathrm{cl}\left(\bigcup_{X \in \mathcal{X}} \mathrm{cl}(X)\right) = \mathrm{cl}(\cup \mathcal{X}) \ .$$

[5] We take this to include empty sums, so $0 \in \mathrm{cl}^+(X)$ for all X.

(ix) *For* $\mathcal{X} \subseteq \mathcal{P}(S)$,

$$(\cup\mathcal{X}) \cdot Y = \bigcup_{X\in\mathcal{X}} (X \cdot Y)$$

and

$$Y \cdot (\cup\mathcal{X}) = \bigcup_{X\in\mathcal{X}} (Y \cdot X) \ .$$

Proof. (i)–(iv) Trivial.

(v) First observe that $\mathrm{cl}{\downarrow}\,(\mathrm{cl}^{+}(X))$ clearly contains 0 and is closed downward. For closure under sum, suppose $y_1, y_2 \in \mathrm{cl}{\downarrow}\,(\mathrm{cl}^{+}(X))$. Then there exist $x_1, x_2 \in \mathrm{cl}^{+}(X)$ such that $y_1 \leq x_1$, $y_2 \leq x_2$. $y_1 + y_2 \leq x_1 + x_2 \in \mathrm{cl}^{+}(X)$, so $y_1 + y_2 \in \mathrm{cl}{\downarrow}\,(\mathrm{cl}^{+}(X))$. Therefore $\mathrm{cl}{\downarrow}\,(\mathrm{cl}^{+}(X))$ is an ideal.

(vi) Follows from (i)–(v).

(vii) For the inclusion $\mathrm{cl}(\mathrm{cl}(X) \cdot Y) \subseteq \mathrm{cl}(X \cdot Y)$, suppose I is an ideal containing $X \cdot Y$. Take any $x_1, \ldots, x_k \in X$ and $y \in Y$; then $x_1y, \ldots, x_ky \in X \cdot Y \subseteq I$, so $(x_1 + \cdots + x_k)y = x_1y + \cdots + x_ky \in I$. Therefore, $\mathrm{cl}^{+}(X) \cdot Y \subseteq I$. Now take any $x \in \mathrm{cl}(X)$ and $y \in Y$; then by (vi), there exists $x_0 \in \mathrm{cl}^{+}(X)$ such that $x \leq x_0$, so $xy \leq x_0y \in I$, so $xy \in I$. Therefore, $\mathrm{cl}(X) \cdot Y \subseteq I$. This gives us $\mathrm{cl}(\mathrm{cl}(X) \cdot Y) \subseteq \mathrm{cl}(X \cdot Y)$, and the inclusion $\mathrm{cl}(X \cdot Y) \subseteq \mathrm{cl}(\mathrm{cl}(X) \cdot Y)$ is trivial, so $\mathrm{cl}(X \cdot Y) = \mathrm{cl}(\mathrm{cl}(X) \cdot Y)$. $\mathrm{cl}(X \cdot Y) = \mathrm{cl}(X \cdot \mathrm{cl}(Y))$ is similar.

(viii) The inclusion follows from the fact that any ideal containing $\cup\mathcal{X}$ contains each $X \in \mathcal{X}$ and hence contains $\mathrm{cl}(X)$ for each $X \in \mathcal{X}$. The equality follows from the inclusion, using (iii) and (iv).

(ix) Trivial, since $X \cdot Y$ is defined pointwise. □

Lemma 39. $\odot$ *distributes over arbitrary suprema; that is, for* $\mathcal{I} \subseteq K_S$ *and* $J \in K_S$,

$$\left(\bigoplus \mathcal{I}\right) \odot J = \bigoplus_{I\in\mathcal{I}} (I \odot J)$$

and

$$J \odot \left(\bigoplus \mathcal{I}\right) = \bigoplus_{I\in\mathcal{I}} (J \odot I) \ .$$

Proof.

$$\begin{aligned}\left(\bigoplus \mathcal{I}\right) \odot J &= \mathrm{cl}(\mathrm{cl}(\cup\mathcal{I}) \cdot J) \\ &= \mathrm{cl}((\cup\mathcal{I}) \cdot J) \\ &= \mathrm{cl}\left(\bigcup_{I\in\mathcal{I}} (I \cdot J)\right) \\ &= \mathrm{cl}\left(\bigcup_{I\in\mathcal{I}} \mathrm{cl}(I \cdot J)\right) \\ &= \mathrm{cl}\left(\bigcup_{I\in\mathcal{I}} (I \odot J)\right) \\ &= \bigoplus_{I\in\mathcal{I}} (I \odot J)\end{aligned}$$

$J \odot (\bigoplus \mathcal{I}) = \bigoplus_{I\in\mathcal{I}}(J \odot I)$ is similar. □

Lemma 40. $(K_S, B_S, \oplus, \odot, \overline{}, (0), (1))$ *forms a complete idempotent semiring with tests, and the map* $x \mapsto (x)$ *embeds* S *into this structure.*

Proof. For $H, I, J \in K_S$, $(1) \odot I = I \odot (1) = I$ is trivial, and

$$\begin{aligned}H \odot (I \odot J) &= \mathrm{cl}(H \cdot \mathrm{cl}(I \cdot J)) \\ &= \mathrm{cl}(H \cdot (I \cdot J)) \\ &= \mathrm{cl}((H \cdot I) \cdot J) \\ &= \mathrm{cl}(\mathrm{cl}(H \cdot I) \cdot J) \\ &= (H \odot I) \odot J\end{aligned}$$

so $(K_S, \odot, (1))$ forms a monoid. The other requirements of Lemma 17 have already been shown, so $(K_S, \oplus, \odot, (0), (1))$ forms a complete semiring.

$x \mapsto (x)$ is trivially an injection. It is easy to verify $(x) \oplus (y) = (x + y)$ and $(x) \odot (y) = (xy)$, and we have $\overline{(b)} = (\overline{b})$ by definition. Thus $x \mapsto (x)$ is an embedding of $(S, B, +, \cdot, \overline{}, 0, 1)$ into $(K_S, B_S, \oplus, \odot, \overline{}, (0), (1))$.

Note that $x \mapsto (x)$ induces a surjection of B onto B_S, and hence forms an isomorphism between $(B, +, \cdot, \overline{}, 0, 1)$ and $(B_S, \oplus, \odot, \overline{}, (0), (1))$, so the latter is in fact a Boolean algebra. Therefore $(K_S, B_S, \oplus, \odot, \overline{}, (0), (1))$ is a complete IST. □

Proof (of Theorem 18). Immediate from Lemma 40. □

A.3 Simple and Semisimple Horn Formulas

Proof (of Lemma 21). Let P, B contain all atomic programs and tests in $E \to s \leq t$. Let $\mathsf{P}' = \mathsf{P} \cup \{p, p'\}$ where p, p' are the fresh program symbols from the coercion.

Take any $K \in \mathsf{KAT}$. $(K \models E \to s \le t) \Rightarrow (K \models E \wedge p \le s \wedge t \le p' \to p \le p')$ is immediate. Now suppose $K \models E \wedge p \le s \wedge t \le p' \to p \le p'$. Given any interpretation $I : \mathsf{RExp}_{\mathsf{P},\mathsf{B}} \to K$, extend I to an interpretation $I' : \mathsf{RExp}_{\mathsf{P}',\mathsf{B}} \to K$ by letting $I'(p) = I(s)$, $I'(p') = I(t)$. If $K, I \models E$, then $K, I' \models E \wedge p \le s \wedge t \le p'$, so $K, I' \models p \le p'$; we then have $K, I \models s \le t$, since $I(s) = I'(p) \le I'(p') = I(t)$. This gives us $K \models E \to s \le t$.

The same argument applies to IST, when $E \to s \le t$ is $*$-free. □

Proof (of Theorem 24). Suppose φ is a semisimple formula $E \to s \le t$. For each $\sigma \in G(s)$, let φ_σ be the simple formula

$$E \to \sigma \le t \ .$$

If $\mathsf{KAT}^* \models \varphi$, then $\mathsf{KAT}^* \models \varphi_\sigma$ for all $\sigma \in G(s)$, since $\mathsf{KAT}^* \models \sigma \le s$ for all $\sigma \in G(s)$.

Now suppose that for all $\sigma \in G(s)$, $\mathsf{KAT}^* \models \varphi_\sigma$. Take any $K \in \mathsf{KAT}^*$, and any interpretation $I : \mathsf{RExp}_{\mathsf{P},\mathsf{B}} \to K$. If $K, I \not\models E$, then $K, I \models \varphi$. If $K, I \models E$, then $K, I \models \sigma \le t$ for all $\sigma \in G(s)$; it follows by Lemma 15 that $K, I \models s \le t$, so $K, I \models \varphi$.

Therefore, $\mathsf{KAT}^* \models \varphi$ iff for all $\sigma \in G(s)$, $\mathsf{KAT}^* \models \varphi_\sigma$. The latter is a Π^0_2 statement, since the predicate $\mathsf{KAT}^* \models \varphi_\sigma$ is Σ^0_1 by Theorem 22, and the set $G(s)$ is computable. This gives the upper bound.

In [8], by encoding the totality of Turing machines, it is shown that the universal Horn theory of KA^*, restricted to formulas whose hypotheses are monoid equations, is Π^0_2-complete, and such formulas are semisimple; this gives the lower bound. (Much like in the proof of Theorem 22, one might worry that these formulas are not truly semisimple, since their conclusions are not necessarily inequalities. However, the reduction in [8] uses formulas that are in fact semisimple. Alternatively, we could expand the definition of semisimple formula to allow any equation as the conclusion.) □

Proof (of Theorem 27). Let φ be the coercion of the formula in question:

$$s_1 \le t_1 \wedge \cdots \wedge s_n \le t_n \wedge p \le s \wedge t \le p' \to p \le p' \ ,$$

with p and p' fresh.

Suppose (i). Then by Lemma 26, $\mathsf{Ax}(\mathsf{IST}) \cup \Gamma_\varphi \models p \le p'$. By first-order compactness, there is some finite $\Gamma' \subseteq \Gamma_\varphi$ such that $\mathsf{Ax}(\mathsf{IST}) \cup \Gamma' \models p \le p'$. Keeping in mind how Γ_φ was defined, and observing that a conjunction of inequalities $x_1 \le y \wedge \cdots \wedge x_k \le y$ is equivalent to the single inequality $x_1 + \cdots + x_k \le y$, there must be some φ' of the form

$$(\Sigma S_1) \le t_1 \wedge \cdots \wedge (\Sigma S_n) \le t_n \wedge p \le s \wedge (\Sigma T) \le p' \to p \le p'$$

such that $\Gamma' = \Gamma_{\varphi'}$ (with $S_i \subseteq G(s_i)$, $T \subseteq G(t)$, all finite). We then have $\mathsf{Ax}(\mathsf{IST}) \cup \Gamma_{\varphi'} \models p \le p'$, which is equivalent to $\mathsf{IST} \models \varphi'$. φ' is the coercion of the formula

$$(\Sigma S_1) \le t_1 \wedge \cdots \wedge (\Sigma S_n) \le t_n \to s \le (\Sigma T) \ ,$$

so we have

$$\mathsf{IST} \models (\Sigma S_1) \leq t_1 \wedge \cdots \wedge (\Sigma S_n) \leq t_n \to s \leq (\Sigma T) \ .$$

This establishes (i)⇒(iii).

(iii)⇒(ii) is immediate since $\mathsf{KAT}^* \subseteq \mathsf{IST}$. (ii)⇒(i) is immediate since $\mathsf{KAT}^* \models (\Sigma S_i) \leq s_i$ and $\mathsf{KAT}^* \models (\Sigma T) \leq t$. □

A.4 Simple and Semisimple Formulas in RKAT

Proof (of Lemma 29). By Lemma 28,

$$\begin{aligned}
\mathcal{R}(|\mathcal{A}|), I^{\mathcal{A}} \models s \leq t &\Leftrightarrow I^{\mathcal{A}}(s) \subseteq I^{\mathcal{A}}(t) \\
&\Leftrightarrow \text{For all } a, a' \in |\mathcal{A}|, \text{ if } (a, a') \in I^{\mathcal{A}}(s), \text{ then } (a, a') \in I^{\mathcal{A}}(t). \\
&\Leftrightarrow \text{For all } a, a' \in |\mathcal{A}|, \text{ if } \mathcal{A} \models \theta_s(a, a'), \text{ then } \mathcal{A} \models \theta_t(a, a'). \\
&\Leftrightarrow \mathcal{A} \models \theta_{s \leq t}
\end{aligned}$$

□

Proof (of Lemma 31). For any $K \in \mathsf{KA}^*$ (which includes relational algebras) and interpretation $I : \mathsf{RExp}_\Sigma \to K$, Lemma 7 gives us

$$K, I \models E \to s \leq \tau \iff (\forall \sigma \in R(s)) \ \ K, I \models E \to \sigma \leq \tau \ .$$

In particular,

$$\begin{aligned}
\mathsf{KA}^* \models E \to s \leq \tau &\iff (\forall \sigma \in R(s)) \ \ \mathsf{KA}^* \models E \to \sigma \leq \tau \ , \\
\mathsf{RKA} \models E \to s \leq \tau &\iff (\forall \sigma \in R(s)) \ \ \mathsf{RKA} \models E \to \sigma \leq \tau \ .
\end{aligned}$$

The two right hand sides above are equivalent by Lemma 11. □

A.5 Bigness and Smallness of $*$

Proof (of Theorem 34). Noting that the proof of Lemma 21 works for any $*$-algebra, we assume without loss of generality that φ is coerced, and write φ as $E \to p \leq p'$.

Suppose $\mathsf{KAT}^* \models \varphi$. Let $K \in \mathsf{BIG}^*$ with interpretation I such that $K, I \models E$. By Theorem 18, K extends to a complete IST K'. Recall that every complete IST forms a KAT^* by letting $x^* = \sup_{n \in \omega} x^n$. However, to avoid confusion with the * operation of K, we let $x^\circ = \sup_{n \in \omega} x^n$, so that K' forms a KAT^* with $^\circ$ in place of *. For $x \in K$, we have $x^\circ \leq x^*$.

Define the interpretation $I' : \mathsf{RExp}_{\mathsf{P},\mathsf{B}} \to K'$ by letting

$$\begin{aligned}
I'(p) &= I(p) \quad\ (p \in \mathsf{P}) \\
I'(b) &= I(b) \quad\ (b \in \mathsf{B}) \\
I'(s^*) &= I'(s)^\circ \\
I'(s + t) &= I'(s) + I'(t) \\
I'(st) &= I'(s) I'(t) \\
I'(\overline{t}) &= \overline{I'(t)} \ .
\end{aligned}$$

In other words, I' is just like I, except that we are using $^\circ$ in place of *, so that I' is a proper KAT^* interpretation, while I might not be. For $t \in \mathsf{RExp}_{\mathsf{P},\mathsf{B}}$, we have $I'(t) \leq I(t)$, with equality if t is $*$-free, by induction on t.

For each hypothesis $s \leq t$ in E, t is $*$-free, and we have

$$I'(s) \leq I(s) \leq I(t) = I'(t) \ ,$$

so $K', I' \models E$. Hence, by assumption, $K', I' \models p \leq p'$, so $K, I \models p \leq p'$ (since $I(p) = I'(p) \leq I'(p') = I(p')$). Therefore $\mathsf{BIG}^* \models \varphi$. This gives us (i)⇒(ii).

(ii)⇒(iii)⇒(i) are immediate since $\mathsf{KAT}^* \subseteq \mathsf{KAT} \subseteq \mathsf{BIG}^*$. □

Proof (of Theorem 36). (i)⇒(ii) is similar to (i)⇒(ii) in the proof of the Theorem 34, except that we now have $x^* \leq x^\circ$; the other inequalities from the proof are subsequently reversed.

(ii)⇒(i) follows from $\mathsf{KAT}^* \subseteq \mathsf{SMALL}^*$.

$\mathsf{KAT} \models \varphi$ cannot be included as an equivalent condition, because the reduction used to show that $\mathcal{H}\mathsf{KAT}^*$ is Π^1_1-complete uses cosimple formulas [8]. So $\mathcal{H}\mathsf{KAT}^*$, even when restricted to cosimple formulas, is still Π^1_1-complete, while $\mathcal{H}\mathsf{KAT}$ is Σ^0_1-complete. □