On Computing Graph Minor Obstruction Sets*

Kevin Cattell
Department of Computer Science
University of Victoria
Victoria, B.C. V8W 3P6, Canada
kcattell@csr.uvic.ca

Michael J. Dinneen
Department of Computer Science
University of Auckland
Auckland, New Zealand
mjd@cs.auckland.ac.nz

Rodney G. Downey
Department of Mathematics
Victoria University
P.O. Box 600
Wellington, New Zealand
rod.downey@vuw.ac.nz

Michael R. Fellows
Department of Computer Science
University of Victoria
Victoria, B.C. V8W 3P6 Canada
mfellows@csr.uvic.ca

Michael A. Langston
Department of Computer Science
University of Tennessee
Knoxville, Tennessee 37996-1301
U.S.A.
langston@cs.utk.edu

22 April 1997

Abstract

The Graph Minor Theorem of Robertson and Seymour establishes nonconstructively that many natural graph properties are characterized by a finite set of forbidden substructures, the *obstructions* for the property. We prove several general theorems regarding the computation of obstruction sets from other information about a family of graphs. The methods can be adapted to other partial orders on graphs, such as the immersion and topological orders. The algorithms are in some cases practical and have been implemented. Two new technical ideas are introduced. The first is a method of computing a stopping signal for search spaces of expanding pathwidth. This allows obstruction sets to be computed for the first time without the necessity of a prior bound on maximum obstruction width. The second idea is that of a *second order congruence* for a graph property. This is an equivalence relation defined on finite sets of graphs that generalizes the recognizability congruence that is defined on single graphs. It is shown that the obstructions for a graph ideal can be effectively computed from an oracle for the canonical second-order congruence for the ideal and a membership oracle for the ideal. It is shown that the obstruction set for a union $\mathcal{F} = \mathcal{F}_1 \cup \mathcal{F}_2$ of minor ideals can be computed from the obstruction sets for $\mathcal{F}_1$ and $\mathcal{F}_2$ if there is at least one tree that does not belong to the intersection of $\mathcal{F}_1$ and $\mathcal{F}_2$. As a corollary, it is shown that the set of intertwines of an arbitrary graph and a tree are effectively computable.

*Some of the results of this paper were presented at the 1989 IEEE Symposium on the Foundations of Computer Science [FL89b].

1 Introduction

The celebrated Graph Minor Theorem (GMT) of Robertson and Seymour [RS83, RS85, RS94] proves the existence of finite obstruction sets for arbitrary minor order ideals, of which there are many natural examples. Planar graphs are famously an ideal for which the obstructions are $K_{3,3}$ and K_5 (Kuratowski's Theorem). The proof of the GMT is not effective, in the sense that knowing only a decision procedure for a lower ideal $\mathcal{F}$ does not provide enough information to be able to compute the obstruction set for $\mathcal{F}$ [FL89a]. For this reason the Graph Minor Theorem is commonly regarded as "nonconstructive" since usually we know at least a decision algorithm for a natural ideal.

The purpose of this paper is to explore the question:

> *What combinations of information about an ideal $\mathcal{F}$ allow us to effectively compute the obstruction set for $\mathcal{F}$?*

We are currently far from having a satisfactory account of this issue. Some of the open problems that remain in this area are both elegant and apparently difficult. Previous work can be summarized as follows.

(1) Fellows and Langston proved in [FL89a] that there is no algorithm that will, provided with only a decision oracle for an ideal $\mathcal{F}$, compute the set of obstructions $\mathcal{O}$ for $\mathcal{F}$.

(2) Fellows and Langston proved in [FL89b] that if we have access to the three pieces of information:
(i) A decision algorithm for $\mathcal{F}$.
(ii) A bound B on the maximum treewidth (or pathwidth) of the $\mathcal{F}$ obstructions.
(iii) A decision algorithm for a finite index congruence that refines the canonical congruence for $\mathcal{F}$ on t-boundaried graphs (for $t = 1, \ldots, B$).
Then $\mathcal{O}$ can be computed. (The full argument is given here for the first time.)

A curious aspect of this result is that given (i) and (iii) as oracles (which is the assumption of the theorem), then it is impossible to calculate in advance *when* the procedure to calculate $\mathcal{O}$ will terminate, although the proof guarantees that the procedure will eventually halt, having correctly computed $\mathcal{O}$. In other words, the *stopping time* of the algorithm is nonconstructive. The proof employs the GMT for finitely edge-colored graphs to establish that the algorithm will halt, and this is the source of the nonconstructivity concerning the stopping time.

(3) Lagergren and Arnborg [LA91, **?**] showed that if we are given (i), (ii) and (iii) as above, and are additionally given:

(iv) A computable function $f(t)$ that bounds the index of the finite congruences of (iii).

Then it is possible to effectively compute in advance a stopping time for the above procedure and to remove the dependence of the termination argument on the GMT. This also means that given (i)–(iv) we can effectively compute a bound on the size of the largest obstruction, and from this information could simply compute $\mathcal{O}$ by exhaustive search.

(4) An important class of lower ideals for which we have the pieces of information (i), (iii) and (iv) are those that we know how to describe in Monadic Second Order (MSO) logic. In other words, if we are given the information:

(v) An MSO expression ϕ that describes the graphs of the lower ideal $\mathcal{F}$.

Then from this we can effectively derive (i), (iii) and (iv). This result is mainly due to Courcelle [Co90].

(5) Other work on the systematic computation of obstruction sets has appeared in [Pr93, APS90, CD94, CDF95, Kin94, KL91]. Some of these results support practical implementations that have led to some significant mechanical or partly-mechanical proofs of new and nontrivial forbidden substructure theorems.

There has been a considerable amount of overlapping work in this area which is sometimes confusing to sort through. The above review is framed from the point of view that we will develop further here. In particular, we are concerned with identifying those combinations of *abstract* information about a lower ideal $\mathcal{F}$ that either do or do not provide enough information to allow us to compute the obstruction set $\mathcal{O}$ for $\mathcal{F}$. It is important to be attentive to exactly how the information about $\mathcal{F}$ is presented. For example, in our Theorem 1 we prove the result (1) above that (i), (ii) and (iii) are enough to effectively compute $\mathcal{O}$. In this theorem, (i) and (iii) are hypothesized to be available only via oracles, i.e., (iii) is not assumed to be concretely available via a finite state machine or dynamic programming algorithm.

Part of our purpose in this paper is to articulate this area of research, which we believe to be an appealing blend of combinatorics and recursion theory. There are clear models of both positive and negative results in this area, with much that remains unresolved. For example, in view of (4), it is natural to ask whether (v) alone is sufficient information about $\mathcal{F}$ to allow us to compute $\mathcal{O}$. Ideally, we should be able to settle this one way or the other, either by proving a positive result along the lines of (2) — perhaps using the new techniques introduced in this paper — or by proving a negative result along the lines of (1).

One of the main ingredients of the positive result (2) is a collection of finite-index congruences. There are several notions of congruence in the literature of this area that in many situations are essentially equivalent or effectively interchangeable. The basic notion that we use is provided by the following definitions.

Definition. A *t-boundaried graph* $G = (V, E, B, f)$ is an ordinary graph $G = (V, E)$ together with:
(1) a distinguished subset of the vertex set $B \subseteq V$ of cardinality t, the *boundary* of G, and
(2) a bijection $f : B \to \{1, \ldots, t\}$.

In some situations, we will forget the boundary. (For example, if G is a boundaried graph and $\mathcal{F}$ is a family of ordinary graphs, we may write $G \in \mathcal{F}$, meaning by this that G belongs to $\mathcal{F}$ when the boundary of G is ignored.) A fundamental operation (denoted $\oplus$) on t-boundaried graphs is that of gluing them together along their boundaries by identifying like-labeled vertices.

Definition. If $G = (V, E, B, f)$ and $G' = (V', E', B', f')$ are t-boundaried graphs, then $G \oplus G'$ denotes the t-boundaried 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 $f(u) = f'(v)$.

In the sequel, we will consider both *large* and *small* universes of t-boundaried graphs. Many of the main issues concern the large universe, which is easier to think about.

Definition. The *large universe* $\mathcal{U}^t_{\text{large}}$ is the set of all t-boundaried graphs.

Definition. If $\mathcal{F}$ is a family of graphs then the *large canonical congruence* for $\mathcal{F}$ is defined for t-boundaried graphs $X, Y \in \mathcal{U}^t_{\text{large}}$ by $X \approx_{\mathcal{F}} Y$ if and only if

$$\forall Z \in \mathcal{U}^t_{\text{large}} : (X \oplus Z \in \mathcal{F}) \iff (Y \oplus Z \in \mathcal{F})$$

The following definition is from Abrahamson and Fellows [AF93].

Definition. A graph family F is *fully cutset regular* if for every t, the large canonical congruence on $\mathcal{U}^t_{\text{large}}$ has finite index.

Courcelle and Lagergren proved in [CL94] that this notion is equivalent to that of *recognizable* graph families introduced in [Co90]. This must be regarded as an extremely pretty idea, as it captures an essential feature of the complexity of the "information flow" across a bounded-size cutset necessary to determining membership in a graph family. We will refer to the large canonical congruence for a graph family $\mathcal{F}$ as the *canonical recognizability congruence* for $\mathcal{F}$.

It follows from the Graph Minor Theorem and Courcelle's Theorem on MSO graph properties [Co90] that every minor order lower ideal is recognizable. It is interesting that only a few natural graph families are presently known *not* to be recognizable [AF93, BFW92, FHW93].

The positive results of [FL89b] (our Theorem A) apply to many (if not most) natural lower ideals. "Normally" we have the information (i) about $\mathcal{F}$. Note that since we are concerned here with the issue of whether $\mathcal{O}$ can be recursively computed, *any* algorithm that correctly decides membership in $\mathcal{F}$ will serve for (i), i.e., the efficiency of the algorithm is not an issue. It is also the case that "usually" we can find (iii) constructively. The exceptions include the ideals associated with the problems Knotless Embedding and Planar Diameter Improvement described in [DF94].

By far the most problematic aspect of Theorem A is the bound (ii) on the maximum obstruction treewidth or pathwidth. For example, although a congruence for torus embedding is relatively easy to produce, a bound on the maximum obstruction width is much more difficult, although a (very large) bound is now known. Tight bounds seem to be beyond current proof techniques in most situations. Thus it is natural to ask whether the information (ii) is really needed for obstruction set computations.

The basic computational machinery that we develop here shows how we can improve on the ideas of Theorem A and effectively compute $\mathcal{O}$ without having to know a prior bound on the maximum obstruction width. In this approach we use a *second order* congruence for a graph family — a finite-index equivalence relation defined on finite sets of boundaried graphs. Instead of having to prove a prior bound on obstruction width, it is necessary that this second-order congruence have an "eventual termination" property. Since termination can be established computationally, we believe this may be a significant breakthrough for implementations of obstruction set theorem-provers.

The Basic Approach

In [FL89b] (our Theorem A) the computation procedure uses (i) and (ii) to compute, for successive width bounds t, the set $\mathcal{O}^t$ of obstructions for $\mathcal{F}$ that have width (pathwidth or treewidth, either of these can be used) at most t. The main argument shows that for each fixed t this is a finite procedure. The role of hypothesis (ii) is simply to supply a bound on the maximum width t that needs to be considered.

Here we extend this procedure by computing $\mathcal{O}^t$ for successively larger t, and tackle the question: "Is $\mathcal{O} = \mathcal{O}^t$?" (i.e., Can we stop now?) computationally. We consider two different reasons for which the answer to the question "Can we stop now?" might be, "No":

(1) A *small* counterexample (to $\mathcal{O} = \mathcal{O}^t$) is an element of $\mathcal{O} - \mathcal{O}^t$ of width less than some known recursive function $f(t)$.

(2) A *large* counterexample is one whose width is more than $f(t)$.

Our computational strategy is based on the fact that large counterexamples can be easier to detect. For a particular recursive $f(t)$, we can determine whether there is a small counterexample to

$\mathcal{O} = \mathcal{O}_t$ by simply computing $\mathcal{O}^{t'}$ for $t' = t+1, ..., f(t)$, and checking that no new obstructions are found in these widths. To determine whether there is a large counterexample we compute an *alarm function* $\alpha : N \to \{0,1\}$. We interpret $\alpha = 1$ as a signal that we cannot stop at width t because there is (or may be) a counterexample. What we would like to have is an alarm function that simply determines whether there is a counterexample of width more than $f(t)$. But even this might not be possible, and so we relax our requirements. The hypothesis we employ instead of (iii) is that we can compute an alarm function α that is:
(a) *reliable*: if there is a counterexample of width more than $f(t)$ then $\alpha(f(t)) = 1$, and
(b) *eventually quiescent*: there is constant t_0 depending only on $\mathcal{F}$ such that $\forall t \geq t_0$, $\alpha(t) = 0$.
These weaker requirements allow for one-sided errors in answering the "t-stopping question" (with errors on the side of continuing the search), while insuring that only finitely many stages of such "false alarms" will be possible. This method is codified in Theorem B.

Theorem C provides a second general computational engine based on an alarm provided by a *terminating second-order congruence* (explained in the next section). This is in some sense a *specialization* of Theorem B. The naturality of the notion of a terminating second-order congruence is established by the following result:

Theorem D. If we have access to an oracle for the canonical second-order recognizability congruence for an ideal $\mathcal{F}$ and an oracle for membership in $\mathcal{F}$, then we can compute the obstruction set for $\mathcal{F}$.

We describe a natural second order congruence for the problem of computing the obstruction set for a union of ideals for which the obstructions are known, and show that this is a terminating congruence if at least one of the constituent ideals excludes a tree. As a corollary, we show that it is possible to effectively compute the topological intertwines of an arbitrary graph and a tree.

The main significance of Theorems B and C is in the new general techniques for obstruction set computation that we introduce. In particular, the notion of a *width stopping signal* seems to be of fundamental importance not only for the emerging theory of recursive aspects of the GMT, but also for practical implementations of obstruction set theorem provers.

This area of algorithmic graph theory has reached a depth where it is no longer possible for a paper to be entirely self-contained. We assume that the reader is familiar with the results of [Co90] and [AF93] and the basics of the theory of graph minors and well-quasiordering [RS85, NW63, FL88].

The plan of the paper is as follows. In the next section we deal with most of the preliminaries. In §3 we prove Theorem A. In §4 we prove Theorems B and C. In §5 we prove that the canonical second-order recognizability congruence terminates, Theorem D. In §6 we address the problem of computing the obstructions for unions and intertwines. In the final section we summarize and discuss some open problems.

2 Preliminaries

All of our discussion concerns finite simple graphs. A graph H is a *minor* of a graph G if a graph isomorphic to H can be obtained from G by a sequence of operations chosen from (i) delete a vertex, (ii) delete an edge, (iii) contract an edge, removing any multiple edges or loops that form. We write $G \geq_{\mathrm{m}} H$ to denote the minor order.

The *topological order* is defined $G \geq_{\mathrm{top}} H$ if and only if G contains a subgraph H' that is isomorphic to a subdivision of H, where a *subdivision* of a graph H is any graph that can be obtained from H by replacing edges by vertex disjoint paths. The topological order can be equivalently defined by using the definition of the minor order, only restricting operation (iii) to edges where at least one vertex has degree 2.

We may use the notation $\leq$ for simplicity where it is clear which order is under discussion. An *ideal* $\mathcal{J}$ in a partial order $(\mathcal{U}, \geq)$ is a subset of $\mathcal{U}$ such that if $X \in \mathcal{J}$ and $X \geq Y$ then $Y \in \mathcal{J}$. The *obstruction set* for $\mathcal{J}$ is the set of minimal elements of $\mathcal{U} - \mathcal{J}$.

If $\sim$ and $\approx$ are equivalence relations on a set $\mathcal{U}$, we say that $\sim$ *refines* $\approx$ if

$$\forall x, y \in \mathcal{U} : x \sim y \Longrightarrow x \approx y$$

We say that $\sim$ has *finite index* on $\mathcal{U}$ if there are a finite number of equivalence classes. The equivalence class of x with respect to $\sim$ is denoted $[x]_\sim$, or perhaps just $[x]$ where the equivalence relation is clear.

Definition. A *tree-decomposition* of a graph $G = (V, E)$ is a tree T together with a collection of subsets T_x of V indexed by the vertices x of T that satisfies:

1. (*Covering*) For every edge uv of G there is some x such that $\{u, v\} \subseteq T_x$.
2. (*Interpolation*) If y is a vertex on the unique path in T from x to z then $T_x \cap T_z \subseteq T_y$.

The *width* of a tree decomposition is the maximum over the vertices x of the tree T of the decomposition of $|T_x| - 1$. A graph G has *treewidth* at most k if there is a tree decomposition of G of width at most k. *Path-decompositions* and *pathwidth* are defined by restricting the tree T to be simply a path. The pathwidth of a graph G will be denoted $\mathrm{pw}(G)$.

There are several *universes* of boundaried graphs that we work with in this theory. The large universe has been defined in §1.

Definition. The *small treewidth universe* $\mathcal{U}^t_{\mathrm{tree}}$ is the set of all t-boundaried graphs having a tree-decomposition of width $t - 1$ for which the set of boundary vertices is the set of vertices indexed by the root of the tree. The *small pathwidth universe* $\mathcal{U}^t_{\mathrm{path}}$ is the set of all t-boundaried graphs having a path-decomposition of width $t - 1$ for which the set of boundary vertices is the last set of the decomposition.

We will write $\mathcal{U}^t_{\mathrm{small}}$ if it is a matter of indifference whether we mean $\mathcal{U}^t_{\mathrm{tree}}$ or $\mathcal{U}^t_{\mathrm{path}}$.

The following easy lemma is left to the reader.

Lemma 2.1 If A and B are t-boundaried graphs in $\mathcal{U}^t_{\mathrm{small}}$ then $A \oplus B$ has width less than or equal to t.

We extend the minor and topological orders to t-boundaried graphs by requiring that the boundary be held fixed in the operations defining the orders, and use the notation $\leq_{\mathrm{m}}$ and $\leq_{\mathrm{top}}$ to denote the boundaried orders (the context will make clear whether the graphs have boundaries or not). If $A \in \mathcal{U}^t_{\mathrm{large}}$, $\mathrm{int}(A)$ denotes the subgraph of A induced by the non-boundary vertices of A.

Definition. The *small canonical congruence* for $\mathcal{F}$ is defined for t-boundaried graphs $X, Y \in \mathcal{U}^t_{\mathrm{small}}$ by $X \sim_{\mathcal{F}} Y$ if and only if

$$\forall Z \in \mathcal{U}^t_{\mathrm{small}} : (X \oplus Z \in \mathcal{F}) \Longleftrightarrow (Y \oplus Z \in \mathcal{F})$$

(Note that there are two flavors, one for pathwidth and one for treewidth.)

Note that both $\approx_{\mathcal{F}}$ and $\sim_{\mathcal{F}}$ are defined on U^t_{small}. Trivially, $\approx_{\mathcal{F}}$ refines $\sim_{\mathcal{F}}$ on the small universe, but the two equivalence relations might not coincide. Courcelle and Lagergren have shown that on $\mathcal{U}^t_{\mathrm{tree}}$ the large canonical congruence has finite index if and only if the small canonical congruence has finite index [CL94].

We will make essential use of yet another kind of finiteness property that is exhibited by graph ideals. To put this notion in a familiar context, suppose that $L \subseteq \Sigma^*$ is a formal language. Then the canonical (Myhill-Nerode) congruence for L is defined: $x \approx_L y$ if and only if $\forall z \in \Sigma^* : [(xz \in L) \Longleftrightarrow (yz \in L)]$. A *test set* for L is a set of words $T \subset \Sigma^*$ such that if we define $x \sim_T y$ if and only if $\forall t \in T : [(xt \in L) \Longleftrightarrow (yt \in L)]$ then we get $x \sim_T y$ if and only if $x \approx_L y$. A language is regular if and only if it has a finite test set.

Now suppose $\mathcal{F}$ is an arbitrary family of graphs. A *t-concrete test set* for $\mathcal{F}$ is a set $\mathcal{T}^t \subseteq \mathcal{U}^t_{\text{large}}$ such that $\forall X, Y \in \mathcal{U}^t_{\text{large}}$ we have $X \approx_{\mathcal{F}} Y$ if and only if

$$\forall T \in \mathcal{T}^t : [(X \oplus T \in \mathcal{F}) \Longleftrightarrow (Y \oplus T \in \mathcal{F})]$$

Note that each concrete test graph T is used to define a predicate.

A *t-abstract test set* for $\mathcal{F}$ is a set of predicates $\mathcal{P}^t$ such that the equivalence relation defined on $\mathcal{U}^t_{\text{large}}$ by:

$$X \sim Y \quad \text{if and only if: } \forall P \in \mathcal{P}^t \;\; P(X) \leftrightarrow P(Y)$$

is a refinement of the canonical second-order congruence.

Definition. The *canonical second order congruence* for an ideal $\mathcal{F}$ (for convenience also denoted $\approx_{\mathcal{F}}$) is defined on finite sets of t-boundaried graphs in $\mathcal{U}^t_{\text{large}}$ by: if $S_1, S_2 \subseteq \mathcal{U}^t_{\text{large}}$ then $S_1 \approx_{\mathcal{F}} S_2$ if and only if

$$\forall Z \in \mathcal{U}^t_{\text{large}} : (\exists X_1 \in S_1 : X_1 \oplus Z \notin \mathcal{F}) \Longleftrightarrow (\exists X_2 \in S_2 : X_2 \oplus Z \notin \mathcal{F})$$

Definition. A (non-canonical) second-order congruence for $\mathcal{F}$ is an equivalence relation $\sim$ defined on finite subsets of $\mathcal{U}^t_{\text{large}}$ for which $S_1 \sim S_2$ implies $S_1 \approx_{\mathcal{F}} S_2$.

Let $A \in \mathcal{U}^t_{\text{large}}$. We will use the notation $S(A)$ to denote all the t-boundaried graphs properly below A in the boundaried minor order.

Definition. A second-order congruence $\sim$ for an ideal $\mathcal{F}$ is called *terminating* if it satisfies the condition: $\exists t_0$ such that $\forall t \geq t_0$, if $A \in \mathcal{U}^t_{\text{small}}$ such that: (1) $\text{pw}(A) \geq t_0$, and (2) $|\text{int}(A)| \geq t_0$, then $\{A\} \sim S(A)$.

In the §5 we will show that the canonical second-order congruence for a lower ideal $\mathcal{F}$ is terminating.

3 The Basic Computational Engine

In this section we prove the basic positive result on obstruction set computation for a fixed bound on the width of the search space. The proof was sketched in the extended abstract [FL89b].

Theorem A. Suppose that $\mathcal{F}$ is an ideal in the minor order of finite graphs and that we have the following three pieces of information about $\mathcal{F}$:
(1) An algorithm to decide membership in $\mathcal{F}$ (of any time complexity).
(2) A bound B on the maximum treewidth of the obstructions for for $\mathcal{F}$.
(3) For $t = 1, ..., B + 1$ a decision algorithm for a finite index right congruence $\sim$ on t-boundaried graphs that refines the small canonical congruence for $\mathcal{F}$.
Then we can effectively compute the obstruction set $\mathcal{O}$ for $\mathcal{F}$.

Proof. The algorithm is outlined as follows. For $t = 1, ..., B + 1$ we generate in a systematic way the t-boundaried graphs of $\mathcal{U}^{\text{tree}}_t$ until a certain *stop signal* is detected. At this point, for a given t, we will have generated a finite set of graphs $\mathcal{G}_t$. Of particular interest among these are the graphs $\mathcal{M}_t \subseteq \mathcal{G}_t$ that are minimal with respect to a certain partial order $\leq$ on t-boundaried graphs. We will prove that $\mathcal{O}$ is a subset of

$$\mathcal{M} = \bigcup_{t=1}^{B+1} \mathcal{M}_t$$

considering the graphs of $\mathcal{M}$ with the boundaries forgotten.

There are three things to be clarified:
(1) how the graphs of the small universe are generated,
(2) the *search ordering* $\leq$, and (3) the nature of the stop signal for width t.

(1) *The order of generation of* $\mathcal{U}_t^{\mathrm{small}}$.

Suppose X is a t-boundaried graph, $X \in \mathcal{U}_t^{\mathrm{tree}}$. By the *size* of X we refer to the number of nodes in a smallest possible indexing tree for a tree decomposition of X. For a given t, we generate the t-boundaried graphs of $\mathcal{U}_t^{\mathrm{small}}$ in order of increasing size. By the jth *generation* we refer to all of those graphs of size j in this process.

(2) *The search ordering* $\leq$.

To define $\leq$, we first extend the minor ordering of ordinary graphs to t-boundaried graphs in the natural way by holding the boundary fixed. In other words, the boundaried minor order is defined by the same local operations as the minor order, except that we are not allowed to delete boundary vertices or to contract edges between two boundary vertices. This can be easily shown to be a wqo on $\mathcal{U}_t^{\mathrm{large}}$ by using the Graph Minor Theorem for edge-colored graphs. Let $\leq_m$ denote the minor order on ordinary graphs and let $\leq_{\partial m}$ denote the boundaried minor order.

For $X, Y \in \mathcal{U}_t^{\mathrm{small}}$ define $X \leq Y$ if and only if $X \leq_{\partial m} Y$ and $X \sim Y$. This is a wqo since there are only finitely many equivalence classes of $\sim$ on $\mathcal{U}_t^{\mathrm{small}}$.

(3) *The Stop Signal.*

The graphs of $\mathcal{U}_t^{\mathrm{small}}$ are generated by size, one generation at a time (where the jth generation consists of all those of size j). We say that there is *nothing new at time j* if none of the t-boundaried graphs of the jth generation are minimal with respect to the search order $\leq$.

A *stop signal is detected at time* $2j$ if there is nothing new at time i for $i = j, ..., 2j-1$.

We have now completely described the algorithm. For $t = 1, ..., B+1$ we generate the t-boundaried graphs in the manner described until a stop signal is detected. We form the set $\mathcal{M}$ and output the list of elements of $\mathcal{M}$ (with boundaries forgotten) that are obstructions for $\mathcal{F}$. Note that having a decision algorithm for $\mathcal{F}$ is sufficient to determine if any particular graph H is an obstruction, just by checking that $H \notin \mathcal{F}$ while each minor of H is in $\mathcal{F}$. This same procedure and the decision algorithm for $\sim$ allow us to compute whether it is time to stop.

The correctness of the algorithm is established by the following claims.

Claim 1. For each value of t a stop signal is eventually detected.

This follows immediately from the fact that $\leq$ is a wqo on $\mathcal{U}_t^{\mathrm{small}}$ and therefore there are only a finite number of minimal elements.

Claim 2. Suppose that for a given t a stop signal is detected at time $2j$. Then no obstruction for $\mathcal{F}$ that can be parsed with the t-boundaried set of operators has size greater than $2j$.

If T is rooted tree, then by a *rooted subtree* T' of T we mean a subtree that is generated by some vertex r of T (the root of T'), together with all of the vertices descended from r in T. For

t-boundaried graphs X and Y, we say that X is a *prefix* of Y if, in a parse tree T for Y, X is parsed by a rooted subtree T' of T. To denote that X is a prefix of Y we write $X \prec Y$.

Now suppose that T is a parse tree of minimum size for a counterexample H to Claim 2. Since all of the operators in the standard set are either binary or unary, there must be a prefix H' of H of size at least j. Since there is nothing new during the times when H' would have been generated, Claim 2 follows from:

Claim 3. A prefix of a graph that is minimal with respect to $\leq$ must also be minimal.

If X is a prefix of Y and X is not minimal then $X \geq_{\partial m} X'$ with $X \neq X'$ and $X \sim X'$. Since $\sim$ is a right congruence $Y \sim Y'$ where Y' is obtained from Y by substituting a parse tree for X' for the subtree that parses X in a parse tree for Y. Since X' is a proper boundaried minor of X, Y' is a proper boundaried minor of Y. This implies that Y is not minimal with respect to $\leq$.

Claim 4. If $X \in \mathcal{O}$ then for some $t \leq B+1$, $X \in \mathcal{M}_t$.

Since the treewidth of X is at most B, $X \in \mathcal{U}^t_{\text{tree}}$ for some $t \leq B+1$. It remains to argue that X is $\leq$ minimal. But this is obvious, since any proper minor is in $\mathcal{F}$ and since $\sim$ refines the canonical $\mathcal{F}$-congruence. □

A pathwidth version of Theorem A can be proved in essentially the same way.

Jens Lagergren has shown that the use of the GMT in proving that the algorithm of Theorem A terminates can be replaced by an explicit calculation of a "stopping time" computable from the index of the congruence $\sim$ [Lag93].

Perhaps surprisingly, Theorem A can be implemented and a number of previously unknown obstruction sets have been mechanically computed [CD94, CDF95]. The "Holy Grail" of such efforts would be a computation of the obstruction set for torus embedding, which probably contains about 2,000 graphs.

Theorem A can also be adapted to other partial orders, including those such as the topological order, that are not a wqo. It can be shown in this case that the (adapted) algorithm will correctly terminate if and only if the ideal $\mathcal{F}$ has a finite obstruction set — thus providing a potentially interesting way to mechanically prove the *existence* of a finite basis for particular ideals in non-wqos.

4 Computational Engines That Stop on Width

In this section, we extend the basic ideas of Theorem A in a couple of different ways.

Let $\mathcal{O}^t$ denote the $\mathcal{F}$ obstructions of pathwidth at most t.

Definition. An *alarm* for a lower ideal $\mathcal{F}$ is a pair of computable functions:
(1) $f_\alpha : N \to N$, and
(2) $\alpha : N \to \{0,1\}$, satisfying:
(a) (*reliability*) $\alpha(t) = 1$ if there is an obstruction $H \in \mathcal{O} - \mathcal{O}^t$ of pathwidth more than $f_\alpha(t)$
(b) (*eventual quiescence*) $\exists t_0$ such that $\forall t \geq t_0$, $\alpha(t) = 0$.

Theorem B. Suppose the following are known for a minor order lower ideal $\mathcal{F}$:
(1) A decision algorithm for membership in $\mathcal{F}$.
(2) A decision algorithm for a finite-index congruence for $\mathcal{F}$. (The congruence can be either large or small.)
(3) Algorithms for computing α and f_α for an alarm for $\mathcal{F}$.
Then the obstruction set $\mathcal{O}$ for $\mathcal{F}$ can be computed.

Proof. For any fixed t, $\mathcal{O}^t$ can be computed using subroutines (1) and (2) by the methods of Theorem A adapted to pathwidth computations.

Define the *t-Stop Signal* to be that:

$$\mathcal{O}^i = \mathcal{O}^t \qquad \text{for } i = t, ..., f_\alpha(t)$$

and

$$\alpha(t) = 0$$

This forms the basis of the procedure that establishes the theorem.

Obstruction Set Computation

(1) $t \leftarrow 0$

(2) Repeat until a t-Stop Signal is detected:

$t \leftarrow t + 1$

Compute $\mathcal{O}^t$ and $\alpha(t)$.

Check for Stop Signals based on everything computed so far.

(3) Output $\mathcal{O} = \mathcal{O}^t$.

To see that this works correctly, it suffices to argue: (1) if $\mathcal{O} \neq \mathcal{O}^t$ then there will be no t-Stop Signal, and (2) eventually there will be a Stop Signal. If $\mathcal{O} \neq \mathcal{O}^t$ then we consider two cases: (i) There is an obstruction $H \in \mathcal{O} - \mathcal{O}^t$ with $\mathrm{pw}(H) \leq f_\alpha(t)$. In this case, the first condition for a t-Stop Signal will fail. (ii) There is an obstruction $H \in \mathcal{O} - \mathcal{O}^t$ with $\mathrm{pw}(H) > f_\alpha(t)$. In this case, the reliability of the alarm implies that $\alpha(t) = 1$ and so the second condition for a t-Stop Signal fails.

If t_1 is the maximum pathwidth of an $\mathcal{F}$ obstruction, then for all $t \geq t_1$, we have $\mathcal{O} = \mathcal{O}^t = \mathcal{O}^{t_1}$. Thus the first condition for a t-Stop Signal will be satisfied for all $t \geq t_1$. Let $t_2 = \max\{t_0, t_1\}$. The eventual quiescence of the alarm insures that the second condition for a t_2-Stop Signal will be met. □

We next prove an obstruction set computation algorithm that employs a terminating second-order congruence as the alarm.

Let B_h denote the complete binary tree of height h. Thus B_1 consists of a root and two children. B_h has $2^h - 1$ vertices, each vertex that is not a leaf has two children and each leaf is at distance h from the root. Let $h(t)$ be the least value of h such that $B_{h(t)}$ has pathwidth more than t, and let $f(t)$ be the number of vertices of $B_{h(t)}$. It can be shown that $f(t) = O(2^{2t})$. We will use the notation $f^{-1}(y)$ to denote the largest positive integer x such that $f(x) \leq y$. The following structural lemma is crucial to the approach. The proof has appeared in [CDF96].

Lemma 4.1 (Wide Factor Lemma) Let H be an arbitrary undirected graph, and let t be a positive integer. One of the following two statements must hold:

(a) The pathwidth of H is at most $f(t) - 1$.

(b) H can be factored: $H = A \oplus B$, where A, B are boundaried graphs with boundary size $f(t)$, the pathwidth of A is greater than t, and $A \in \mathcal{U}_{\text{path}}^{f(t)}$.

Furthermore, if $f(t+1) > t' > f(t)$, then one of the following must hold:

(c) The pathwidth of H is at most $t' - 1$.

(d) $H = A \oplus B$, $A \in \mathcal{U}_{\text{path}}^{t'}$, $B \in \mathcal{U}_{\text{large}}^{t'}$, and $pw(A) > t$. □

Proof Sketch. We suppose that we have a set of $2^{h(t)} - 1$ tokens corresponding to the vertices of $B_{h(t)}$. By a procedure for pebbling the graph with these tokens, we can either: (1) completely pebble the graph, in which case the sets of vertices occupied by pebbles at times $t = 0, 1, 2, ...$ yields a path-decomposition of width at most $2^{h(t)} - 2$, or (2) we get stuck (by running out of pebbles). In this case, at the stuck point, all of the pebbles are on the graph, and are linked in such a way that they provide a proof that the graph contains $B_{h(t)}$ topologically. We remark that the proof of this Lemma (which we use here only structurally) is significant for providing the first simple linear-time algorithm for obtaining an approximate path-decomposition of a graph. □

We remark that the Wide Factor Lemma appears to be a bit "thin" in the sense that a "best possible" lower bound on the pathwidth of the t-factor should probably be closer to $t/4$ than $\log t$. No analog for treewidth is currently known.

The Wide Factor Lemma is part of our method of recursively detecting large counterexamples to the hypothesis: $\mathcal{O} = \mathcal{O}^t$. The form proved above is the most natural, in some sense, since it is allied with an efficient approximate path decomposition algorithm. The factor A that it produces has the weakness, however, that *all* of the vertices of A may be boundary vertices. We next prove a form that is probably better suited to establishing termination properties of second-order congruences. We give this variation a similar name.

Lemma 4.2 (Fat Factor Lemma) There is a (known) recursive function $f(t) = O(2^{2^t})$ such that if H is an arbitrary undirected graph then one of the following three statements must hold:
(1) The pathwidth of H is at most $f(t) - 1$.
(2) H can be factored: $H = A \oplus B$, where A and B are $f(t)$-boundaried graphs, the pathwidth of A is greater than t, A has at least t internal vertices, and $A \in \mathcal{U}_{\text{path}}^{f(t)}$.
(3) H topologically contains the complete graph on t vertices.

Proof. We make use of a theorem of Mader [Mad72] that constructively identifies a function $g(t) = O(2^t)$ such that any graph with minimum degree $g(t)$ contains topologically the complete graph on t vertices. We use the same proof technique as for the Wide Factor Lemma, except that we preface the pebbling procedure of that proof with an attempt at the following pebbling moves that require at most $t \cdot g(t) + 1$ additional pebbles:

Repeat t times:

If there is a vertex v of H of degree at most $g(t)$ (possibly pebbled) then:
Pebble $N[v]$.
Remove the pebble from v.

If we are unable to complete this preface, then H has a subgraph of minimum degree $g(t)$, and therefore H topologically contains the complete graph K_t. If we complete the preface, then the argument for the Wide Factor Lemma shows that one of the other two alternatives must hold. □

Theorem C. Suppose the following are known for a minor order lower ideal $\mathcal{F}$:
(1) A decision algorithm for membership in $\mathcal{F}$.
(2) A decision algorithm for a terminating second-order congruence $\approx$ for $\mathcal{F}$.
Then the obstruction set $\mathcal{O}$ for $\mathcal{F}$ can be computed by an algorithm that uses (1) and (2) as subroutines.

Proof. We may assume that $\mathcal{F}$ is nontrivial (i.e., has at least one obstruction) because this can be easily determined using the subroutine for (2) with $t = 2$. (Note that the algorithm (2) allows us to decide a large (first-order) congruence for $\mathcal{F}$ for elements of $\mathcal{U}_{\text{large}}^t$ by considering singleton sets, and that this congruence necessarily refines the canonical first-order congruence for $\mathcal{F}$.) Let $\mathcal{O}^t$ denote the $\mathcal{F}$ obstructions of pathwidth at most t. Let $\mathcal{M}^t$ denote the minimal elements of $\mathcal{U}^t_{\text{path}}$ in the partial order that is the intersection of the large (first-order) congruence available from (2) and the boundaried minor order. Since the congruence has finite index, $\mathcal{M}^t$ is finite by the GMT. For any fixed t, the sets $\mathcal{O}^t$ and $\mathcal{M}^t$ can be computed using subroutines (1) and (2) by the methods of Theorem A adapted to pathwidth computations.

Let $m(t)$ denote the maximum order of an obstruction of pathwidth at most t. Let f be the function in the Fat Factor Lemma. Let $t' > f(m(t))$. We say that the *t-Stop Signal* is *witnessed at* t' if: (1) $\mathcal{O}^i = \mathcal{O}^t \quad$ for $i = t, ..., t'$, and
(2) $\forall A \in \mathcal{M}^{t'}$ with $\text{pw}(A) > f^{-1}(t') \geq m(t)$ and $|\text{int}(A)| > f^{-1}(t') \geq m(t):\quad \{A\} \approx S(A)$.
A t-Stop Signal *occurs* if there is a $t' > t$ as above at which it is witnessed.

Obstruction Set Computation

(1) $t \leftarrow 0$
(2) Repeat until a t-Stop Signal occurs:
$t \leftarrow t + 1$
Compute $\mathcal{O}^t$ and M^t.
Check for Stop Signals based on everything computed so far.
(3) Output $\mathcal{O} = \mathcal{O}^t$.

We argue that if $\mathcal{O} \neq \mathcal{O}^t$ then there will be no t-Stop Signal. Let $H \in \mathcal{O} - \mathcal{O}^t$, and suppose that the t-Stop Signal is witnessed at t'. If $\text{pw}(H) \leq t'$ then clearly there will be no Stop Signal. So suppose $\text{pw}(H) > t'$ where $t' > f(m(t))$. By Lemma 4.2, one of two cases must hold:
Case 1: There is a factorization $H = A \oplus B$ where $A \in \mathcal{U}^{t'}_{\text{path}}$, $B \in \mathcal{U}^{t'}_{\text{large}}$, $\text{pw}(A) > m(t)$ and $|\text{int}(A)| > m(t)$. Since A is a factor of an obstruction, we have that for every A' properly below A in the boundaried minor order, $A' \oplus B \in \mathcal{F}$ and therefore $A \not\approx_{\mathcal{F}} A'$. Consequently for each such A' we have $A \not\approx A'$ and thus $A \in \mathcal{M}^{t'}$. We also have that $\{A\} \not\approx_{\mathcal{F}} S(A)$ and therefore $\{A\} \not\approx S(A)$, a contradiction.
Case 2: A topologically contains the complete graph on $m(t)$ vertices. But in this case A topologically contains any obstruction in $\mathcal{O}^t$, which contradicts that it is a factor of an obstruction in $\mathcal{O} - \mathcal{O}^t$.

If $\mathcal{O} = \mathcal{O}^t$ (this must eventually hold, since $\mathcal{O}$ is finite), then a t-Stop Signal will be witnessed at $t' = \max\{f(t_0), f(m(t))\}$, where t_0 is the termination constant. $\square$

Some Remarks on Implementations

A proof that an obstruction set can be computed that uses Theorem C (or Theorem B) leaves us in an interesting situation. For a concrete example, suppose we believe (and are correct) that all of the obstructions have been found for $t = 4$. We know by Theorem B that if we are wrong, either we will find a new obstruction at $t' = 5$, or a factor of a large obstruction will cause the second part of the $t = 4$ Stop Signal not to occur at $t' = 5$. However, there is no converse implication for the second part of the Stop Signal; it may fail at $t' = 5$ even if $\mathcal{O} = \mathcal{O}^4$. All that we are guaranteed is that there exists a t' at which the $t = 4$ Stop Signal will be witnessed. Whether in practice much "waiting" is required for a particular obstruction set computation is an interesting question (which of course will depend on the particular congruence employed).

Furthermore, suppose that the procedure of Theorem C is applied using a second-order congruence for $\mathcal{F}$ that is *not known* to be terminating. If a t-Stop Signal is observed at at some t' then the proof of the theorem shows that $\mathcal{O} = \mathcal{O}^t$. For implementations, this is likely to be a valuable *sufficient* condition for obstruction set identification. We conjecture that many "natural" second-order congruences terminate rapidly, while our ability to prove termination appears to be much weaker. A natural second-order congruence for the union of ideals is described in §6.

The implementation of an obstruction set computation engine at the University of Victoria and Los Alamos National Laboratories (described in [CD94, CDF95]) is based, for a fixed pathwidth bound t, on the exploration of a tree whose root is the empty t-boundaried graph, and whose nodes correspond to the elements of $\mathcal{M}^t$, the minimal elements with respect to a known (first-order) congruence $\sim$ for $\mathcal{F}$. An element of $\mathcal{M}^t$ is characterized by the property: $\forall A' \in S(A) : A \not\sim A'$. In the proof of Theorem C we use the stronger property satisfied by a factor A of an obstruction relative to a second-order congruence $\approx$ for $\mathcal{F}$: $A \not\approx S(A)$. This can provide the basis for an improved search strategy that explores only the subtree generated by that subset of $\mathcal{M}^t$ that satisfies this more stringent minimality criterion.

5 The Canonical Second-Order Congruence

In this section, we show that Theorem C is natural, by establishing that the canonical second-order congruence for an ideal necessarily terminates.

Lemma 5.1 Let $G = (V, E)$ be an ordinary graph. If $\mathrm{pw}(G) = w$ then any subdivision of G has pathwidth at most $w + 2$.

Proof. Let $H =$ be a subdivision of G. Thus for each edge uv of G we have a (possibly empty) set of vertices of H that subdivide uv. Let S_{uv} denote this set of vertices and suppose that the vertices of S_{uv} are indexed in the order in which they occur between between u and v, starting from either end (this is not important)

$$S_{uv} = \{s[u, v, i] : 1 \leq i \leq m_{uv}\}$$

Let $(P_1, ..., P_m)$ be a path decomposition for G of width w. Thus each set of vertices P_i has at most $w-1$ members. For each edge uv of G choose a set $P_{i_{uv}}$ of the decomposition such that $\{u, v\} \subseteq P_{i_{uv}}$. We may assume that the choices are all distinct, just by assuming that any set of the decomposition of G is repeated sufficiently many times. We can obtain a path decomposition of H of the width required by replacing each set $P_{i_{uv}}$ in the decomposition of G by the sequence of sets:

$$(P_{i_{uv}},\quad P_{i_{uv}} \cup \{s[u,v,1]\},\quad P_{i_{uv}} \cup \{s[u,v,1], s[u,v,2]\},\quad P_{i_{uv}} \cup \{s[u,v,2], s[u,v,3]\},\quad \cdots$$
$$P_{i_{uv}} \cup \{s[u,v,m_{uv}-1], s[u,v,m_{uv}]\},\quad P_{i_{uv}} \cup \{s[u,v,m_{uv}]\},\quad P_{i_{uv}})$$

That this satisfies the definition of a path decomposition for H is easily checked. □

Definition. A t-boundaried graph $A \in \mathcal{U}^t_{\text{large}}$ is a *minimal topological factor* of an ordinary graph H with respect to a fixed t-boundaried graph $B \in \mathcal{U}^t_{\text{large}}$ if:
(1) $A \oplus B \geq_{\text{top}} H$, and
(2) For every A' properly below A in the boundaried topological order, $A \oplus B$ is not above H in the topological order.

Lemma 5.2 Suppose A, B and B are t-boundaried graphs with A a minimal topological factor of H with respect to B and B a minimal topological factor of H with respect to A. Then $A \oplus B$ is a subdivision of H where the only subdivisions are boundary vertices. □

Definition. Given an ordinary graph H, define $\text{parts}^t(H)$ to be the set of all t-boundaried graphs $A \in \mathcal{U}^t_{\text{large}}$ for which there is a t-boundaried graph $B \in \mathcal{U}^t_{\text{large}}$ such that A is a minimal topological factor of H with respect to B.

Lemma 5.3 $\mathrm{parts}^t(H)$ is computable from H.

Proof. Let $H^{(t)}$ denote the graph obtained from H by subdividing each edge t times, or equivalently, replacing each edge with a path having t internal vertices. Let $P(H)$ denote the set of t-boundaried graphs obtained from $H^{(t)}$ by the following procedure:
(1) In all possible ways: specify a boundary set V' of size t.
(2) In all possible ways: partition the edge components of $H - V'$ into two sets and thus obtain two factors A and B of $H^{(t)}$.
(3) For each such factorization $A \oplus B = H^{(t)}$ compute the minimal factors of H with respect to B that are below A in the boundaried topological order.
(4) Let $P(H)$ be the union of the sets computed in Step (3).

Clearly $P(H) \subseteq \mathrm{parts}^t(H)$ by definition. To prove the inclusion in the other direction, suppose A is a minimal topological factor of H with respect to B, where $A, B \in \mathcal{U}^t_{\mathrm{large}}$. Let B' be a minimal topological factor of H with respect to A. From the definition, A is also a minimal topological factor of H with respect to B'. By Lemma 5.3, $A \oplus B'$ is a subdivision of H in which the only subdividing vertices are boundary vertices. The lemma follows. □

The next lemma follows easily from the definitions.

Lemma 5.4 Let X and Y be t-boundaried graphs. Then $X \oplus Y \geq_{\mathrm{top}} H$ if and only if there are $X', Y' \in \mathrm{parts}^t(H)$ for which: $X \geq_{\mathrm{top}} X'$, $Y \geq_{\mathrm{top}} Y'$, $X \oplus Y' \geq_{\mathrm{top}} H$, $X' \oplus Y \geq_{\mathrm{top}} H$, and $X' \oplus Y' \geq_{\mathrm{top}} H$. □

Theorem D. The canonical second-order congruence $\approx_{\mathcal{F}}$ for a lower ideal $\mathcal{F}$ terminates.

Proof. By the GMT, $\mathcal{F}$ has a finite set of obstructions in the minor order, and therefore also a finite set of obstructions $\mathcal{O} = \{H_1, ..., H_s\}$ in the topological order. Let m be the maximum pathwidth of the H_i. Take $t_0 = m + 3$. Suppose $A \in \mathcal{U}^t_{\mathrm{small}}$ for $t \geq t_0$, and suppose $\mathrm{pw}(A) \geq t_0$. If $\{A\} \not\approx_{\mathcal{F}} S(A)$ then there is a t-boundaried graph $B \in \mathcal{U}^t_{\mathrm{large}}$ such that $A \oplus B \notin \mathcal{F}$ but for every $A' \in S(A)$ we have $A' \oplus B \in \mathcal{F}$. Thus $A \oplus B \geq_{\mathrm{top}} H_i$ for some i, $1 \leq i \leq m$. Suppose that B is a minimal element in the t-boundaried topological order on $\mathcal{U}^t_{\mathrm{large}}$ for which this is so. By Lemma 5.4, it must be the case that $A \in \mathrm{parts}^t(H_i)$. By Lemma 5.2, $A \oplus B$ is a subdivision of H_i. By Lemma 5.1, we have $\mathrm{pw}(A \oplus B) \leq m + 2$. But this contradicts that $\mathrm{pw}(A) \geq m + 3$. □

Note that the property of termination for a second-order congruence is a *finiteness property*, and thus is amenable to powerful tools such as the GMT. The GMT is used implicitly in the above proof of termination.

6 A Second-Order Congruence for the Union Problem

A consequence of Theorem A is that the obstructions for a union $\mathcal{F} = \mathcal{F}_1 \cup \mathcal{F}_2$ of lower ideals is computable from $\mathcal{O}_1$ and $\mathcal{O}_2$ if $\mathcal{O}_1$ and $\mathcal{O}_2$ each contain a planar graph, since in this case it is possible to effectively calculate a bound on the maximum treewidth of an $\mathcal{F}$-obstruction, and since knowledge of $\mathcal{O}_1$ and $\mathcal{O}_2$ allow us to compute a refinement of the canonical recognizability congruence for $\mathcal{F}$ (see Lemma 6.1 below). Alternatively, using the results of Lagergren and Arnborg [LA91], we have enough information to compute a bound on the maximum number of vertices in an $\mathcal{F}$-obstruction, and can then compute $\mathcal{O}$ by exhaustive search. If we are only interested in planar graphs, then such a bound on the maximum size of an obstruction can also be computed by the different method of Gupta and Impagliazzo [GI91] which may give a better bound.

In this section, we show how to use Theorem C to compute the obstruction set for a union

$\mathcal{F} = \mathcal{F}_1 \cup \mathcal{F}_2$ of ideals with known obstruction sets $\mathcal{O}_1$ and $\mathcal{O}_2$, in the case where one of these contains a tree.

Clearly, knowing $\mathcal{O}_1$ and $\mathcal{O}_2$ allows us to decide membership in $\mathcal{F}$, and thus we have the first ingredient for applying Theorem C. We next describe a decision algorithm for a second order congruence for $\mathcal{F}$ based on a set of abstract tests.

For each positive integer t, the set of predicates (abstract tests) is indexed and defined as follows, where X is the t-boundaried graph to which the predicate is applied:

Index: (B_1, B_2) where $B_1 \in \text{parts}^t(H_1)$, $B_2 \in \text{parts}^t(H_2)$, $H_1 \in \mathcal{O}_1$ and $H_2 \in \mathcal{O}_2$.
Question: Is there a choice of $i \in \{1, 2\}$ such that $X \oplus B_i \in \mathcal{F}$?

We say that X *fails* the test $\tau = (B_1, B_2)$ if the answer to the question is "no".

Definition. If $\mathcal{T} = \{\mathcal{T}^t\}$ is a collection of sets of tests on t-boundaried graphs, we define the second-order congruence *induced by* $\mathcal{T}$ by $S_1 \approx S_2$ if and only if

$$\forall \tau \in \mathcal{T}^t : (\exists A_1 \in S_1 : \; A_1 \text{ fails } \tau) \Longleftrightarrow (\exists A_2 \in S_2 : \; A_2 \text{ fails } \tau)$$

Note that if the sets of tests are finite, then the induced congruence has finite index on subsets of $\mathcal{U}^t_{\text{large}}$.

Lemma 6.1 The second-order congruence $\approx$ induced by the set of tests described above is a refinement of the canonical second-order congruence for $\mathcal{F} = \mathcal{F}_1 \cup \mathcal{F}_2$.

Proof. Suppose S_1 and S_2 are sets of t-boundaried graphs with $S_1 \not\approx_{\mathcal{F}} S_2$. Then (w.l.o.g.) $\exists Z \in \mathcal{U}^t_{\text{large}}$ and $\exists X \in S_1$ with $X \oplus Z \notin \mathcal{F}$ but $\forall Y \in S_2$, $Y \oplus Z \in \mathcal{F}$.

So we have $X \oplus Z \notin \mathcal{F}_1$ and $X \oplus Z \notin \mathcal{F}_2$. Let $Z \geq_{\text{top}} Z_1$ where Z_1 is minimal in the boundaried topological order, such that $X \oplus Z_1 \notin \mathcal{F}_1$ and similarly, suppose $X \oplus Z_2 \notin \mathcal{F}_2$ where $Z \geq_{\text{top}} Z_2$ and Z_2 is minimal. Then for some $H_1 \in \mathcal{O}_1$ and for some $H_2 \in \mathcal{O}_2$, we have $Z_i \in \text{parts}^t(H_i)$ for $i = 1, 2$. So $\tau = (Z_1, Z_2)$ is a test failed by X.

For all $Y \in S_2$, either $Y \oplus Z \in \mathcal{F}_1$ or $Y \oplus Z \in \mathcal{F}_2$, and therefore, either $Y \oplus Z_1 \in \mathcal{F}_1$ or $Y \oplus Z_2 \in \mathcal{F}_2$, so Y passes τ. Thus $S_1 \not\approx S_2$. □

We conjecture that the above congruence *always* terminates, but for now we have only the following weaker result.

Theorem E. Suppose that $\mathcal{F}$ is a union of ideals in the minor order

$$\mathcal{F} = \mathcal{F}_1 \cup \mathcal{F}_2$$

where the obstruction sets $\mathcal{O}_1$ and $\mathcal{O}_2$ for $\mathcal{F}_1$ and $\mathcal{F}_2$ are known and suppose that $\mathcal{O}_1$ contains at least one tree T. Then the obstruction set $\mathcal{O}$ for $\mathcal{F}$ can be effectively computed.

Proof. Let $\mathcal{O}'_i$ be the set of topological obstructions for $\mathcal{F}_i$, $i = 1, 2$. These are easily computed from the sets $\mathcal{O}_i$. Choose t_0 to be larger than the maximum number of vertices of any graph in $\mathcal{O}_1 \cup \mathcal{O}_2$, and large enough so that any graph of pathwidth greater than or equal to t_0 contains topologically a complete binary tree T_1 of sufficient size so that any forest T'_1 obtained from T_1 by contracting or deleting a single edge still has the obstruction tree T as a minor.

Now suppose $t \geq t_0$ and $A \in \mathcal{U}^t_{\text{path}}$ such that:
(1) $\text{pw}(A) \geq t_0$, and
(2) $|\text{int}(A)| \geq t_0$
and suppose that A fails the test $\tau = (B_1, B_2)$. Thus $A \oplus B_1 \notin \mathcal{F}_1$ and $A \oplus B_2 \notin \mathcal{F}_2$. Choose $H \in \mathcal{O}'_2$ so that $A \oplus B_2 \geq_{\text{top}} H$, and fix attention on:

- A subgraph S_1 of A that is a subdivision of T_1.
- A subgraph S_2 of $A \oplus B_2$ that is a subdivision of H.

The vertices of S_2 are of two different kinds: (i) those that correspond to vertices of H, and (ii) those that correspond to subdivisions of edges of H. Let $u \in \mathrm{int}(A)$ be a vertex in the interior of A that is not of the kind (i). If u has degree 0, then $A' = A - u$ fails τ and we are done. Otherwise, there is an edge uv in A. Let A' be obtained from A by contracting uv. We have $A' \geq_{\mathrm{m}} T$, so $A' \oplus B_1 \notin \mathcal{F}_1$, and we have $A' \oplus B_2 \geq_{\mathrm{top}} H$, so $A' \oplus B_2 \notin \mathcal{F}_2$. Thus $A' \in S(A)$ fails τ, which shows that $\approx$ terminates. □

As *intertwine* of two graphs G and H is a graph that contains both G and H topologically, and that is minimal for this in the topological ordering. As a corollary of Theorem E, we have the following concerning the computation of the (necessarily finite, by the GMT) set of intertwines of two graphs.

Corollary. The set of intertwines of an arbitrary graph G and a tree T can be effectively computed.

Proof. Let $\mathcal{O}_1$ be the set of graphs that are minimal in the minor order (equivalently, the topological order) on the universe U of graphs of maximum degree, among those graphs that have G as a minor. Let $\mathcal{O}_2$ similarly be the set of graphs that are minimal in the minor order (equivalently, the topological order) on U, among those graphs (i.e., trees) that have T as a minor. These sets can be computed by considering all possible ways of splitting vertices of degree greater than 3.

The procedures of Theorems A, B and C can be restricted to recursive subsets of the set of all graphs (in the manner of Consequence 2 of [FL88]). Using Theorem C restricted in this way to U, compute the U-intertwines of each pair of graphs (H_1, H_2) with $H_1 \in \mathcal{O}_1$ and $H_2 \in \mathcal{O}_2$, and let $\mathcal{O}$ denote the union of all of these sets of U-intertwines. It is easy to show that if H is an intertwine of G and T, then H is a minor of some $H' \in \mathcal{O}$. The set of intertwines of G and T can therefore be computed by searching exhaustively among the minors of the graphs in $\mathcal{O}$. □

7 Summary and Open Problems

To what extent can the Graph Minor Theorem be made effective? It seems to us that much further progress on this general question should be possible, in part because powerful results (such as the GMT itself), can be brought to bear on such questions. It is sometimes assumed that anything having to do with well-quasiordering is hopelessly impractical, but the successful implementation of obstruction set theorem-provers belie this and must be regarded as a notable development, given the important role of forbidden substructure theorems in graph theory.

The main contribution of this paper has been to establish methods for computing obstruction sets that do not require a prior bound on maximum obstruction width. We believe our results to be of practical significance for implementations of obstruction set theorem-provers. The following three basic questions stand out for attention.

(1) Is it possible to compute the obstruction set for a minor ideal $\mathcal{F}$ from an oracle for $\mathcal{F}$ membership and an oracle for the canonical recognizability congruence for $\mathcal{F}$?

(2) Is it possible to compute the obstruction set for a minor ideal $\mathcal{F}$ from a MSO description of $\mathcal{F}$?

(3) Is it possible to compute the obstructions for a union of ideals $\mathcal{F} = \mathcal{F}_1 \cup \mathcal{F}_2$ from the obstruction sets for $\mathcal{F}_1$ and $\mathcal{F}_2$?

By standard arguments, it is not hard to show that an answer of "yes" to question (i) implies "yes" to question $(i+1)$ for $i = 1, 2$. Our limited positive result on (3) could be extended to the case where one of the ideals excludes a planar graph if there is a positive resolution of the following question.

(4) Is there a treewidth analog of the Fat Factor Lemma?

Acknowledgement. We thank Bruno Courcelle and Jan Arne Telle for stimulating conversations and suggestions in connection with this work.

References

[AF93] K. R. Abrahamson and M. R. Fellows. Finite automata, bounded treewidth and well-quasiordering. In *Graph Structure Theory, Contemporary Mathematics vol. 147*, pp. 539–564. American Mathematical Society, 1993.

[APS90] S. Arnborg, A. Proskurowski and D. Seese. Monadic second-order logic: tree automata and forbidden minors. Technical Report, University of Oregon, Dept. of Computer and Information Sciences UO-CIS-TR-90/23, 1990.

[BFW92] H. P. Bodlaender, M. R. Fellows and T. J. Warnow. Two strikes against perfect phylogeny. *Proc. 19th International Colloquium on Automata, Languages and Programming (ICALP '92)*, Springer Verlag, Lecture Notes in Computer Science vol. 623 (1992), 273–283.

[CD94] K. Cattell and M. J. Dinneen. A characterization of graphs with vertex cover up to five. *Proceedings ORDAL'94*, Springer Verlag, Lecture Notes in Computer Science vol. 831 (1994), 86–99.

[CDF95] K. Cattell, M. J. Dinneen and M. R. Fellows. Obstructions to within a few vertices or edges of acyclic. *Proceedings WADS'95*, Springer-Verlag, Lecture Notes in Computer Science vol. 955 (1995), 415–427.

[CDF96] K. Cattell, M. J. Dinneen and M. R. Fellows. A simple linear time algorithm for finding path decompositions of small width. *Information Processing Letters* 57 (1996), 197–203.

[CL94] B. Courcelle and J. Lagergren. Equivalent definitions of recognizability for graphs of bounded tree-width. *Mathematical Structure in Computer Science*, to appear.

[Co90] B. Courcelle. The monadic second order logic of graphs I: Recognizable sets of finite graphs. *Information and Computation* 85 (1990), 12–75.

[DF94] R. Downey and M. Fellows. Parameterized computational feasibility. In *Proceedings of the Second Cornell Workshop on Feasible Mathematics*, (eds. P. Clote and J. Remmel), Birkhauser (1985), 219–244.

[FHW93] M. R. Fellows, M. T. Hallett and T. Wareham. DNA physical mapping: three ways difficult. *Proc. 1st European Symposium on Algorithms (ESA '93)*. Springer Verlag, Lecture Notes in Computer Science vol. 726 (1993), 157–168.

[FL88] M. R. Fellows and M. A. Langston. Nonconstructive tools for proving polynomial-time decidability. *J. ACM* 35 (1988), 727–739.

[FL89a] M. R. Fellows and M. A. Langston. On search, decision and the efficiency of polynomial-time algorithms. *Proc. Symposium on the Theory of Computing* (STOC), ACM Press (1989), 501–512.

[FL89b] M. R. Fellows and M. A. Langston. An analogue of the Myhill-Nerode theorem and its use in computing finite-basis characterizations. *Proc. Symposium on the Foundations of Computer Science* (FOCS), IEEE Press (1989), 520–525.

[GI91] A. Gupta and R. Impagliazzo. Computing planar intertwines. *Proc. 1991 Symposium on the Foundations of Computer Science (FOCS '91)*, IEEE Press (1991).

[Kin94] N. G. Kinnersley. Constructive obstruction set isolation for min cut linear arrangement. Technical Report, 1994, Computer Science Department, Univ. of Iowa, Ames, Iowa.

[KL91] N. G. Kinnersley and M. A. Langston. Obstruction set isolation for the gate matrix layout problem. *Discrete Appl. Math.* 54 (1994), 169–213.

[LA91] J. Lagergren and S. Arnborg. Finding minimal forbidden minors using a finite congruence. *Proc. 18th International Colloquium on Automata, Languages and Programming* (ICALP), Springer-Verlag, Lecture Notes in Computer Science vol. 510 (1991), 533–543.

[Lag93] J. Lagergren. An upper bound on the size of an obstruction. In *Graph Structure Theory, Contemporary Mathematics vol. 147*, pp. 601–622. American Mathematical Society, 1993.

[Mad72] W. Mader. Hinreichende dedingungen fur die existenz von teilgraphen, die zu einem vollstandigen graphen homomorph sind. *Abh. Math, Sem. Hamburg Univ.* 37 (1972), 86–97.

[NW63] C. St. J. A. Nash-Williams. On well-quasi-ordering finite trees. *Proc. Camb. Philos. Soc.* 59 (1963), 833–835.

[Pr93] A. Proskurowski. Graph reductions, and techniques for finding minimal forbidden minors. In *Graph Structure Theory, Contemporary Mathematics vol. 147*, pp. 591–600. American Mathematical Society, 1993.

[RS83] N. Robertson and P. D. Seymour. Graph minors I: Excluding a forest. *J. Comb. Theory Series B* 35 (1983), 39–61.

[RS85] N. Robertson and P. D. Seymour. Graphs minors — a survey. In *Surveys in Combinatorics*, I. Anderson, Ed. Cambridge University Press, 1985, pp. 153–171.

[RS94] N. Robertson and P. D. Seymour. Graph minors XVI: Wagner's conjecture. *J. Comb. Th. Ser. B*, to appear.