On the Complexity of the Horn Theory of REL

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

Dexter Kozen
Department of Computer Science
Cornell University
Ithaca, New York 14853-7501, USA
kozen@cs.cornell.edu

May 8, 2003

Abstract

We show that the universal Horn theory of relational Kleene algebras is Π^1_1-complete.

1 Introduction

Kleene algebra (KA) is fundamental and ubiquitous in computer science. Since its invention by Kleene in 1956, it has arisen in various forms in program logic and semantics, relational algebra, automata theory, and the design and analysis of algorithms. Many authors have contributed to the development of Kleene algebra over the years (see [11] and references therein).

On the practical side, KA provides a natural and effective tool for equational specification and verification. It has recently been used successfully in numerous applications involving basic safety analysis, low-level program transformations, compiler optimization, and concurrency control [1, 2, 3, 10, 12].

The equational theory of KA has been well studied. The equational theory alone is *PSPACE*-complete [14], and this is as efficient as one could expect. However, in practice, one often needs to reason in the presence of assumptions of various forms. For example, a commutativity condition $pq = qp$ models the fact that the programs p and q can be executed in either order with the same result, and the condition $p = pb$, where b is a test, models the fact that the execution of the program p causes b to hold immediately afterward. Such assumptions are needed to reason about basic program transformations such as constant propagation and moving static computations out of loops. Several examples of this style of reasoning are given in [1, 10].

Thus the *universal Horn theory* is of interest. A *universal Horn formula* is an implication $E \to s = t$, 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, although we are usually only interested in specific substitution instances. In typical applications, the set E postulates basic assumptions about the interaction of atomic programs and tests such as $pq = qp$ or $p = pb$, and the conclusion $s = t$ represents the equivalence of the optimized and unoptimized program. The *universal Horn theory* of a class of structures $\mathcal{C}$ is the set of universal Horn formulas valid under all interpretations over structures in $\mathcal{C}$. The equational theory is the restricted case in which E is empty. The universal Horn theory of Kleene algebras is a natural consideration, since the axiomatization of KA is itself of this form.

Whereas the equational theories of various natural subclasses of Kleene algebras coincide, their Horn theories do not. For example, consider the *-continuous algebras (KA^*). A Kleene algebra is **-continuous* if it satisfies the infinitary condition

$$pq^*r \quad = \quad \sup_{n \geq 0} pq^n r,$$

where the supremum is with respect to the natural order in the Kleene algebra. Not all Kleene algebras are *-continuous, but all known naturally occurring ones are. Although *-continuity often provides a convenient shortcut in equational proofs, there are no more equations provable with it than without it; that is, the equational theories of KA and KA^* coincide [9]. However, it was shown in [11] that the universal Horn theory of KA^* is Π^1_1-complete, whereas that of KA is recursively enumerable, since it has a finitary complete first-order axiomatization. Thus the universal Horn theories of KA and KA^* diverge. Despite this fact, there is no known *natural* example of a universal Horn formula that is valid over *-continuous interpretations but not valid in general.

One important class of *-continuous interpretations is the family of *relational models*. In these models, elements are binary relations on a set X and the KA operators have standard binary relation-theoretic interpretations: the operator $\cdot$ is interpreted as relational composition $\circ$, $+$ as union $\cup$, 0 and 1 as the empty relation $\varnothing$ and the identity relation $\{(u, u) \mid u \in X\}$ on X, respectively, and * as reflexive transitive closure. The class of all relational Kleene algebras is denoted REL. This class is important because it is the preferred class of interpretations for applications in program semantics and verification.

Again, the equational theory of REL coincides with that of KA and KA^* [13], but the Horn theories diverge. Every relational model is *-continuous, so the inclusion holds in one direction; however, $p \leq 1 \to p^2 = p$ is an example of a relationally valid formula that does not hold in all *-continuous algebras. In particular, it does not hold in the *-continuous min,+ (tropical) algebra used in shortest path algorithms.

Given the importance of relational models in program semantics and verification, it is interesting to characterize their universal Horn theory. Two interesting questions are:

(i) What is the complexity of deciding whether a given universal Horn formula in the language of Kleene algebra is valid over all relational interpretations?

(ii) Is it possible to characterize this theory axiomatically?

Our main result is a solution to problem (i). We show that the universal Horn theory of REL is Π^1_1-complete. Although related to the results and constructions of [11], neither the upper nor the lower bound follows from results of [11]; both require new constructions.

2 Kleene Algebra

Kleene algebra (KA) is the algebra of regular expressions [4, 8]. The axiomatization used here is from [9]. A *Kleene algebra* is an algebraic structure $(K, +, \cdot, ^*, 0, 1)$ that is an idempotent semiring under $+, \cdot, 0, 1$ such that p^*q is the $\leq$-least solution to $q + px \leq x$ and qp^* is the $\leq$-least solution to $q + xp \leq x$. Here $\leq$ refers to the natural partial order on K: $p \leq q \overset{\text{def}}{\iff} p + q = q$. This is a universal Horn axiomatization. A Kleene algebra is **-continuous* if it satisfies the stronger infinitary property $pq^*r = \sup_n pq^n r$. The family of *-continuous Kleene algebras is denoted KA^*. It is a proper subclass of the Kleene algebras, but all naturally occurring Kleene algebras are *-continuous.

The axioms for * say essentially that * behaves like the Kleene asterate operator of formal language theory or the reflexive transitive closure operator of relational algebra.

Kleene algebra is a versatile system with many useful interpretations. Standard models include the family of regular sets over a finite alphabet; the family of binary relations on a set; and the family of $n \times n$ matrices over another Kleene algebra. Other more unusual interpretations include the min,+ algebra, also known as the *tropical semiring*, used in shortest path algorithms, and models consisting of convex polyhedra used in computational geometry.

If P is a set of atomic program symbols, a *regular expression* over P is just a term over the signature $+, \cdot, ^*, 0, 1$ of KA with atomic symbols in P. The set of all regular expressions over P is denoted RExp_P. Given an interpretation $I : \mathsf{RExp}_\mathsf{P} \to K$ over a Kleene algebra K and a quantifier-free equational Horn formula φ, we write $K, I \vDash \varphi$ if φ is true under the interpretation I under the usual semantics of first-order logic. We write $\mathsf{KA} \vDash \varphi$ and say that φ is *valid* if it is true under all interpretations. We write $\mathsf{KA}^* \vDash \varphi$ if φ is true under all interpretations over *-continuous algebras. We write $\mathsf{REL} \vDash \varphi$ and say that φ is *relationally valid* if it is true under all relational interpretations.

Let Reg_P denote the Kleene algebra of regular sets of strings over the alphabet P. The standard interpretation $R : \mathsf{RExp}_\mathsf{P} \to \mathsf{Reg}_\mathsf{P}$ mapping p to $\{p\}$, $p \in \mathsf{P}$, is universal for the equational theory of Kleene algebra; that is, $\mathsf{KA} \vDash s = t$ iff $\mathsf{Reg}_\mathsf{P}, R \vDash s = t$. Thus Reg_P is the free Kleene algebra on generators P [9]. This equational theory also coincides with the equational theories of KA^* and REL [13]. Thus if φ is a valid equation, we can write $\vDash \varphi$, omitting the KA, KA^*, or REL before the symbol $\vDash$.

3 Main Results

In this section we prove the main result of this paper: deciding the relational validity of Horn formulas of Kleene algebra is Π^1_1-complete (Corollary 3.10). The lower bound depends partially on encoding Turing machine computations as monoid equations. This part of this construction is more or less standard (see [5]) and similar to [11], but the actual reduction to the Horn theory of REL is new.

3.1 Restricted Turing Machines

Without loss of generality, we consider only total deterministic Turing machines M that conform to the following restrictions.

- M has input alphabet $\{a\}$ and finite tape alphabet Γ containing a and two special blank symbols $\rhd$ and $\lhd$ distinct from each other and from a. The alphabet Γ may contain other symbols as well.

- M has a finite set of states Q disjoint from Γ containing at least a start state s, an accept state t, and a reject state r, all distinct. There are no transitions into the start state s and no transitions out of t or r. Thus, once M enters a halt state, it cannot proceed.

- Transitions of M are of the form $((p, b), (q, c, d))$, where $p, q \in Q$, $b, c \in \Gamma$, and $d \in \{\text{left,right}\}$, indicating that when M is in state p scanning symbol b, it writes c on the current tape cell, moves its tape head one cell in direction d, and enters state q. For every (p, b) with $p \notin \{t, r\}$, there is exactly one (q, c, d) such that $((p, b), (q, c, d))$ is a transition of M.

- M has a single two-way-infinite read-write tape padded on the left by infinitely many blanks $\rhd$ and on the right by infinitely many blanks $\lhd$. M never writes $\rhd$ to the right of a nonblank symbol or $\lhd$ and never writes $\lhd$ to the left of a nonblank symbol or $\rhd$. Thus the tape always contains a unique finite contiguous string (possibly null) of nonblank symbols surrounded by infinitely many blank symbols $\rhd$ on the left and $\lhd$ on the right.

- If M either reads or writes $\rhd$, it must move right, and if it either reads or writes $\lhd$, it must move left. Thus M never moves more than one cell away from the nonblank portion of the tape.

- Inputs to M are pairs $(m,n) \in \omega^2$, represented as a pair of strings a^m, a^n. On input (m,n), M starts in state s with a^{m+n} written on its tape and its head scanning the $m+1^{\text{st}}$ symbol of a^{m+n}, or the $\lhd$ immediately following a^m if $n = 0$. If M accepts (m,n), then it does so by entering state t with a^n written on its tape and its head scanning the first $\lhd$ following the a^n. If M rejects, it erases its tape and enters state r with its head scanning the first $\lhd$.

Let $\Delta \stackrel{\text{def}}{=} \Gamma \cup Q$. A *configuration* is a string in Δ^* of the form $\rhd xqy\lhd$ or $q\rhd y\lhd$, where $x,y \in (\Gamma - \{\rhd,\lhd\})^*$ and $q \in Q$. Configurations describe instantaneous global descriptions of M in the course of some computation. In the configuration $\rhd xqy\lhd$, the current state is q, the tape currently contains the nonblank string xy surrounded by infinitely many blanks $\rhd$ on the left and $\lhd$ on the right, and M is scanning the first symbol of y. If y is null, then M is scanning the first $\lhd$ to the right of x. In the configuration $q\rhd y\lhd$, the current state is q, and M is scanning the blank symbol $\rhd$ immediately to the left of y. The start configuration of M on input (m,n) is $\rhd a^m s a^n\lhd$. If M accepts (m,n), the accept configuration is $\rhd a^n t\lhd$, and if M rejects, the reject configuration is $\rhd r\lhd$.

Let

$$\begin{aligned}
\mathsf{Config} &\stackrel{\text{def}}{=} \{\text{configurations of } M\},\\
\mathsf{Sub} &\stackrel{\text{def}}{=} \{\text{substrings of configurations of } M\},\\
\mathsf{Pre} &\stackrel{\text{def}}{=} \{\text{prefixes of configurations of } M\} \subseteq \mathsf{Sub}.
\end{aligned}$$

Note that both Sub and Pre are closed under the prefix relation.

3.2 A Rewrite Relation

Now we define a rewrite relation $\xrightarrow[M]{}$ and an associated set of equations E_M that describe the operation of M. The rewrite relation $\xrightarrow[M]{}$ consists of the following rules:

(i) for each transition $((p,\rhd),(q,c,\text{right}))$, $c \neq \rhd$, the rule $p\rhd \xrightarrow[M]{} \rhd cq$;

(ii) for each transition $((p,b),(q,\rhd,\text{right}))$, $b \neq \rhd$, the rule $\rhd pb \xrightarrow[M]{} \rhd q$;

(iii) for any other transition of the form $((p,b),(q,c,\text{right}))$ not covered by (i) or (ii), the rule $pb \xrightarrow[M]{} cq$;

(iv) for each transition $((p,\lhd),(q,c,\text{left}))$, $c \neq \lhd$, and each $e \in \Gamma - \{\lhd\}$, the rule $ep\lhd \xrightarrow[M]{} qec\lhd$;

(v) for each transition $((p,b),(q,\lhd,\text{left}))$, $b \neq \lhd$, and each $e \in \Gamma - \{\lhd\}$, the rule $epb\lhd \xrightarrow[M]{} qe\lhd$;

(vi) for any other transition of the form $((p,b),(q,c,\text{left}))$ not covered by (iv) or (v), and each $e \in \Gamma - \{\lhd\}$, the rule $epb \xrightarrow[M]{} qec$.

By the restrictions above, these cases are exhaustive. Let E_M be the set of equations

$$E_M \stackrel{\text{def}}{=} \{x = y \mid x \xrightarrow[M]{} y \text{ according to (i)–(vi) above}\}. \tag{1}$$

The relation $\xrightarrow[M]{}$ can be used to rewrite configurations in a way that mimics the computation of M. Thus we write $uxv \xrightarrow[M]{} uyv$ whenever $x \xrightarrow[M]{} y$ according to (i)–(vi) above. Note that every element of Config, Pre, or Sub has at most one redex, and rewriting by $\xrightarrow[M]{}$ preserves membership/nonmembership in Config, Pre, and Sub. Moreover, every element of Config except those containing t or r has exactly one redex.

Let $\xrightarrow[M]{*}$ denote the reflexive transitive closure of $\xrightarrow[M]{}$. Since M is assumed to be total, either M accepts (m,n), in which case $\rhd a^m s a^n \lhd \xrightarrow[M]{*} \rhd a^n t \lhd$, or M rejects (m,n), in which case $\rhd a^m s a^n \lhd \xrightarrow[M]{*} \rhd r \lhd$.

3.3 A Lower Bound

Consider a recursive relation $R \subseteq \omega^2$. One can think of R as the set of edges of a directed graph on vertices ω. The relation R is said to be *well-founded* from vertex n if all R-paths starting from n are finite. Given such an R, say by a total Turing machine M of the form described in Section 3.1, the question of whether R is well-founded from any given vertex is a well known Π^1_1-complete problem (see [6]). We will reduce this problem to the universal Horn theory of REL, thereby showing that the theory is Π^1_1-hard. We will give a separate argument in Section 3.4 to show that the theory is in Π^1_1.

Define

$$\mathsf{WF} \stackrel{\text{def}}{=} \{n \in \omega \mid R \text{ is well-founded from } n\}.$$

If we denote by $R(m)$ the set of R-successors of m,

$$R(m) \stackrel{\text{def}}{=} \{n \mid (m,n) \in R\},$$

then WF is the $\subseteq$-least solution of the recursive set equation

$$\mathsf{WF} = \{m \mid R(m) \subseteq \mathsf{WF}\}.$$

Let f be a choice function that for any $m \notin \mathsf{WF}$ gives $f(m) \in R(m) - \mathsf{WF}$. Such an $f(m)$ must exist if $m \notin \mathsf{WF}$. Define $f(m) = m$ for $m \in \mathsf{WF}$. Thus if $m \in \mathsf{WF}$, then

$(m, f(m)) \notin R$, and if $m \notin \mathsf{WF}$, then $(f^i(m), f^{i+1}(m)) \in R$ for all $i \geq 0$, therefore $m, f(m), f^2(m), \ldots$ is an infinite R-path through the graph.

Let M be a total Turing machine of the form described in Section 3.1 accepting R. Let E_M be the finite set of equations (1), and let

$$E \quad \stackrel{\text{def}}{=} \quad E_M \cup \{t \leq sa^*\}. \tag{2}$$

We define the relation $\xrightarrow[t]{}$ on Sub by

(i) $\rhd a^n ty \xrightarrow[t]{} \rhd a^n sa^{f(n)}y$ for $n \in \omega$ and any y, and

(ii) $xty \xrightarrow[t]{} xsy$ for $x \in \mathsf{Pre}$ not of the form $\rhd a^n$ and any y.

Let

$$\xrightarrow[M,t]{} \quad \stackrel{\text{def}}{=} \quad \xrightarrow[M]{} \cup \xrightarrow[t]{},$$

and let $\xrightarrow[M,t]{*}$ denote the reflexive transitive closure of $\xrightarrow[M,t]{}$.

Lemma 3.1 *For any $x \in \mathsf{Sub}$, there is at most one y such that $x \xrightarrow[M,t]{} y$.*

Proof. It suffices to show this for $x \in \mathsf{Config}$, since substrings of x can contain no more redexes than x. It is true for the relation $\xrightarrow[M]{}$, since M is deterministic, and true for the relation $\xrightarrow[t]{}$ by construction. For the union $\xrightarrow[M,t]{}$, if t occurs in x, then x contains no $\xrightarrow[M]{}$-redex, since M has no transitions out of state t. If t does not occur in x, then x contains no $\xrightarrow[t]{}$-redex. $\square$

Let $\equiv$ be the string congruence on Sub generated by $\xrightarrow[M,t]{}$; that is, the smallest reflexive, symmetric, and transitive relation respecting concatenation and containing $\xrightarrow[M,t]{}$.

Lemma 3.2 *The following are equivalent:*

(i) $x \equiv y$*;*

(ii) *there exists z such that $x \xrightarrow[M,t]{*} z$ and $y \xrightarrow[M,t]{*} z$.*

Proof. Certainly if (ii) holds, then $x \equiv y$. For the other direction, we observe that the relation on x, y defined by (ii) is a congruence containing $\xrightarrow[M,t]{}$ (transitivity following from Lemma 3.1), therefore contains the least such congruence $\equiv$. $\square$

The rewrite relations $\xrightarrow[M]{}$ and $\xrightarrow[t]{}$ clearly preserve membership in Config, Sub, and Pre. It is easily argued that they preserve nonmembership in Config, Sub, and Pre as well. It follows that $\xrightarrow[M,t]{}$ and $\equiv$ also preserve membership/nonmembership in Config, Sub, and Pre.

Define

$$\begin{array}{rcl} [x] & \stackrel{\text{def}}{=} & \{y \mid x \equiv y\} \\ \mathsf{Pre}/\!\equiv & \stackrel{\text{def}}{=} & \{[x] \mid x \in \mathsf{Pre}\}. \end{array}$$

Let K be the Kleene algebra of all binary relations on $\mathsf{Pre}/\!\equiv$.

For each $p \in \Delta$, define

$$I(p) \quad \stackrel{\text{def}}{=} \quad \{([x],[xp]) \mid xp \in \mathsf{Pre}\},$$

and extend I homomorphically to an interpretation $I : \mathsf{RExp}_\Delta \to K$. We will show below that if

$$K, I \quad \models \quad E \to \ \triangleright a^m t \triangleleft \ \leq \ \triangleright r \triangleleft,$$

then $m \in \mathsf{WF}$.

Lemma 3.3 *For any $y \in \Delta^*$,*

$$I(y) \quad = \quad \{([x],[xy]) \mid xy \in \mathsf{Pre}\}.$$

Proof. The proof is by induction on the length of y. For the empty string ε, we have

$$I(\varepsilon) \quad = \quad I(1) \quad = \quad \{([x],[x]) \mid x \in \mathsf{Pre}\} \quad = \quad \{([x],[x\varepsilon]) \mid x\varepsilon \in \mathsf{Pre}\}.$$

Now assume the lemma holds for y. For yp, where $p \in \Delta$,

$$\begin{array}{rcll} I(yp) & = & I(y) \circ I(p) & \\ & = & \{([x],[xy]) \mid xy \in \mathsf{Pre}\} \circ \{([w],[wp]) \mid wp \in \mathsf{Pre}\} & (3) \\ & = & \{([x],[wp]) \mid xy \equiv w \text{ and } wp \in \mathsf{Pre}\} & (4) \\ & \subseteq & \{([x],[wp]) \mid xyp \equiv wp \text{ and } wp \in \mathsf{Pre}\} & (5) \\ & = & \{([x],[xyp]) \mid xyp \in \mathsf{Pre}\}. & \end{array}$$

Step (3) follows from the induction hypothesis and the definition of $I(p)$. In step (4), requiring $xy \in \mathsf{Pre}$ is redundant, since Pre is closed under prefix and $\equiv$ preserves membership in Pre. Step (5) follows from the fact that $\equiv$ is a congruence.

Conversely, since Pre is closed under prefix,

$$\begin{aligned}
&\{(\texttt{[}x\texttt{]},\texttt{[}xyp\texttt{]}) \mid xyp \in \mathsf{Pre}\} \\
&\quad\subseteq\quad \{(\texttt{[}x\texttt{]},\texttt{[}xy\texttt{]}) \mid xy \in \mathsf{Pre}\} \circ \{(\texttt{[}xy\texttt{]},\texttt{[}xyp\texttt{]}) \mid xyp \in \mathsf{Pre}\} \\
&\quad=\quad I(y) \circ I(p).
\end{aligned}$$

□

Lemma 3.4 $K, I \vDash E$.

Proof. If $y \underset{M}{\longrightarrow} z$, we have $y \equiv z$, therefore $\texttt{[}xy\texttt{]} = \texttt{[}xz\texttt{]}$ for all x, since $\equiv$ is a congruence. Moreover, $xy \in \mathsf{Pre}$ iff $xz \in \mathsf{Pre}$, since $\equiv$ preserves membership in Pre. By Lemma 3.3,

$$I(y) \;=\; \{(\texttt{[}x\texttt{]},\texttt{[}xy\texttt{]}) \mid xy \in \mathsf{Pre}\} \;=\; \{(\texttt{[}x\texttt{]},\texttt{[}xz\texttt{]}) \mid xz \in \mathsf{Pre}\} \;=\; I(z).$$

Thus $K, I \vDash y = z$ for any equation $y = z$ in E_M.

Now consider $t \le sa^*$. For $(\texttt{[}z\texttt{]},\texttt{[}zt\texttt{]}) \in I(t)$, $zt \in \mathsf{Pre}$, let

$$\ell \;=\; \begin{cases} f(k), & \text{if } z = \triangleright a^k, \\ 0, & \text{otherwise.} \end{cases}$$

By definition of $\underset{t}{\longrightarrow}$, $zt \equiv zsa^\ell$, thus

$$(\texttt{[}z\texttt{]},\texttt{[}zt\texttt{]}) \;=\; (\texttt{[}z\texttt{]},\texttt{[}zsa^\ell\texttt{]}) \;\in\; I(sa^\ell) \;\subseteq\; I(sa^*),$$

therefore $I(t) \subseteq I(sa^*)$ and $K, I \vDash t \le sa^*$. □

Lemma 3.5 *If* $K, I \vDash \triangleright a^m t \triangleleft \le \triangleright r \triangleleft$*, then* $m \in \mathsf{WF}$.

Proof. Suppose $K, I \vDash \triangleright a^m t \triangleleft \le \triangleright r \triangleleft$. By Lemma 3.3,

$$I(\triangleright r \triangleleft) \;=\; \{(\texttt{[}z\texttt{]},\texttt{[}z\triangleright r \triangleleft\texttt{]}) \mid z\triangleright r\triangleleft \in \mathsf{Pre}\} \;=\; \{(\texttt{[}\varepsilon\texttt{]},\texttt{[}\triangleright r\triangleleft\texttt{]})\},$$

since $z\triangleright r\triangleleft \in \mathsf{Pre}$ only if $z = \varepsilon$. Then

$$\begin{aligned}
(\texttt{[}\varepsilon\texttt{]},\texttt{[}\triangleright a^m t\triangleleft\texttt{]}) \quad&\in\quad I(\triangleright a^m t \triangleleft) \qquad \text{by Lemma 3.3} \\
&\subseteq\quad I(\triangleright r \triangleleft) \\
&=\quad \{(\texttt{[}\varepsilon\texttt{]},\texttt{[}\triangleright r\triangleleft\texttt{]})\},
\end{aligned}$$

so $\triangleright a^m t \triangleleft \equiv \triangleright r \triangleleft$. By Lemma 3.2, there is a z such that $\triangleright a^m t \triangleleft \overset{*}{\underset{M,t}{\longrightarrow}} z$ and $\triangleright r \triangleleft \overset{*}{\underset{M,t}{\longrightarrow}} z$. But since $\triangleright r \triangleleft$ contains no $\underset{M,t}{\longrightarrow}$-redexes, we must have $\triangleright a^m t \triangleleft \overset{*}{\underset{M,t}{\longrightarrow}} \triangleright r \triangleleft$.

Now if $m \notin \mathsf{WF}$, then

$$\triangleright a^m t \triangleleft \xrightarrow[t]{} \triangleright a^m s a^{f(m)} \triangleleft \xrightarrow[M]{*} \triangleright a^{f(m)} t \triangleleft \xrightarrow[t]{} \triangleright a^{f(m)} s a^{f^2(m)} \triangleleft \xrightarrow[M]{*} \cdots ,$$

contradicting Lemma 3.1 and the fact that $\triangleright a^m t \triangleleft \xrightarrow[M,t]{*} \triangleright r \triangleleft$. Thus $m \in \mathsf{WF}$. □

Theorem 3.6 *The following are equivalent:*

(i) $m \in \mathsf{WF}$*;*

(ii) $\mathsf{KA}^* \vDash E \to \triangleright a^m t \triangleleft \leq \triangleright r \triangleleft$*;*

(iii) $\mathsf{REL} \vDash E \to \triangleright a^m t \triangleleft \leq \triangleright r \triangleleft$.

Proof. The argument for (i) ⇒ (ii) is the same as in [11], *mutatis mutandis*. The implication (ii) ⇒ (iii) is a direct consequence of the inclusion $\mathsf{REL} \subseteq \mathsf{KA}^*$. Finally, (iii) ⇒ (i) is immediate from Lemmas 3.4 and 3.5 and the fact that $K \in \mathsf{REL}$. □

Corollary 3.7 *The universal Horn theory of* REL *is* Π^1_1*-hard.*

Proof. Our construction of E from M is effective, therefore constitutes a reduction from the well-foundedness problem to the Horn theory of REL. □

3.4 An Upper Bound

It remains to show that the universal Horn theory of REL is Π^1_1. We first show that it suffices to restrict our attention to countable models.

Lemma 3.8 *Let φ be an arbitrary first-order sentence in the language of Kleene algebra. The following are equivalent:*

(i) *φ is valid over all relational Kleene algebras;*

(ii) *φ is valid over all countable relational Kleene algebras over countably many states.*

Proof.

The implication (i) ⇒ (ii) is immediate.

For (ii) ⇒ (i), suppose (i) fails. Then there is a relational Kleene algebra and interpretation I over that algebra satisfying $\neg\varphi$. By the downward Löwenheim-Skolem theorem, that algebra has a countable elementary substructure K containing the image of I. Then $K, I \vDash \neg\varphi$. Let S be the set of states of K. Although K is countable, S need not be. However, we can pare S down to a countable set of states S' while maintaining the algebraic structure of K. Specifically, the map $x \mapsto x \restriction S'$ will be an injective homomorphism of K into the algebra of binary relations on S', where $x \upharpoonright S'$ denotes $x \cap (S' \times S')$.

For $x, y \in K$ such that $x \not\leq y$, let $(s_{xy}, t_{xy}) \in x - y$. For $x, y \in K$ and $(s, t) \in xy$, let $u_{xyst} \in S$ such that $(s, u_{xyst}) \in x$, $(u_{xyst}, t) \in y$. The pair (s_{xy}, t_{xy}) witnesses the fact that $x \not\leq y$, and u_{xyst} witnesses the fact that $(s, t) \in xy$. Let S' be the smallest set of states containing s_{xy} and t_{xy} for $x, y \in K$, $x \not\leq y$, and closed under the addition of u_{xyst} for $s, t \in S'$, $(s, t) \in xy$. Note that S' is countable, since it can be constructed as the union of a countable chain of countable sets.

Now let K' be the relational structure on S' consisting of elements $x \upharpoonright S'$ for $x \in K$. We claim that this structure is a relational Kleene algebra and that the map $x \mapsto x \upharpoonright S'$ is an isomorphism. The map is surjective by definition and injective since S' contains s_{xy} and t_{xy}.

To argue that relational composition works correctly, note that for $x, y \in K$, for any $(s, t) \in xy \upharpoonright S'$, we have $(s, u_{xyst}) \in x \upharpoonright S'$ and $(u_{xyst}, t) \in y \upharpoonright S'$, therefore $(s, t) \in (x \upharpoonright S') \cdot (y \upharpoonright S')$. The reverse inclusion is straightforward, therefore $xy \upharpoonright S' = (x \upharpoonright S') \cdot (y \upharpoonright S')$.

That $0 \upharpoonright S' = \varnothing$, $1 \upharpoonright S' = \{(u, u) \mid u \in S'\}$, and $(x + y) \upharpoonright S' = x \upharpoonright S' + y \upharpoonright S'$ are all straightforward. That $x^* \upharpoonright S' = (x \upharpoonright S')^*$ follows from the fact that the map $x \mapsto x \upharpoonright S'$ respects relational composition and arbitrary union.

We have constructed a countable relational model K' on countably many states satisfying $\neg\varphi$. Thus for any φ, φ is valid over all relational models iff it is valid over all countable relational models on countably many states. $\square$

Theorem 3.9 *The universal Horn theory of* REL *is in* Π^1_1.

Proof. We will express the validity of a KA sentence φ as a Π^1_1 sentence of arithmetic. This will involve an arithmetic encoding of sentences of KA. Validity is expressed using second-order universal quantification over all countable relational Kleene algebras over states ω, which is sufficient by Lemma 3.8.

Let $\langle\ ,\ \rangle : \omega^2 \to \omega$ be a standard arithmetic pairing function. To interpret a set $X \subseteq \omega$ as a countable relational Kleene algebra, we interpret X as a set of triples $\langle s, t, m\rangle \in \omega^3$, indicating that the pair (s, t) is in the m^{th} element of the algebra. We let $\overline{m} = \{(s, t) \mid \langle s, t, m\rangle \in X\}$. In formulas, we will write $(s, t) \in \overline{m}$ as shorthand for $\langle s, t, m\rangle \in X$.

The formula $\overline{n} = \overline{m}^*$ can be translated as

$$\forall s, t\ \ ((s, t) \in \overline{n}\ \leftrightarrow\ \exists\langle s_0, s_1, \ldots, s_k\rangle\ (s_0 = s \wedge s_k = t \wedge \forall i < k\ (s_i, s_{i+1}) \in \overline{m}));$$

that is, $(s, t) \in \overline{n}$ iff $(s, t) \in \overline{m}^k$ for some $k \geq 0$. The quantification over arbitrary finite sequences of natural numbers can be coded using Gödel's β function (see [7, p. 238]).

The translations of $\overline{n} = 0$, $\overline{n} = 1$, $\overline{n} = \overline{\ell} + \overline{m}$, and $\overline{n} = \overline{\ell}\,\overline{m}$ are similar.

We now define the predicate $\mathsf{IsModel}(X)$:

$$\forall \ell, m\ \exists n_0, n_1, n_2, n_3, n_4\ \ (\overline{n}_0 = 0 \wedge \overline{n}_1 = 1 \wedge \overline{n}_2 = \overline{\ell} + \overline{m} \wedge \overline{n}_3 = \overline{\ell}\,\overline{m} \wedge \overline{n}_4 = \overline{\ell}^*).$$

This says that X does in fact encode a relational model. Note that $\mathsf{IsModel}(X)$ does not require $\overline{m} \neq \overline{n}$ for $m \neq n$.

Using the coding above, for any first-order sentence φ in the language of Kleene algebra, we can effectively construct a predicate $\mathsf{Models}_\varphi(X)$ that says that X models φ. The formula $\mathsf{Models}_\varphi(X)$ uses only first order quantifiers. Instead of quantifying over interpretation functions separately, we adopt the convention that the n constant symbols appearing in φ will be interpreted as the first n elements of X, so that X is really a model paired with an interpretation. For example, if φ is $\forall x\ (x + c^* = c)$, and if we wish to interpret the constant c as $\overline{0}$, then $\mathsf{Models}_\varphi(X)$ would be

$$\forall x\ \exists y_1, y_2\ \ \overline{y}_1 = \overline{0}^*\ \wedge\ \overline{y}_2 = \overline{x} + \overline{y}_1\ \wedge\ \overline{y}_2 = \overline{0}.$$

The validity of φ over relational models can then be expressed

$$\forall X\ \ \mathsf{IsModel}(X)\ \rightarrow\ \mathsf{Models}_\varphi(X),$$

which is Π^1_1. □

Corollary 3.10 *The Horn theory of* REL *is* Π^1_1*-complete.*

Acknowledgments

This work was supported in part by NSF grant CCR-0105586 and ONR Grant N00014-01-1-0968. The views and conclusions contained herein are those of the authors 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] Adam Barth and Dexter 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] Ernie Cohen. Lazy caching. Available as ftp://ftp.telcordia.com/pub/ernie/research/homepage.html, 1994.

[3] Ernie Cohen. Using Kleene algebra to reason about concurrency control. Available as ftp://ftp.telcordia.com/pub/ernie/research/homepage.html, 1994.

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

[5] Martin Davis. *Computability and Unsolvability*. McGraw-Hill, New York, 1958.

[6] David Harel, Dexter Kozen, and Jerzy Tiuryn. *Dynamic Logic*. MIT Press, Cambridge, MA, 2000.

[7] S. C. Kleene. *Introduction to Metamathematics*. D. van Nostrand, 1952.

[8] Stephen C. Kleene. Representation of events in nerve nets and finite automata. In C. E. Shannon and J. McCarthy, editors, *Automata Studies*, pages 3–41. Princeton University Press, Princeton, N.J., 1956.

[9] Dexter Kozen. A completeness theorem for Kleene algebras and the algebra of regular events. *Infor. and Comput.*, 110(2):366–390, May 1994.

[10] Dexter Kozen. Kleene algebra with tests. *Transactions on Programming Languages and Systems*, 19(3):427–443, May 1997.

[11] Dexter Kozen. On the complexity of reasoning in Kleene algebra. *Information and Computation*, 179:152–162, 2002.

[12] Dexter Kozen and Maria-Cristina Patron. Certification of compiler optimizations using Kleene algebra with tests. In John Lloyd, Veronica Dahl, Ulrich Furbach, Manfred Kerber, Kung-Kiu Lau, Catuscia Palamidessi, Luis Moniz Pereira, Yehoshua Sagiv, and Peter 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] Dexter Kozen and Frederick Smith. Kleene algebra with tests: Completeness and decidability. In D. van Dalen and M. Bezem, editors, *Proc. 10th Int. Workshop Computer Science Logic (CSL'96)*, volume 1258 of *Lecture Notes in Computer Science*, pages 244–259, Utrecht, The Netherlands, September 1996. Springer-Verlag.

[14] L. J. Stockmeyer and A. R. Meyer. Word problems requiring exponential time. In *Proc. 5th Symp. Theory of Computing*, pages 1–9, New York, 1973. ACM.