An Analogue of the Myhill-Nerode Theorem and Its Use in Computing Finite-Basis Characterizations*

(Extended Abstract)

Michael R. Fellows
Dept of Computer Science
University of Idaho
Moscow, ID 83843

Michael A. Langston
Dept of Computer Science
University of Tennessee
Knoxville, TN 37996

Abstract.

Advances in the theory of well-partially-ordered sets now make it possible to prove the existence of low-degree polynomial-time decision algorithms for a vast assortment of natural problems, many of which seem to resist more traditional means of complexity classification. Surprisingly, these proofs are *nonconstructive*, based on the promise of an unknown but finite obstruction set.

Recent progress has yielded constructivization strategies that, for most applications, allow the desired decision (and search) algorithms to be *known* [FL3], despite the nonconstructive nature of the underlying mathematical tools on which the existence of these algorithms is based. These constructivizations produce algorithms that rely on the finiteness of an obstruction set, yet they ensure no means for computing or even verifying a candidate set.

The main purpose of this paper is to prove a theorem that is a graph-theoretic analogue of the Myhill-Nerode characterization of regular languages. We employ this result to establish that, for many applications, obstruction sets *are* computable by known algorithms.

* This research is supported in part by the National Science Foundation under grant MIP-8603879, by the Office of Naval Research under contracts N00014-88-K-0343 and N00014-88-K-0456, and by the National Aeronautics and Space Administration under engineering research center grant NAGW-1406.

1. Introduction

The primary result of this paper is a theorem that is strikingly analogous to the Myhill-Nerode characterization of regular languages, which we use to construct algorithms for computing the obstruction sets (finite bases) for many minor-closed and immersion-closed graph families. Also, as a byproduct of our methods, we achieve decision algorithms that are asymptotically faster than those previously reported for most of these problems.

A major motivation for the approach we take is a desire to mechanize obstruction set identification and verification, thereby avoiding the need for (100 page plus!) heroic case analyses. See, for example, [Ar, Ki, Sa].

We note that, of course, as long as a decision algorithm is known for a closed family, the obstruction set can be enumerated. The challenge is to determine when the entire set has been identified (when to halt the enumeration). The algorithms we present here accomplish just that, halting after an obstruction set has been captured. Yet they still possess an intriguing element of nonconstructivity: we can prove no bound on their halting time.

The remainder of this paper is organized as follows. In the next section, we state our main theorem (modulo notation and definitions to be developed later). Preliminaries are introduced and a weak version of the main theorem is proved in Section 3. In Section 4, we complete the proof of the main theorem. We illustrate its use with sample applications in Section 5. The final section consists of a few concluding remarks.

CH2806-8/89/0000/0520/$01.00 © 1989 IEEE

2. Overview of the Main Theorem

Given $L \subseteq \Sigma^*$, the canonical equivalence relation $\sim_L$ is: $x \sim_L y$ if and only if, $\forall z \in \Sigma^*, xz \in L \iff yz \in L$. An equivalence relation $\sim$ is a *congruence* if, $\forall z \in \Sigma^*, x \sim y \Longrightarrow xz \sim yz$ and $zx \sim zy$. The theorem below follows immediately from the well-known work of Myhill and Nerode.

Theorem A (Myhill-Nerode). Let L denote a subset of Σ^*. If an equivalence relation $\sim$ that is decidable with a known algorithm satisfies:

(1) $\sim$ has finite index,

(2) $\sim$ is a congruence with respect to concatenation, and

(3) $x \sim y \Longrightarrow x \sim_L y$,

then (L is regular and) an algorithm is known to compute a DFA that recognizes L.

The theorem that follows is a formal analogue of Theorem A. We state the result in terms of the minor order for notational convenience, but observe that it immediately applies to other RS posets as well (for example, the immersion order). Also, we postpone the formal definition of t-boundaried graphs and t-terminal operators until later, noting for now only that the former roughly corresponds to graphs with a fixed set of t distinguished, labeled vertices and that the latter provides a means for describing the recursive generation of the set of all graphs of tree-width at most t.

Theorem 2 (Main Result). Let F denote a minor-closed family, with a known membership algorithm, and a known bound on the tree-width of its obstruction set O. If an equivalence relation $\sim$ on t-boundaried graphs that is decidable with a known algorithm satisfies:

(1) $\sim$ has finite index,

(2) $\sim$ is a congruence with respect to any t-terminal operation, and

(3) $x \sim y \Longrightarrow x \sim_F y$,

then (F can be recognized in cubic time and) an algorithm is known to compute O.

3. Preliminary Results

Definition. A t-boundaried graph $G = (V, E, B, f)$ is an ordinary graph $G = (V, E)$ that is allowed to have loops and multiple edges, together with (1) a distinguished subset of the vertex set $B \subseteq V$ with $|B| = t$, and (2) a bijection $f : B \to \{1, \ldots, t\}$. We term B the boundary of G. If G is a boundaried graph, we may also simply write G to denote the underlying ordinary graph $G = (V, E)$. Where there is any possibility of confusion, we will write $\underline{G}$ to denote the underlying ordinary graph.

Definition. If $G = (V, E, B, f)$ and $G' = (V', E', B', f')$ are t-boundaried graphs, then $G \oplus G'$ denotes the ordinary graph obtained from the disjoint union of the graphs $G = (V, E)$ and $G' = (V', E')$ by identifying each vertex $u \in B$ with the vertex $v \in B'$ for which $f(u) = f'(v)$.

Informally, a t-boundaried graph is just an ordinary graph with a set of t distinguished vertices uniquely labeled with $\{1, \ldots, t\}$, and for boundaried graph G and H, $G \oplus H$ is obtained by gluing G and H together according to the labeling of the boundaries. Note that $G \oplus H \cong H \oplus G$. The lemma below follows immediately from the definition. (We write $X \geq_m Y$ if X and Y are ordinary graphs and X is greater than or equal to Y in the minor ordering.)

Lemma 1. If X, Y are t-boundaried graphs with $\underline{X} \geq_m Z$ for an ordinary graph Z, then $X \oplus Y \geq_m Z$.

If F is an arbitrary family of (ordinary) graphs, we define the following canonical equivalence relation on the set of t-boundaried graphs.

Definition. $X \sim_F Y$ if and only if, for every t-boundaried graph Z, $X \oplus Z \in F \iff Y \oplus Z \in F$.

This yields a natural decision problem.

F-CONGRUENCE

Input: t-boundaried graphs X and Y

Question: Is $X \sim_F Y$?

We use in an essential way a version of the minor ordering appropriate to t-boundaried graphs.

Definition. If $G = (V, E, B, f)$ and $G' = (V', E', B', f')$ are t-boundaried graphs, then $G \geq_m G'$ if and only if a t-boundaried graph isomorphic to G' can be obtained from G by a sequence of operations chosen from (1) delete an edge, (2) delete a vertex of $V - B$, and (3) contract an edge uv for which $|\{u, v\} \cap B| \leq 1$ (if, without loss of generality, $u \in B$, then the resulting vertex retains the label of u assigned by f).

Thus the minor order for boundaried graphs is defined just like the ordinary minor order, except that we require the boundary set of vertices B to be held fixed. The following lemma is an immediate consequence of the definition above.

Lemma 2. If G', G, H', H are t-boundaried graphs for which $G' \geq_m G$ and $H' \geq_m H$ then $G' \oplus H' \geq_m G \oplus H$. □

We make use of the theory of graph families generated by t-terminal composition operations (see [Wi]).

Definition. A binary t-terminal composition operator $\otimes$ is defined by (1) a t-boundaried graph $T_\otimes = (V_\otimes, E_\otimes, B_\otimes, f_\otimes)$ and (2) injective maps $f_i : \{1, \ldots, t\} \to V_\otimes$, for $i = 1, 2$. If $G_i = (V_i, E_i, B_i, f_i)$ is a pair of t-boundaried graphs, then $G_1 \otimes G_2$ is defined to be the t-boundaried graph for which the ordinary underlying graph is formed from the disjoint union of G_1, G_2 and $T_\otimes$ by identifying each vertex u of B_i (for $i = 1, 2$) with its image $f_i(u)$ in $V_\otimes$. The boundary set and labeling for $G_1 \otimes G_2$ is given by $B_\otimes$ and $f_\otimes$.

For a simple example, note that $\oplus$ is the composition operator with underlying graph $\overline{K}_t$ and with $f_1 = f_2$ the identity function. Crucial to our arguments are the theorems below.

Theorem B (Wimer). The family of boundaried graphs of tree-width at most k is recursively generated from a finite set of boundaried graphs with boundary size at most $k+1$ by a finite set of composition operators.

Theorem C (Robertson-Seymour). The set of all t-boundaried graphs is well-partially-ordered by $\geq_m$, and there are known polynomial-time order tests for $\geq_m$.

The next lemma generalizes Lemma 2 to arbitrary composition operators.

Lemma 3. If $\otimes$ is a composition operator and if G', G, H', H are t-boundaried graphs with $G' \geq_m G, H' \geq_m H$, then $G' \otimes H' \geq_m G \otimes H$.

Let F be a closed family in the minor order (on ordinary graphs). We define an associated partial order on t-boundaried graphs in terms of F-congruence. A key step in our argument will be to show that this associated order is a well-partial-order (for fixed t).

Definition. If G, H are t-boundaried graphs, then $G \geq_F H$ if and only if $G \geq_m H$ and $G \sim_F H$.

Lemma 4. If a decision algorithm is known for F-CONGRUENCE, then an algorithm is known for determining whether a t-boundaried graph G is minimal with respect to $\geq_F$.

Proof. Exhaustively check for each t-boundaried graph H below G in the order $\geq_m$ (these are easily enumerated and there are a finite number) whether there is a t-boundaried graph K with $G \oplus K \notin F, H \oplus K \in F$, using the decision algorithm for F-CONGRUENCE. □

Lemma 5. If G is a t-boundaried graph that is not minimal with respect to $\geq_F$, for F a closed family in the minor order of ordinary graphs, then $\underline{G} \notin O$ the obstruction set for F.

Proof. If G is not minimal, then $G \geq_m H, G \neq H$, for some H for which $G \oplus K \notin F$ implies $H \oplus K \notin F$. It follows by taking $K = \overline{K_t}$ that $\underline{G} = G \oplus K \geq_m H \oplus K = \underline{H}$, and $\underline{G} \notin F$ implies $\underline{H} \notin F$ so $\underline{G}$ is not a minimal element of the complement of F. □

The next lemma shows that $\geq_F$ is congruent with composition.

Lemma 6. If $G' \geq_F G$ and $H' \geq_F H$, then $G' \otimes H' \geq_F G \otimes H$.

Proof. By the definition of $\geq_F$ and Lemma 3, $G' \otimes H' \geq_m G \otimes H$. Suppose $(G' \otimes H') \oplus K \notin F$, but $(G \otimes H) \oplus K \in F$. By the definition of composition operators it is straightforward to describe K' such that $(G \otimes H) \oplus K = G \oplus K'$

and $(G' \otimes H) \oplus K = G' \oplus K'$. Then $G \oplus K' \in F$ implies $G' \oplus K' \in F$ since $G' \geq_F G$. Similarly, it is easy to describe K'' such that $(G' \otimes H) \oplus K = H \oplus K''$ and $(G' \otimes H') \oplus K = H' \oplus K''$. Then $G' \oplus K' = H' \oplus K'' \in F$ implies $H' \oplus K'' \in F$ since $H' \geq_F H$, a contradiction. □

Corollary. If G is not $\geq_F$ minimal, then for all $H, G \otimes H$ is not $\geq_F$ minimal.

Lemma 7. For each fixed t, $\geq_F$ is a well-partial-ordering of the set of t-boundaried graphs for F a closed family in $\geq_m$.

Proof. Clearly, there are no infinite descending chains. Suppose A is an infinite antichain. By a standard argument (see [Mi]) we can obtain an infinite ascending chain in the minor order $\geq_m$ of elements of A. Let $X_0, X_1, \ldots$ denote this chain. For $i < j$ there must be a t-boundaried graph Y_{ij} such that $X_j \oplus Y_{ij} \notin F$, while $X_i \oplus Y_{ij} \in F$, else it is contradicted that A is an infinite antichain. Consider the sequence of such witnesses $Y_{01}, Y_{12}, Y_{23}, \ldots$. By Theorem C we must have indices $i < j \leq k < \ell$ with $Y_{ij} \leq_m Y_{k\ell}, j = i+1, \ell = k+1$. But then we have $X_i \oplus Y_{ij} \in F$, $X_j \oplus Y_{ij} \notin F$, $X_k \oplus Y_{k\ell} \in F$ and $X_j \leq_m X_k$. By Lemma 2 we must have $X_j \oplus Y_{ij} \leq_m X_k \oplus Y_{k\ell}$, contradicting that F is a closed family in the minor order of ordinary graphs. □

For the remainder of this section let F denote a closed family under $\geq_m$ for ordinary graphs with obstruction set O. Let t be a fixed positive integer so that, with Lemma 7, $\geq_F$ is a well-partial-ordering of the boundaried graphs with boundary size at most t. For the set of all such boundaried graphs, let M denote the set of $\geq_F$ minimal elements. Let Ω be a finite set of t-terminal composition operators that, from a finite set A of t-boundaried graphs, generates a class C with $O \subseteq \underline{C}$.

Definition. A set R of boundaried graphs is *stable* with respect to F and Ω if $\forall \otimes \in \Omega$ and $\forall G, H \in R$ such that $G \otimes H \notin R$, $G \otimes H$ is not $\geq_F$ minimal.

The motivation for the above definition is as follows. Knowing a bound on the tree-width of graphs in O, we use Theorem B to describe a process of generating *all* graphs of the relevant tree-width. If at some point in this generation process the set of generated (boundaried) graphs R is stable (which can be determined if we can decide $\geq_F$ minimality), then we know that all of the obstructions for F have been captured. We prove this in the next lemma.

Lemma 8. If R is a set of boundaried graphs, with $A \subseteq R$, that is stable with respect to F and Ω, then $O \cap \underline{R} = O$.

Proof. The inclusion $\subseteq$ is trivial. For the reverse, suppose $K \in O, K \notin \underline{R} \cap O$. By our assumptions, there is a boundaried graph L such that $\underline{L} = K$ and L can be expressed as a finite product with operators chosen from Ω and with arguments in R. Since $L \notin R$ this a nontrivial factorization, and so by Lemmas 3, 5 and the Corollary to Lemma 6 either $K \geq_m H \in O \cap \underline{R}$ and $K \neq H$, or L is not $\geq_F$ minimal. In either case we reach a contradiction. □

The main point is just that when the generation process stabilizes, it is useless to go any further in searching for F-obstructions. Anything further that can be generated is useless (and useless for any further applications of the operators) because it must be $\geq_F$ nonminimal. Since no further obstructions can be generated, they have all been captured. The next lemmas shows that at some point the generation process must stabilize.

Lemma 9. If $M \subseteq R$, then R is stable with respect to F and Ω.

Proof. Let $G, H \in R$ and suppose $G \otimes H \notin R$. Since $M \subseteq R$ we have $G \otimes H \geq_F K \in M \subseteq R$ and so $G \otimes H \neq K$ and $G \otimes H$ is not minimal. □

We conclude this section with a weak but relatively simple and illustrative version of our main result.

Theorem 1. If F is a minor-closed family of graphs for which the following are known:

(1) a decision algorithm for recognizing F,

(2) a decision algorithm for F-CONGRUENCE, and

(3) a bound on the tree-width of the obstructions for F,

then an algorithm is known that will compute the obstruction set for F.

Proof. Using Theorem B, starting from a finite set of t-boundaried graphs with boundary sizes at most some constant t we can generate a class C that contains O. By (1) we can recognize obstructions when they are found. By (2) and Lemma 4 we can determine whether the generation procedure has stabilized. By Lemma 9 it must eventually stabilize, and by Lemma 8 at that point all of the obstructions have been found. □

4. Proof of the Main Theorem

Theorem 2. Let F denote a minor-closed family, with a known membership algorithm, and a known bound on the tree-width of its obstruction set O. If an equivalence relation $\sim$ on t-boundaried graphs that is decidable with a known algorithm satisfies:

(1) $\sim$ has finite index,

(2) $\sim$ is a congruence with respect to any t-terminal operation, and

(3) $x \sim y \Longrightarrow x \sim_F y$,

then (F can be recognized in quadratic time and) an algorithm is known to compute O.

Proof. Define $X \geq Y$ if and only if $X \geq_m Y$ and $X \sim Y$. Because $\geq_m$ and $\sim$ are decidable by known algorithms, we know an algorithm to decide $\geq$. Suppose A is an infinite antichain in the order $\geq$. By (1), there is an infinite subset A' of A with all members of A' in the same equivalence class of $\sim$. Since $\geq_m$ is a well-partial-order, there must be $X, Y \in A'$ for which $X \geq_m Y$, and therefore $X \geq Y$, a contradiction. So $\geq$ is a well-partial-order. Let M be the (finite) set of minimal elements of the set of all t-boundaried graphs under $\geq$, where t is appropriate to the known bound on the tree-width of O. Thus, from a finite set of t-boundaried graphs, we can generate using operators in a finite set Ω a class C that contains O. If X is not minimal with respect to $\geq$, then by (3) it is not minimal with respect to $\geq_F$, and so by Lemma 5 its underlying graph is not an obstruction. We define stability as before, replacing $\geq_F$ with $\geq$. Since we can decide $\geq$, we can decide stabilization. By (2) and an easy argument analogous to that used in Lemma 8, at stabilization the obstruction set is captured. When M has been generated, stabilization is assured (recall Lemma 9). Because M is finite, stabilization is inevitable. □

5. Sample Applications

Just as is the case when attempting to apply Theorem A to demonstrate the regularity of some $L \subseteq \Sigma^*$, we find that the problem of establishing the computability (by a known algorithm) of an obstruction set is now reduced to that of finding a satisfactory equivalence relation (that is, one that meets the conditions required by Theorem 2).

For example, consider the union of two minor-closed families F_1 and F_2. Clearly, such a union is minor-closed as well, guaranteeing a finite-basis characterization. But surprisingly, and quite unlike the case for intersection, proving the computability (by a known algorithm) of the union's obstruction set from knowledge of the sets for F_1 and F_2 has resisted all previously-known attempts.

Theorem 3. If F_1 and F_2 are minor-closed families with known obstruction sets O_1 and O_2, and if a bound is known on the tree-width of the obstructions for $F = F_1 \cup F_2$, then an algorithm is known that will compute the obstruction set for F.

Proof. Using Theorem 2, we define $X \sim Y$ if and only if $X \sim_{F_1} Y$ and $X \sim_{F_2} Y$. Because O_1 and O_2 are known, a decision algorithm for F is known. Since $\sim_{F_1}$ and $\sim_{F_2}$ are decidable with known methods, $\sim$ is decidable with a known method. The index of $\sim$ is finite, because it is the intersection of the equivalence relations $\sim_{F_1}$ and $\sim_{F_2}$, each of finite index. That $\sim$ is a congruence with respect to t-terminal operators follows from Lemma 6 and the fact that the intersection of congruences is a congruence. To complete the application of Theorem 2, our only remaining task is to show that $X \sim Y$ implies $X \sim_F Y$. Suppose otherwise for some $X \sim Y$. Then (without loss of generality) there exists a Z such that $X \oplus Z \in F$ and $Y \oplus Z \notin F$. This means, by the definition of

F, that either (1) $X \oplus Z \in F_1$ and $Y \oplus Z \notin F_1$, or (2) $X \oplus Z \in F_2$ and $Y \oplus Z \notin F_2$. But (1) contradicts the assumption that $X \sim_{F_1} Y$ and (2) contradicts the assumption that $X \sim_{F_2} Y$. □

Other well-studied problems for which we can now employ Theorem 2, along with an appropriate equivalence relation, to prove obstruction set computability (by a known algorithm) include: k-VERTEX COVER, k-FEEDBACK VERTEX SET, k-MIN CUT LINEAR ARRANGEMENT, k-DISK DIMENSION [FL1], WITHIN-k VERTICES OF F [FL2], g-UNIT CYCLE AVOIDANCE [FL2], k-VERTEX INTEGRITY and k-EDGE INTEGRITY [CEF].

Moreover, only the lack of knowledge of an appropriate tree-width bound (there always *is* such a bound) at this time precludes the application of Theorem 2 to prove the computability (by a known algorithm) of Kuratowski characterizations for problems such as GENUS, LINKLESSNESS and several others.

6. Concluding Remarks

We emphasize that our concern here has focused exclusively on just what is computable (by a known algorithm) in principle, as opposed to what is computable in practice. We remind the unwary reader that this general line of investigation is fraught with truly astronomical constants of proportionality [RS1].

Also, we note that the need to know a bound on an obstruction set's tree-width cannot be completely eliminated. Specifically, Theorem 16 of [FL3] can be extended to show that there is no algorithm to compute, from a finite description of a minor-closed family F (as represented by a Turing machine that accepts precisely the graphs in F), *any* bound on the tree-width of the obstruction set for F.

Clearly, the results we have reported here have the potential to open up new avenues of inquiry into the limits of nonconstructivity *inherent* in well-partial-order-based complexity tools. We find it especially satisfying to see that such a classic theorem from formal language theory has a useful analogue in this novel research domain.

References

[Ar] D. Archdeacon, "A Kuratowski Theorem for the Projective Plane," Ph.D. Thesis, Ohio State University, 1980.

[CEF] L. Clark, R. C. Entringer and M. R. Fellows, "Computational Complexity of Integrity," *J. Combin. Math. Combin. Comput.* 3 (1988), 549–560.

[FL1] M. R. Fellows and M. A. Langston, "Nonconstructive Advances in Polynomial Time Complexity," *Info. Proc. Letters* 26 (1987), 157–162.

[FL2] ——————, "Nonconstructive Tools for Proving Polynomial-Time Decidability," *J. of the ACM* 35 (1988), 727–739.

[FL3] ——————, "On Search, Decision and the Efficiency of Polynomial-Time Algorithms," *Proc. 21st ACM STOC* (1989), 501–512.

[Ki] N. G. Kinnersley, "Obstruction Set Isolation for Layout Permutation Problems," Ph. D. Thesis, Washington State University, 1989.

[Mi] E. Milner, "Basic WQO and BQO Theory," in Graphs and Orders (I. Rival, ed.), Reidel, Amsterdam, 1985.

[RS1] N. Robertson and P.D. Seymour, "Graph Minors V. Excluding a Planar Graph," *J. Comb. Th. Ser. B* 41 (1986), 92–114.

[RS2] ——————, "Graph Minors XIII. The Disjoint Paths Problem," to appear.

[RS3] ——————, "Graph Minors XVI. Wagner's Conjecture," to appear.

[Sa] H. Saran, "Constructive Results in Graph Minors: Linkless Embeddings," Ph. D. Thesis, University of California, Berkeley, 1989.

[Wi] T.V. Wimer, "Linear Algorithms on k-Terminal Graphs," Ph.D. Thesis, Clemson University, 1987.