

Partial Automata and Finitely Generated Congruences: An Extension of Nerode's Theorem*

Dexter Kozen
University of Aarhus and Cornell University[†]

For Anil Nerode, on the occasion of his 60th birthday

June 4, 1992

Abstract

Let T_Σ be the set of ground terms over a finite ranked alphabet Σ . We define *partial automata* on T_Σ and prove that the finitely generated congruences on T_Σ are in one-to-one correspondence (up to isomorphism) with the finite partial automata on Σ with no inaccessible and no inessential states. We give an application in term rewriting: every ground term rewrite system has a canonical equivalent system that can be constructed in polynomial time.

*Research supported in part by the Danish Research Academy, the National Science Foundation, the John Simon Guggenheim Foundation, and the U.S. Army Research Office through the ACSyAM branch of the Mathematical Sciences Institute of Cornell University, contract DAAL03-91-C-0027.

[†]Address until July 31, 1992: Computer Science Department, University of Aarhus, Ny Munkegade, DK-8000 Aarhus C, Denmark. Address from August 1, 1992: Computer Science Department, Cornell University, Ithaca, New York 14853, USA. Email: kozen@cs.cornell.edu

1 Introduction

The *Myhill-Nerode Theorem* is a classic result in the theory of finite automata. It dates to work of Myhill [13] and Nerode [14] in the late 1950s, but is still today considered one of the most important results in the subject. It has numerous applications, especially in showing that certain sets are regular or certain apparently stronger types of automata are really no more powerful than finite automata. Nevertheless, its statement and proof are elementary enough that it can be taught in introductory courses.

The Myhill-Nerode Theorem exploits a fundamental connection between combinatorics and algebra to give a particularly satisfying characterization of the regular sets over a finite alphabet. As presented in a standard undergraduate text [8], it states:

Myhill-Nerode Theorem [13, 14] *Let R be a set of strings over a finite alphabet Σ . The following three propositions are equivalent:*

- (i) *R is accepted by a finite automaton*
- (ii) *R is a union of classes of a right-invariant equivalence relation of finite index*
- (iii) *the relation \equiv_R is of finite index, where $x \equiv_R y$ iff*

$$\forall z \in \Sigma^* \quad xz \in R \leftrightarrow yz \in R .$$

The equivalence of (i) and (ii) is generally established using the following lemma:

Correspondence Lemma *Up to isomorphism, there is a one-to-one correspondence between the right-invariant equivalence relations of finite index on Σ and deterministic finite automata over Σ with no inaccessible states.*

Essentially, the states correspond to the equivalence classes, and the property of right invariance allows the deterministic transition function to be defined unambiguously on equivalence classes.

The Myhill-Nerode Theorem generalizes in a straightforward way to automata on finite trees. This generalization first came to light in the late

1960s, ten years after Myhill and Nerode’s work, and can be attributed to a combination of results of Brainerd [2, 3], Eilenberg and Wright [5], and Arbib and Give’on [1], although one must also credit Thatcher and Wright [15] in this context with the development of the algebraic approach to automata on finite trees, which allows “conventional finite automata theory [to go] through for the generalization—and. . . quite neatly” [15]. A particularly easy proof of this generalization in the style of [8] can be found in [11].

In the Thatcher-Wright approach to automata on finite trees, the elements of Σ are assigned finite *arities*, and instead of strings one works with the ground terms T_Σ over Σ . A *deterministic finite tree automaton* over Σ is just a finite Σ -algebra \mathcal{A} , consisting of a finite carrier $|\mathcal{A}|$ and a distinguished n -ary function $f^{\mathcal{A}} : |\mathcal{A}|^n \rightarrow |\mathcal{A}|$ for each n -ary symbol $f \in \Sigma$. This definition includes the nullary case ($n = 0$), in which the function symbol is called a *constant* and interpreted as an element of $|\mathcal{A}|$. By analogy with the combinatorial treatment of [8], we call elements of $|\mathcal{A}|$ *states*.

Since T_Σ is the free Σ -algebra on the empty set of generators, there exists a unique Σ -algebra homomorphism

$$\delta : T_\Sigma \rightarrow \mathcal{A} .$$

This map assigns a unique state $\delta(t)$ to each term t in an inductive fashion, and is analogous to “running” the automaton on input t . A state is said to be *accessible* if it is $\delta(t)$ for some term t .

An equivalence relation R on T_Σ is said to be *recognized* by the automaton \mathcal{A} if the kernel of δ (*i.e.*, the relation $\{s \equiv t \mid \delta(s) = \delta(t)\}$) refines R . In other words, R is recognized by \mathcal{A} if for any terms $s, t \in T_\Sigma$, if $\delta(s) = \delta(t)$, then sRt . The special case of regular sets discussed above corresponds to an R with two equivalence classes. If R is recognized by \mathcal{A} , it is possible to partition the states of \mathcal{A} such that the inverse image of the partition under δ coincides with R ; this partition of the states corresponds to the specification of a set of final or accept states in the special case of regular sets.

For a given equivalence relation $R \subseteq T_\Sigma$ (recognizable or not), define $s \equiv_R t$ if for all terms u with exactly one occurrence of a variable x and no other variables,

$$u[x/s] R u[x/t],$$

where $u[x/s]$ denotes the term obtained by substituting s for x in u . The relation \equiv_R generalizes the relation on strings of the same name mentioned

above.

Myhill-Nerode Theorem for trees [3, 5, 1] *Let R be an equivalence negation on T_Σ . The following three propositions are equivalent:*

- (i) *R is recognizable*
- (ii) *there exists a congruence on T_Σ of finite index refining R*
- (iii) *the relation \equiv_R is of finite index.*

The Myhill-Nerode theorem for strings corresponds to the special case of a single nullary operator and several unary operators.

In the algebraic approach, the tree version of the Correspondence Lemma reduces to an elementary fact of universal algebra: up to isomorphism, the homomorphic images of T_Σ and the congruences on T_Σ are in one-to-one correspondence. The correspondence is given by the quotient construction

$$\equiv \mapsto T_\Sigma / \equiv ,$$

in which it is readily observed that the quotient is finite iff the corresponding congruence is of finite index.

In [9, 10], we investigated the complexity of various decision problems in Σ -algebras presented by finite sets of ground equations over T_Σ ; that is, quotients of T_Σ modulo finitely generated congruences on T_Σ . We showed, among other results, that every such algebra has a minimal canonical presentation that is unique up to isomorphism.

This result has an interesting interpretation in terms of the Myhill-Nerode Theorem. First, we note that every congruence \equiv on T_Σ of finite index is finitely generated. To see this, let $U \subseteq T_\Sigma$ be a complete set of representatives for the \equiv -classes, and consider the finite subrelation consisting of all pairs in \equiv of the form

$$f u_1 \dots u_n \equiv u \tag{1}$$

for $u_1, \dots, u_n, u \in U$ and $f \in \Sigma_n$. The relation generated by the equations (1) is surely contained in \equiv ; conversely, an easy inductive argument shows that every term is equivalent to the $u \in U$ in its \equiv -class under the congruence generated by the equations (1).

However, not every finitely generated congruence is of finite index: for example, the identity relation on T_Σ is of infinite index (assuming Σ has at least one constant and at least one symbol of higher arity), but is generated by the empty relation.

The question thus arises as to whether there is a more general version of the Myhill-Nerode theorem with “finitely generated” in place of “finite index”.

The answer to this question is mixed. On the positive side, we formulate and prove a version of the Correspondence Lemma in this more general setting. On the other hand, we construct an equivalence relation R that has no minimal refining finitely generated congruence.

In order to formulate the first result, we need a combinatorial structure that is to finitely generated congruences as finite tree automata are to congruences of finite index. The appropriate notion is a *finite partial automaton* on T_Σ . Simply stated, a finite partial automaton is just a finite partial Σ -algebra, where a partial Σ -algebra is like a Σ -algebra except the distinguished operations need not be everywhere defined. We will show how a finite partial automaton \mathcal{A} uniquely determines a possibly infinite set of “states”. This is done formally by a universal algebraic construction giving a certain total extension $\hat{\mathcal{A}}$ of \mathcal{A} called its *free total extension*.

Finally, we give an application to term rewriting. We show that every ground term rewrite system has a canonical equivalent system which is unambiguous and in which all rules are of the form $f q_1 \dots q_n \rightarrow q$, where q_1, \dots, q_n, q are auxiliary constants. By *canonical* we mean that the system is minimal and unique up to isomorphism. The canonical system can be obtained effectively from the original system in polynomial time. This allows us to test the equivalence of ground term rewrite systems over a signature of bounded arity in polynomial time. When the arity is unbounded, the equivalence problem for ground term rewrite systems is equivalent to graph isomorphism.

Although the notions of partial automaton and free total extension and the formulation of this result in automata-theoretic terms are apparently new, much of the essential content is more or less implicit in [9, 10]

2 Partial Algebras and Partial Automata

Let Σ be an arbitrary but fixed finite ranked alphabet. The rank of $f \in \Sigma$ is called its *arity*. The set of n -ary elements of Σ is denoted Σ_n . The set of ground terms over Σ is denoted T_Σ .

A *congruence* on T_Σ is an equivalence relation \equiv such that $fs_1 \dots s_n \equiv ft_1 \dots t_n$ whenever $f \in \Sigma_n$ and $s_i \equiv t_i, 1 \leq i \leq n$. If Γ is a binary relation on T_Σ , the congruence *generated by* Γ is the smallest congruence on T_Σ containing Γ . For $s, t \in T_\Sigma$, we write $s \equiv t (\Gamma)$ and say s and t are *congruent modulo* Γ if s and t are equivalent modulo the congruence generated by Γ . A congruence \equiv is *finitely generated* if it is generated by a finite subrelation.

An equivalence relation \equiv is of *finite index* if there are only finitely many \equiv -classes. An equivalence relation R *refines* another equivalence relation S if each S -class is a union of R -classes; equivalently, if sRt implies sSt .

Definition 1 A *partial Σ -algebra* (or just *partial algebra* for short) is a structure

$$\mathcal{A} = (|\mathcal{A}|, \cdot^{\mathcal{A}})$$

where $|\mathcal{A}|$ is a set, called the *carrier* of \mathcal{A} , and $\cdot^{\mathcal{A}}$ assigns a partial n -ary function

$$f^{\mathcal{A}} : |\mathcal{A}|^n \rightarrow |\mathcal{A}|$$

to each n -ary function symbol f of Σ . By *partial* we mean that $f^{\mathcal{A}}$ need not be everywhere defined. We identify nullary functions

$$c^{\mathcal{A}} : |\mathcal{A}|^0 \rightarrow |\mathcal{A}|$$

with elements of $|\mathcal{A}|$. Nullary symbols c are often called *constants*. We usually use c, d, \dots for constants and f, g, \dots for function symbols in Σ of any arity. Like functions of higher arity, $c^{\mathcal{A}}$ may be undefined in a partial algebra \mathcal{A} .

The partial algebra \mathcal{A} is said to be *total* if all functions $f^{\mathcal{A}}$ are everywhere defined. It is said to be *finite* if $|\mathcal{A}|$ is a finite set. \square

Definition 2 Let \mathcal{A} and \mathcal{B} be two partial Σ -algebras. A (total) function

$$h : \mathcal{A} \rightarrow \mathcal{B}$$

is a *partial Σ -algebra homomorphism* (or just *partial homomorphism* for short) if, whenever $q_1, \dots, q_n \in \mathcal{A}$, $f \in \Sigma_n$, and $f^{\mathcal{A}}(q_1, \dots, q_n)$ is defined, then $f^{\mathcal{B}}(h(q_1), \dots, h(q_n))$ is defined and equal to $h(f^{\mathcal{A}}(q_1, \dots, q_n))$. We emphasize that partial homomorphisms are always total functions.

We write $\mathcal{A} \sqsubseteq \mathcal{B}$ and say that \mathcal{A} is a *partial subalgebra* of \mathcal{B} and that \mathcal{B} is an *extension* of \mathcal{A} if $|\mathcal{A}| \subseteq |\mathcal{B}|$ and the inclusion map $\mathcal{A} \rightarrow \mathcal{B}$ is a partial homomorphism.

A partial subalgebra \mathcal{A} of \mathcal{B} is said to be the *induced partial subalgebra* of \mathcal{B} on $Q \subseteq |\mathcal{B}|$ if $|\mathcal{A}| = Q$ and for all $q_1, \dots, q_n \in Q$ and $f \in \Sigma_n$,

$$f^{\mathcal{A}}(q_1, \dots, q_n) = f^{\mathcal{B}}(q_1, \dots, q_n)$$

whenever the right hand side is defined and in Q . □

Definition 3 If \mathcal{A} is a partial algebra, let $T_{\Sigma \cup |\mathcal{A}|}$ be the set of ground terms over the disjoint union $\Sigma \cup |\mathcal{A}|$, where we assign elements of $|\mathcal{A}|$ arity 0. The set of formal equations

$$\Delta_{\mathcal{A}} = \{q \equiv f q_1 \dots q_n \mid q_1, \dots, q_n, q \in |\mathcal{A}|, f \in \Sigma_n, \\ f^{\mathcal{A}}(q_1, \dots, q_n) \text{ exists and is equal to } q\}$$

is called the *diagram* of \mathcal{A} . □

The term *partial automaton* is synonymous with partial algebra. When thinking automata-theoretically, we often call elements of $|\mathcal{A}|$ *states*.

A conventional tree automaton over Σ in the sense of Thatcher and Wright is just a finite total Σ -algebra \mathcal{A} . Informally, such an automaton takes a ground term in T_{Σ} as input. It starts at the leaves and moves upward, associating a state with each subterm inductively. If the immediate subterms t_1, \dots, t_n of the term $f t_1 \dots t_n$ are labeled with states q_1, \dots, q_n respectively, then the term $f t_1 \dots t_n$ will be labeled with state $f^{\mathcal{A}}(q_1, \dots, q_n)$. Note that the basis of the induction is included here: the state labeling the term c is $c^{\mathcal{A}}$.

Formally, the labeling function is just the unique Σ -algebra homomorphism

$$\delta : T_{\Sigma} \rightarrow \mathcal{A}$$

from the free Σ -algebra T_{Σ} to \mathcal{A} . By considerations of universal algebra, this homomorphism exists and is unique. A state of \mathcal{A} is said to be *accessible* if

it is in the image of T_Σ under δ , *inaccessible* otherwise. Thus we would say that the automaton \mathcal{A} has no inaccessible states if the map δ is onto.

This definition extends the usual definition of automata on finite strings in a natural way: we can think of an automaton on strings over a finite alphabet Σ as a tree automaton over $\Sigma \cup \{\square\}$ turned on its side, where \square is a new constant and elements of Σ are assigned arity 1.

Equivalently, we can define finite tree automata as term rewrite systems. This is the approach taken for example in [7]. Given an algebra \mathcal{A} , we can consider $\Delta_{\mathcal{A}}$ as a ground term rewrite system on $T_{\Sigma \cup |\mathcal{A}|}$ in which the equations are ordered from right to left. This system is unambiguous and terminating, thus normal forms exist and are unique [4]. By elementary considerations of term rewrite theory, the terms s and t are congruent modulo $\Delta_{\mathcal{A}}$ iff they have the same normal form. For a total algebra \mathcal{A} , the $\Delta_{\mathcal{A}}$ -normal form of term t is $\delta(t) \in |\mathcal{A}|$.

3 Free Total Extensions

A partial automaton runs inductively on a ground term in the same way as a total automaton. However, the reader is probably already asking the obvious questions what happens when it reaches a situation from which it cannot continue because the appropriate $f^{\mathcal{A}}(q_1, \dots, q_n)$ is undefined? Informally, whenever it encounters such a situation, it *creates* a new state *symbolically* and moves to it. In this way a finite partial automaton \mathcal{A} gives rise to a possibly infinite set $\widehat{\mathcal{A}}$ of symbolic states that would be created in this way. The construction of $\widehat{\mathcal{A}}$ from \mathcal{A} is analogous to the construction of algebraic extensions of fields or of the rational numbers from the integers where we wish to extend the structure in the freest possible way so that certain functions are defined. We formalize this idea by the notion of *free total extension* of a partial algebra.

Formally, free total extensions are defined in terms of their most salient property, a universality property similar to that of free algebras.

Definition 4 The *free total extension* of a partial algebra \mathcal{A} is defined to be a total extension $\widehat{\mathcal{A}}$ of \mathcal{A} such that for any total algebra \mathcal{B} and partial Σ -algebra homomorphism $h : \mathcal{A} \rightarrow \mathcal{B}$, there is a unique Σ -algebra homomor-

phism $\hat{h} : \hat{\mathcal{A}} \rightarrow \mathcal{B}$ such that the diagram

$$\begin{array}{ccc} & \hat{\mathcal{A}} & \\ \sqsubseteq \uparrow & \xrightarrow{\hat{h}} & \mathcal{B} \\ \mathcal{A} & \xrightarrow{h} & \end{array}$$

(2)

commutes. □

Theorem 5 *The free total extension $\hat{\mathcal{A}}$ of a partial Σ -algebra \mathcal{A} exist and is unique up to isomorphism. Moreover, \mathcal{A} is the induced partial subalgebra of $\hat{\mathcal{A}}$ on $|\mathcal{A}|$.*

Proof. Let $\Delta_{\mathcal{A}}$ be the diagram of \mathcal{A} (Definition 3) and take $\hat{\mathcal{A}} = T_{\Sigma \cup |\mathcal{A}|} \setminus \Delta_{\mathcal{A}}$. Let $\nu(t)$ denote the $\Delta_{\mathcal{A}}$ -normal form of $t \in T_{\Sigma \cup |\mathcal{A}|}$ and let $[t]$ denote the congruence class of t modulo $\Delta_{\mathcal{A}}$. The canonical map $t \mapsto [t]$ restricted to domain $|\mathcal{A}|$ constitutes a partial homomorphism $\mathcal{A} \rightarrow \mathcal{A}$, since if $f^{\mathcal{A}}(q_1, \dots, q_n) = q$, then $q \equiv f q_1 \dots q_n \in \Delta_{\mathcal{A}}$, therefore

$$f^{\hat{\mathcal{A}}}([q_1], \dots, [q_n]) = [f q_1 \dots q_n] = [q] . \quad (3)$$

This map is also one-to-one on \mathcal{A} since distinct elements of \mathcal{A} have distinct normal forms ($\nu(q) = q$ for $q \in |\mathcal{A}|$), therefore occupy distinct $\Delta_{\mathcal{A}}$ -congruence classes. By a slight abuse, we may thus consider $\mathcal{A} \sqsubseteq \hat{\mathcal{A}}$.

The partial algebra \mathcal{A} is the induced partial subalgebra of $\hat{\mathcal{A}}$ on $|\mathcal{A}|$, since if (3) holds with $q_1, \dots, q_n, q \in \mathcal{A}$, then

$$\nu(f q_1 \dots q_n) = \nu(q) = q ,$$

thus $q \equiv f q_1 \dots q_n \in \Delta_{\mathcal{A}}$, therefore $f^{\mathcal{A}}(q_1, \dots, q_n)$ exists and is equal to q .

If $h : \mathcal{A} \rightarrow \mathcal{B}$ is a partial Σ -algebra homomorphism from \mathcal{A} to any total algebra \mathcal{B} then let h' denote the unique homomorphism $T_{\Sigma \cup |\mathcal{A}|} \rightarrow \mathcal{B}$ such that $h'(q) = h(q)$ for $q \in \mathcal{A}$. We wish to show that h' factors through $\hat{\mathcal{A}}$ giving the following commutative diagram:

$$\begin{array}{ccc}
T_{\Sigma \cup |\mathcal{A}|} & \xrightarrow{h'} & \mathcal{B} \\
\downarrow [\] & \searrow \hat{h} & \\
\hat{\mathcal{A}} & \xrightarrow{\hat{h}} & \mathcal{B} \\
\sqsubseteq \uparrow & \nearrow h & \\
\mathcal{A} & &
\end{array}$$

For this purpose it suffices to show that if $s \equiv t(\Delta_{\mathcal{A}})$ then $h'(s) = h'(t)$. For any equatian $q \equiv f q_1 \dots q_n \in \Delta_{\mathcal{A}}$, we have that $f^{\mathcal{A}}(q_1, \dots, q_n)$ exists and is equal to q . Then

$$\begin{aligned}
h'(q) &= h(q) \\
&= h(f^{\mathcal{A}}(q_1, \dots, q_n)) \\
&= f^{\mathcal{B}}(h(q_1), \dots, h(q_n)) \\
&= f^{\mathcal{B}}(h'(q_1), \dots, h'(q_n)) \\
&= h'(f q_1 \dots q_n) .
\end{aligned}$$

Since $\Delta_{\mathcal{A}}$ is contained in the kernel of h' , so is the congruence generated by $\Delta_{\mathcal{A}}$. Thus $s \equiv t(\Delta_{\mathcal{A}})$ implies $h'(s) = h'(t)$, and we have a unique map $\hat{h} : \hat{\mathcal{A}} \rightarrow \mathcal{B}$ that agrees with h on \mathcal{A} .

The uniqueness of $\hat{\mathcal{A}}$ up to isomorphism follows directly from the universality property (2): if $\hat{\mathcal{A}}$ and $\hat{\mathcal{A}}'$ are two free total extensions of \mathcal{A} , then there are unique homomorphisms between $\hat{\mathcal{A}}$ and $\hat{\mathcal{A}}'$ in either direction, and these must be inverses. \square

We have actually shown that the construction $\mathcal{A} \mapsto \hat{\mathcal{A}}$ constitutes a left adjoint to the inclusion functor from the category of total Σ -algebras and Σ -algebra homomorphisms to the category of partial Σ -algebras and partial Σ -algebra homomorphisms.

4 Essential Elements

To get a one-to-one correspondence in the Correspondence Lemma, we had to delete inaccessible states from the automaton. We will have to do that here as well, but we will also have to delete other states that are inessential for the construction of the free total extension.

Intuitively, an element of a total Σ -algebra \mathcal{A} is *essential* if it is a source of nonfreeness. For example, q is essential if $q = f^{\mathcal{A}}(p) = g^{\mathcal{A}}(r)$ and $f \neq g$, or if $q = f^{\mathcal{A}}(q)$. This will imply that q must be contained in any partial subalgebra of \mathcal{A} having \mathcal{A} as its free total extension. Moreover, we will show that under a mild restriction on how \mathcal{A} is generated, the induced partial subalgebra of \mathcal{A} on the set of its essential elements has \mathcal{A} as its free total extension. Thus the induced partial subalgebra on the essential elements of \mathcal{A} is the unique minimal partial subalgebra of \mathcal{A} having \mathcal{A} as its free total extension.

A unary function $|\mathcal{A}| \rightarrow |\mathcal{A}|$ is said to be *definable* (in \mathcal{A}) if it is of the form $\lambda x.t$ where $x \notin \Sigma$ is a nullary variable, t is a term over $\Sigma \cup \{x\}$, and the function symbols $f \in \Sigma$ occurring in t are interpreted as $f^{\mathcal{A}}$.

Definition 6 Let \mathcal{A} be a total Σ -algebra. An element $q \in \mathcal{A}$ is said to be *essential* if any of the following five conditions hold:

- (i) $q \neq f^{\mathcal{A}}(q_1, \dots, q_n)$ for any $n \geq 0$, $f \in \Sigma_n$ and $q_1, \dots, q_n \in \mathcal{A}$
- (ii) $q = f^{\mathcal{A}}(p_1, \dots, p_m) = g^{\mathcal{A}}(q_1, \dots, q_n)$ and $f \neq g$
- (iii) $q = f^{\mathcal{A}}(p_1, \dots, p_n) = f^{\mathcal{A}}(q_1, \dots, q_n)$ and $p_i \neq q_i$ for some i , $1 \leq i \leq n$
- (iv) $q = F(q)$ for some definable unary function $F = \lambda x.t$ on \mathcal{A} , and $t \neq X$
- (v) $p = F(q)$ for some definable unary function F on \mathcal{A} and p is essential.

(Note that the definition is inductive because of this clause.)

We define \mathcal{EA} to be the induced partial subalgebra of \mathcal{A} on the set of essential elements of \mathcal{A} . The partial algebra \mathcal{EA} is called the *essential subalgebra* of \mathcal{A} . An element of a partial algebra \mathcal{A} is said to be *essential* if it is an essential element of $\widehat{\mathcal{A}}$. (This definition does not conflict if \mathcal{A} is total, since in this case $\widehat{\mathcal{A}} \cong \mathcal{A}$.) \square

Definition 7 Let \mathcal{A} be a total Σ -algebra. A subset $Q \subseteq |\mathcal{A}|$ is a *generating set* if the canonical map $T_{\Sigma \cup Q} \rightarrow \mathcal{A}$ is onto. The set Q is a *minimal generating set* if it is a generating set and no subset of Q is a generating set.

\square

If \mathcal{A} is a partial algebra, then the null set is a generating set of $\widehat{\mathcal{A}}$ exactly when there are no inaccessible elements of \mathcal{A} , *i.e.*, when the canonical map $T_\Sigma \rightarrow \widehat{\mathcal{A}}$ is onto. Of course, in this case the null set is also a minimal generating set. Any algebra with a finite generating set has a minimal generating set. The integers with successor give an example of an algebra with no minimal generating set.

Lemma 8 *Let \mathcal{A} be a total Σ -algebra possessing a minimal generating set Q . Then every element of Q is essential.*

Proff. Let

$$\delta : T_{\Sigma \cup Q} \rightarrow \mathcal{A}$$

be the canonical map in which $\delta(q) = q$ for $q \in Q$. For any $q \in Q$, if the only term $t \in T_{\Sigma \cup Q}$ with $q = \delta(t)$ is q itself, then q is essential by Definition 6(i). Otherwise, there exists an n -ary function symbol f for some $n \geq 0$ and terms $t_1, \dots, t_n \in T_{\Sigma \cup Q}$ such that $q = \delta(ft_1 \dots t_n)$. If q occurs some term t_i , then q is essential by Definition 6(iv). If not, then $Q - \{q\}$ is a generating set, contradicting the assumption that Q was minimal. \square

The next theorem justifies the term “essential”. It shows that the essential elements of a total algebra \mathcal{B} must be contained in any partial subalgebra having \mathcal{B} as its free total extension.

Theorem 9 *Any partial algebra \mathcal{A} contains all essential elements of $\widehat{\mathcal{A}}$. Moreover, the partial algebra $\mathcal{E}\widehat{\mathcal{A}}$ is the induced partial subalgebra of \mathcal{A} on the set of essential elements of $\widehat{\mathcal{A}}$.*

Proof. Let $\mathcal{E} = \mathcal{E}\widehat{\mathcal{A}}$, let $t \mapsto [t]$ be the canonical map $T_{\Sigma \cup |\mathcal{A}|} \rightarrow \widehat{\mathcal{A}}$, and let $\nu(t)$ denote the $\Delta_{\mathcal{A}}$ -normal form of $t \in T_{\Sigma \cup |\mathcal{A}|}$. We show first that $|\mathcal{E}| \subseteq |\mathcal{A}|$. For any $e \in |\mathcal{E}|$, let $t \in T_{\Sigma \cup |\mathcal{A}|}$ be the unique term in $\Delta_{\mathcal{A}}$ -normal form with $[t] = e$.

If $e \in |\mathcal{E}|$ because of Definition 6(i), then t must be e itself. Thus $e \in |\mathcal{A}|$.

If $e \in |\mathcal{E}|$ because of Definition 6(ii), then there exist terms $fs_1 \dots s_m$ and $gt_1 \dots t_n$ with

$$\nu(fs_1 \dots s_m) = \nu(gt_1 \dots t_n) = t .$$

Since these two terms have distinct head symbols but the same normal form, we must have $t = e \in |\mathcal{A}|$.

If $e \in |\mathcal{E}|$ because of Definition 6(iii), then there exist terms $fs_1 \dots s_n$ and

$ft_1 \dots t_n$ with

$$\nu(fs_1 \dots s_n) = \nu(ft_1 \dots t_n) = t$$

but

$$\nu(s_i) \neq \nu(t_i)$$

for some i , $1 \leq i \leq n$. Again, in order for $fs_1 \dots s_n$ and $ft_1 \dots t_n$ to have the same normal form, we must have $t = e \in |\mathcal{A}|$.

If $e \in |\mathcal{E}|$ because of Definition 6(iv)7 then there exists a term s with exactly one occurrence of a variable x , but not x itself, such that

$$\nu(s[x/t]) = t .$$

Since s is not x itself, the depth of $s[x/t]$ is strictly greater than the depth of t . In order to reduce $s[x/t]$ to t , since t is in normal form, the occurrence of t in $s[x/t]$ must be an element of $|\mathcal{A}|$ and this element must be e .

Finally, if $e \in |\mathcal{E}|$ because of Definition 6(v), then there exists a term s with one occurrence of a variable x such that $s[x/t] \equiv p$ ($\Delta_{\mathcal{A}}$) and p is essential. By the induction hypothesis, $p \in |\mathcal{A}|$, so $\nu(s[x/t]) = p$. Therefore the occurrence of t in $s[x/t]$ must be an element of $|\mathcal{A}|$, and this element must be e .

We have shown that $|\mathcal{E}| \subseteq |\mathcal{A}|$. Since \mathcal{E} is the induced partial subalgebra of $\widehat{\mathcal{A}}$ on $|\mathcal{E}|$ and \mathcal{A} is the induced partial subalgebra of $\widehat{\mathcal{A}}$ on $|\mathcal{A}|$ (Theorem 5), it follows that the inclusion map $\mathcal{E} \rightarrow \mathcal{A}$ is a partial Σ -algebra homomorphism and that \mathcal{E} is the induced partial subalgebra of \mathcal{A} on $|\mathcal{E}|$. \square

Theorem 10 *Let \mathcal{A} be a total Σ -algebra with essential subalgebra $\mathcal{E} = \mathcal{E}\mathcal{A}$. Then $\widehat{\mathcal{E}}$ is embedded isomorphically in \mathcal{A} . Moreover, if \mathcal{A} contains a minimal generating set, then $\widehat{\mathcal{E}}$ and \mathcal{A} are isomorphic.*

Proof. By definition, $\mathcal{E} \sqsubseteq \mathcal{A}$. By Theorem 5, there exists a unique homomorphism $h : \widehat{\mathcal{E}} \rightarrow \mathcal{A}$ with h the identity on \mathcal{E} . We wish to show that h is injective.

Let $h' : T_{\Sigma \cup |\mathcal{E}|} \rightarrow \mathcal{A}$ be the canonical map with $h'(q) = q$ for $q \in |\mathcal{E}|$. We have the following commutative diagram:

$$\begin{array}{ccc}
T_{\Sigma \cup \mathcal{E}} & \xrightarrow{h'} & \\
\downarrow [\] & & \searrow \\
\hat{\mathcal{E}} & \xrightarrow{h} & \mathcal{A} \\
\sqsubseteq \uparrow & & \nearrow \\
\mathcal{E} & \xrightarrow{\sqsubseteq} &
\end{array}$$

We wish to show that for any $s, t \in T_{\Sigma \cup \mathcal{E}}$, if $h'(s) = h'(t)$ then $s \equiv t$ ($\Delta \mathcal{E}$).

We show first that if $t \in T_{\Sigma \cup \mathcal{E}}$ is in $\Delta_{\mathcal{E}}$ -normal form and $h'(t) = q \in |\mathcal{E}|$, then $t = q$. Suppose for a contradiction that $t = ft_1 \dots t_n$, $f \in \Sigma_n$, and t is of minimum depth. Since t is in $\Delta_{\mathcal{E}}$ -normal form, so are the t_i , $1 \leq i \leq n$, and

$$\begin{aligned}
q &= h'(ft_1 \dots t_n) \\
&= f^{\mathcal{A}}(h'(t_1), \dots, h'(t_n)) .
\end{aligned}$$

By Definition 6(v), $h'(t_i) \in |\mathcal{E}|$, say $h'(t_i) = q_i$. Since t was of minimum depth, $t_i = q_i$, $1 \leq i \leq n$. We thus have

$$q = f^{\mathcal{A}}(q_1, \dots, q_n) ,$$

thus

$$q \equiv fq_1 \dots q_n \in \Delta_{\mathcal{E}} ,$$

contradicting the assumption that t was in normal form.

Now let $s, t \in T_{\Sigma \cup \mathcal{E}}$ be in $\Delta_{\mathcal{E}}$ -normal form, and suppose $h'(s) = h'(t)$. We proceed by induction on the form of s and t .

If $s = q \in |\mathcal{E}|$, then $h'(s) = h'(t) = q$, thus $s = t = q$. The argument is similar for $t \in |\mathcal{E}|$. Otherwise, assume neither s nor t is in $|\mathcal{E}|$.

If $s = fs_1 \dots s_m$ and $t = gt_1 \dots t_n$ and $f \neq g$, then

$$\begin{aligned}
f^{\mathcal{A}}(h'(s_1), \dots, h'(s_m)) &= h'(fs_1 \dots s_m) \\
&= h'(gt_1), \dots, t_n) \\
&= g^{\mathcal{A}}(h'(t_1), \dots, h'(t_n)) ,
\end{aligned}$$

and $h'(s) \in |\mathcal{E}|$ by Definition 6(ii), contradicting the assumption that $h'(s) \notin |\mathcal{E}|$.

If $s = fs_1 \dots s_n$ and $t = ft_1 \dots t_n$, and if some $h'(s_i) \neq h'(t_i)$, then we obtain a contradiction as in the previous case, using Definition 6(iii).

Thus we are left with the case $s = fs_1 \dots s_n$, $t = ft_1 \dots t_n$, and $h'(s_i) = h'(t_i)$, $1 \leq i \leq n$. By the induction hypothesis, $s_i \equiv t_i (\Delta_{\mathcal{E}})$, $1 \leq i \leq n$, therefore: $s \equiv t (\Delta_{\mathcal{E}})$.

If \mathcal{A} contains a minimal generating set Q , then $Q \subseteq \mathcal{E}$ by Lemma 8, thus \mathcal{E} is also a generating set. Since \mathcal{E} also generates $\widehat{\mathcal{E}}$ the map h is onto in this case.

Corollary 11 *Let \mathcal{A} be a total Σ -algebra possessing a minimal generating set. Up to isomorphism, the essential \mathcal{EA} of \mathcal{A} is the unique minimal partial algebra having free total extension \mathcal{A} .*

The corollary is not true in general for algebras not possessing a minimal generating set. For example, consider a nonstandard model of the natural numbers with 0 and successor and the usual Peano axioms over this signature. There is no minimal set generating the nonstandard elements, and there are no essential elements. Thus the free total extension of the essential subalgebra consists of the standard natural numbers.

5 Partial Automata and Congruences

The following theorem our generalized version of the Correspondence Lemma.

Theorem 12 *Up to isomorphism, there is a one-to-one correspondence between (finitely generated) congruences on T_{Σ} and (finite) partial automata over T_{Σ} with no inaccessible and no inessential states.*

Proof. We establish a one-to-one correspondence between congruences on T_{Σ} and partial Σ -algebras with no inaccessible and no inessential elements, and show that a congruence is finitely generated iff its corresponding partial algebra is finite.

For a congruence \equiv on T_{Σ} , let $\mathcal{E} = \mathcal{E}(T_{\Sigma}/\equiv)$ be the essential subalgebra

of the quotient T_Σ/\equiv . Since the canonical map $T_\Sigma \rightarrow T_\Sigma/\equiv$ is onto, T_Σ/\equiv has minimal generating set \emptyset . By Theorem 10,

$$\widehat{\mathcal{E}} \cong T_\Sigma/\equiv ,$$

therefore \mathcal{E} has no inessential or inaccessible elements. Thus the map

$$\equiv \mapsto \mathcal{E}(T_\Sigma/\equiv) \tag{4}$$

takes congruences on T_Σ to partial Σ -algebras with no inaccessible and no inessential elements.

Conversely, let \mathcal{A} be a partial Σ -algebra with no inaccessible and no inessential elements, and let $\sim_{\mathcal{A}}$ be the kernel of the canonical map $\delta : T_\Sigma \rightarrow \widehat{\mathcal{A}}$. This construction gives a map

$$\mathcal{A} \mapsto \sim_{\mathcal{A}} \tag{5}$$

from partial Σ -algebras with no inaccessible and no inessential elements to congruences on T_Σ .

We now show that the maps (4) and (5) are inverses up to isomorphism. For any congruence \equiv on T_Σ , let $\mathcal{E} = \mathcal{E}(T_\Sigma/\equiv)$. Then \equiv and $\sim_{\mathcal{E}}$ are the same relations since δ is the unique homomorphism

$$\delta : T_\Sigma \rightarrow \widehat{\mathcal{E}} \cong T_\Sigma/\equiv .$$

Conversely, for any partial Σ -algebra \mathcal{A} with no inaccessible or inessential elements, we wish to show that \mathcal{A} and $\mathcal{E} = \mathcal{E}(T_\Sigma/\sim_{\mathcal{A}})$ are isomorphic. We have by Theorem 9 that $\mathcal{E}\widehat{\mathcal{A}}$ is the induced partial subalgebra of \mathcal{A} on $|\mathcal{E}|$. Since \mathcal{A} has no inessential elements,

$$\mathcal{A} \cong \mathcal{E}\widehat{\mathcal{A}} .$$

Since \mathcal{A} has no inaccessible elements, the canonical map $\delta : T_\Sigma \rightarrow \widehat{\mathcal{A}}$ is onto, thus

$$\widehat{\mathcal{A}} \cong T_\Sigma/\sim_{\mathcal{A}} ,$$

therefore

$$\mathcal{E}\widehat{\mathcal{A}} \cong \mathcal{E}(T_\Sigma/\sim_{\mathcal{A}}) .$$

Finally, we show

(i) if \mathcal{A} is finite, then $\sim_{\mathcal{A}}$ is finitely generated

(ii) if Γ is a finite relation on T_{Σ} then $\mathcal{E}(T_{\Sigma}/\Gamma)$ is finite.

First (i). If \mathcal{A} is finite, then so is $\Delta_{\mathcal{A}}$. Since $\delta : T_{\Sigma} \rightarrow \widehat{\mathcal{A}}$ is onto, for each $q \in |\mathcal{A}|$ there exists a $\eta(q) \in T_{\Sigma}$ such that $\delta(\eta(q)) \equiv q$ ($\Delta_{\mathcal{A}}$). The map η extends uniquely to a homomorphism $\eta : T_{\Sigma \cup |\mathcal{A}|} \rightarrow T_{\Sigma}$, and by uniqueness of the maps we have that the diagram

$$\begin{array}{ccc} T_{\Sigma \cup |\mathcal{A}|} & \xrightarrow{[\]} & \\ \eta \downarrow & \searrow & T_{\Sigma \cup |\mathcal{A}|} / \Delta_{\mathcal{A}} \cong \widehat{\mathcal{A}} \\ T_{\Sigma} & \xrightarrow{\delta} & \end{array}$$

commutes. Thus for $s, t \in T_{\Sigma \cup |\mathcal{A}|}$,

$$\begin{aligned} s \equiv t(\Delta_{\mathcal{A}}) &\leftrightarrow [s] = [t] \\ &\leftrightarrow \delta(\eta(s)) = \delta(\eta(t)) \\ &\leftrightarrow \eta(s) \sim_{\mathcal{A}} \eta(t) . \end{aligned}$$

We now show that $\sim_{\mathcal{A}}$ is generated by the finite relation

$$\eta(\Delta_{\mathcal{A}} = \{\eta(s) \equiv \eta(t) \mid s \equiv t \in \Delta_{\mathcal{A}}\})$$

on T_{Σ} . Certainly the congruence on T_{Σ} generated by $\eta(\Delta_{\mathcal{A}})$ is contained in $\sim_{\mathcal{A}}$ since $\eta(\Delta_{\mathcal{A}})$ is, and a straightforward inductive argument shows that for any $s, t \in T_{\Sigma \cup |\mathcal{A}|}$,

$$s \equiv t(\Delta_{\mathcal{A}}) \rightarrow \eta(s) \equiv \eta(t) \ (\eta(\Delta_{\mathcal{A}})) .$$

In particular for $s, t \in T_{\Sigma}$, we have $s = \eta(s)$ and $t = \eta(t)$, thus

$$\begin{aligned} s \sim_{\mathcal{A}} t &\leftrightarrow s \equiv t(\Delta_{\mathcal{A}}) \\ &\leftrightarrow s \equiv t(\eta(\Delta_{\mathcal{A}})) . \end{aligned}$$

To show (ii), let Γ be a finite relation on T_{Σ} . Define a finite partial Σ -algebra \mathcal{A} as follows. Let $t \mapsto [t]$ be the canonical map $T_{\Sigma} \rightarrow T_{\Sigma}/\Gamma$. Call the

term t present in Γ if t is a subterm of some u or v appearing in an equation $u \equiv v \in \Gamma$. Let \mathcal{A} be the induced partial subalgebra of T_Σ/Γ on the set

$$\{[t] \mid t \text{ is present in } \Gamma\} .$$

By Theorem 5, the inclusion map $\mathcal{A} \rightarrow T_\Sigma/\Gamma$ extends uniquely to a homomorphism $h : \widehat{\mathcal{A}} \rightarrow T_\Sigma/\Gamma$. Let δ be the canonical map $T_\Sigma \rightarrow \widehat{\mathcal{A}}$. We have the commutative diagram

$$\begin{array}{ccc} T_\Sigma & \xrightarrow{[\]} & T_\Sigma/\Gamma \\ \delta \downarrow & \searrow h & \\ \widehat{\mathcal{A}} & \xrightarrow{h} & T_\Sigma/\Gamma \\ \sqsubseteq \uparrow & \nearrow \sqsubseteq & \\ \mathcal{A} & & \end{array}$$

We show that h is an isomorphism. It is certainly onto, since $[\]$ is. To show that it is one-to-one, it suffices to show that δ is onto and for $s, t \in T_\Sigma$, $s \equiv t (\Gamma)$ implies $\delta(s) = \delta(t)$.

A straightforward inductive argument shows that $\delta(t) = [t]$ for t present in Γ : if $ft_1 \dots t_n$ is present in Γ then

$$[ft_1 \dots t_n] \equiv f[t_1] \dots [t_n] \in \Delta_{\mathcal{A}} ,$$

therefore

$$\begin{aligned} \delta(ft_1 \dots t_n) &= f^{\widehat{\mathcal{A}}}(\delta(t_1), \dots, \delta(t_n)) \\ &= f^{\widehat{\mathcal{A}}}([t_1], \dots, [t_n]) \\ &= [ft_1 \dots t_n] . \end{aligned}$$

Since $\widehat{\mathcal{A}}$ is generated by $|\mathcal{A}|$, δ is onto. Now if $s \equiv t \in \Gamma$, then $[s] = [t] \in |\mathcal{A}|$, and $\delta(s) = \delta(t) = [s]$. Since the relation Γ is contained in the kernel of δ so is the congruence generated by Γ . Thus $s \equiv t (\Gamma)$ implies $\delta(s) = \delta(t)$.

By Theorem 9, the essential subalgebra $\mathcal{E}(T_\Sigma/\Gamma)$ is contained in \mathcal{A} and is therefore finite. \square

The following theorem was essentially proved in [9] and [10, Lemma 25], to which we refer the reader for the algorithm and proof of correctness.

Theorem 13 ([9, 10]) *Given any finite relation Γ on T_Σ , the diagram $\Delta_\mathcal{E}$ of $\mathcal{E} = \mathcal{E}(T_\Sigma/\Gamma)$ can be produced from Γ in polynomial time.*

By Corollary 11, $\Delta_\mathcal{E}$ gives a canonical presentation of the finitely presented algebra T_Σ/Γ .

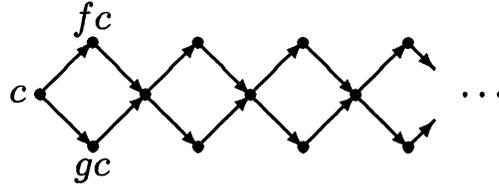
6 A Counterexample

Let R be an equivalence relation on T_Σ . Although the relation \equiv_R is the coarsest congruence refining R , it may not be finitely generated, even though there always exists a finitely generated congruence refining R (namely the identity). Thus the analog of clause (iii) in the statement of the Myhill-Nerode Theorem fails for partial automata.

It suffices to construct a congruence R on T_Σ that is not finitely generated (then \equiv_R and R coincide). Suppose we have a single nullary operator c and two unary operators f and g . Define $|c| = 0$ and $|ft| = |gt| = 1 + |t|$. Let Γ be the set

$$\Gamma = \{s \equiv t \mid |s| = |t| \text{ and } |s| \text{ is even}\}$$

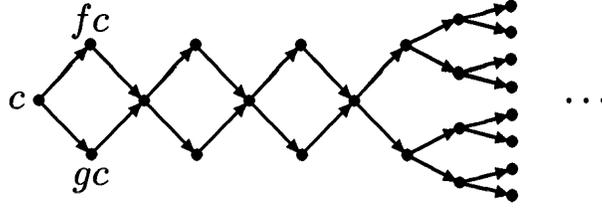
and let R be the congruence generated by Γ . Then T_Σ/Γ looks like this:



The congruence R is not finitely generated, since any finite subrelation Δ of R is contained in the congruence generated by some Γ_n , where

$$\Gamma_n = \{s \equiv t \mid |s| = |t| \leq n \text{ and } |s| \text{ is even}\} ,$$

thus T_Σ/Δ is a homomorphic preimage of T_Σ/Γ_n , which looks like this:



7 Applications to Term Rewrite Systems

Theorems 12 and 13 have the following application to term rewrite systems. Suppose we are given a ground term rewrite system over Σ . Let Q be a new set of auxiliary constants disjoint from Σ . Let us call a ground term rewrite system over $\Sigma \cup Q$ *simple* if

- all rules are of the form $f q_1 \dots q_n \rightarrow q$, where $q_1 \dots q_n, q, \in Q$ and $f \in \Sigma_n$;
- the system is unambiguous in the sense that there are no overlapping redexes.

A system over $\Sigma \cup Q$ is said to be *equivalent* to the original system over Σ if they induce the same congruence on T_Σ .

Theorems 12 and 13 have the following interpretation in this context:

Corollary 14 *For every ground term rewrite system Γ over Σ , there is a unique minimal simple system Γ' equivalent to Γ . Moreover, Γ' can be constructed from Γ in polynomial time.*

The system Γ' is of course just $\Delta_{\mathcal{E}}$, where \mathcal{E} is the essential subalgebra of T_Σ/Γ .

It was shown in [9, 10] that the problem of isomorphism of finitely presented algebras is equivalent to the problem of graph isomorphism. Essentially, Corollary 11 says that a finitely presented Σ -algebra is uniquely represented by its essential subalgebra, which is uniquely represented by its diagram,

which in turn can be represented as a labeled graph in a straightforward way. Conversely, the graph isomorphism problem is easily encoded as a problem of isomorphism of finitely presented algebras [9, 10].

In the construction given in [9, 10], it is readily observed that the degree of the graph is linear in the maximum arity in Σ ; thus using a result of Luks [12], there is a polynomial time algorithm to decide equivalence of ground term rewriting systems over Σ of bounded arity. In the case of unbounded arity, the problem is as hard as determining the isomorphism of graphs of unbounded degree.

References

- [1] M. A. Arbib and Y. Give'on, "Algebra automata I: parallel programming as a prolegomena to the categorical approach," *Inf. and Control* 12 (1968), 331–345.
- [2] W. S. Brainerd. *Tree Generating Systems and Tree Automata*. PhD thesis, Purdue University, 1967.
- [3] W. S. Brainerd, "The minimalization of tree automata," *Inf. and Control* 13 (1968), 484–491.
- [4] N. Dershowitz and J.-P. Jouannaud, "Rewrite Systems," in: J. van Leeuwen (ed.), *Handbook of Theoretical Computer Science*, vol. B, North Holland, Amsterdam, 1990, 243–320.
- [5] S. Eilenberg and J. B. Wright, "Automata in general algebra," *Inf. and Control* 11 (1967), 452–470.
- [6] J. Engelfriet. *Tree automata and tree grammars*. Technical Report FN-10, Computer Science Department, University of Aarhus, Denmark, April 1975.
- [7] Z. Fülöp Ed S. Vágvölgyi, "Congruential tree languages are the same as recognizable tree languages—a proof for a theorem of D. Kozen," *Bull. Eur. Assoc. Theor. Comput. Sci.* 39 (October 1989), 175–184.
- [8] J. E. Hopcroft and J. D. Ullman. *Introduction to Automata Theory, Languages, and Computation*. Addison-Wesley, 1979.

- [9] D. Kozen, “Complexity of finitely presented algebras,” PhD thesis, Cornell University, May 1977.
- [10] D. Kozen, “Complexity of finitely presented algebras,” *Proch 9th ACM Symp. Theory of Comput.*, May 1977, 164–177.
- [11] D. Kozen, “On the Myhill-Nerode Theorem for Trees,” *Bull. Assoc. Theor. Comput. Sci.*, submitted.
- [12] Ed Luks, “Isomorphism of graphs of bounded valence can be tested in polynomial time,” *J. Comput. Syst. Sci.* 25 (1982), 42–65.
- [13] J. Myhill, “Finite automata and the representation of events,” WADC Tech. Rep. 57-264, 1957.
- [14] A. Nerode, “Linear automata transformations,” *Proc, Amer. Math. Soc.* 9 (1958), 541–544.
- [15] J. W. Thatcher and J. B. Wright, “Generalized finite automata theory with an application to a decision problem of second order logic,” *Math. Syst. Theory* 2 (1968), 57–81.