DIMENSION AND STRUCTURE FOR A POSET OF GRAPH MINORS

NOAH STREIB AND STEPHEN J. YOUNG

Abstract. Given a graph G with labeled vertices, define $\mathrm{MP}(G)$, the *labeled minor-poset* of G, to be the poset whose elements are the minors of G with $G_1 \leq G_2$ if and only if $G_1 \preceq G_2$. In this paper we study the dimension this poset, which is a minor-monotone graph parameter. We provide important structural results which yield non-trivial bounds on this parameter for cycles, complete graphs, trees, and the minor-closed class of graphs which exclude $K_{2,4}$-minors. We also state a conjecture that characterizes of the class of $K_{2,t}$-minor free graphs. Lastly, we consider two multigraph models and provide direction for future research.

1. Introduction

Over the past 30 years there have been many significant results that relate structural graph theory and the structure of partially ordered sets. From the primarily-graph-theory perspective, the most notable example is the acclaimed graph minors theorem of Robertson and Seymour [17], which has vaulted the study of a graph's minors to the central role in the study of structural graph theory. The theorem states that any poset $\mathbf{P} = (X, P)$, in which X is a possibly infinite set of finite, undirected, unlabeled graphs and $G \leq_P H$ if and only if G is a minor of H, has no infinite antichain. In the theory of partially ordered sets, structure is often related to the concept of dimension (e.g. see [20]). A collection of papers by Schnyder [18], Babbai and Duffus [3], Brightwell and Trotter [5] [6], and Felsner, Li and Trotter [11] examines this perspective by associating to each graph (or graph drawing) a poset which encapsulates the incidences of the graph. They show that the structure of the graphs is intimately related to the structure of the corresponding posets.

In this work we combine these viewpoints — to each graph G we associate a poset whose elements, the minors of G, are ordered by the graph-minor relation. As in the previous works, we study the relationship between the structural properties of graphs and that of the corresponding posets. The goal is to use tools from graph theory to prove dimension-related results abouts posets, and to use tools from the combinatorics of dimension to prove results about graphs. In this paper we discuss results that accomplish the former. However, we believe that deep results in graph theory may be attainable with the further development of this line of research.

Along with the exposition above, this work was also motivated by the following theorem.

Theorem 1. [21] *Let $G = (V, E)$ be a connected graph with $|V| \geq 2$. Form a poset $\mathbf{P} = (X, P)$ where X is the set of all connected induced subgraphs of G and $G_1 \leq_P G_2$ if and only if G_1 is an induced subgraph of G_2. Then $\dim(P)$ is the number of non-cutvertices in G.*

If such a statement is true for the poset of induced subgraphs, what can be said about the analogous poset of minors? We view our work as a natural extension of this idea. While no statement as simple as that of Theorem 1 can be made, our work has deeper structural significance.

The paper is organized as follows. In the next section we discuss previous work related to this problem and introduce relevant terminology. In Section 3 we introduce define our model and prove some basic results. In Section 4 we proof statements concerning cycles and complete graphs. In Section 5 we analyze posets corresponding to trees. In Section 6 we derive bounds on the dimension of posets corresponding to certain minor-closed classes of graphs and state our main conjecture. In Section 7 we discuss extensions of our work to multigraphs. And finally, in Section 8 we discuss future directions of this research and offer a conclusion.

Date: October 20, 2010.
1991 *Mathematics Subject Classification.* 06A07, 05C83.
Key words and phrases. poset; dimension; graph; labeled graph; graph minor; $K_{2,t}$-minor; excluded minors.

2. Background and Terminology

Given a poset $\mathbf{P} = (X, P)$, a set of linear extensions whose intersection is $\mathbf{P}$ is called a *realizer* of $\mathbf{P}$. That is, a set of linear extensions realizes $\mathbf{P}$ if for every pair of incomparable elements x and y, there exists a linear extensions L_1 and L_2 in the set with $x <_{L_1} y$ and $y <_{L_2} x$. The *dimension* of a poset $\mathbf{P} = (X, P)$, a concept introduced in [10], is the size of the smallest realizer of $\mathbf{P}$. The term "dimension" comes from the fact that a poset of dimension t can be embedded in $\mathbb{R}^t$, with elements of X and relations in P corresponding to t-dimensional vectors to coordinate-wise vector comparison, respectively, but cannot embed in $\mathbb{R}^s$ for any $s < t$.

A useful property of dimension, and one that we shall use throughout this paper, is that it is monotonic over subposets. That is, if $\mathbf{Q}$ is a subposet of $\mathbf{P}$, then $\dim(\mathbf{Q}) \leq \dim(\mathbf{P})$. To see this, simply observe that any realizer of $\mathbf{P}$, when restricted to the elements of $\mathbf{Q}$, is a realizer for $\mathbf{Q}$.

From a computational perspective, dimension is a very difficult parameter to compute. The decision problem "Given a poset $\mathbf{P}$ and a natural number k, is $\dim(\mathbf{P}) \leq k$?" is NP-Complete for $k \geq 3$ [23]. In fact, the same question is NP-Complete for the class of height two posets when $k \geq 4$. Determining the computational complexity for height two posets when $k = 3$ remains an open problem. Moreover, approximation is also hard; approximating the dimension of an n-element poset in polynomial time within a factor of $\sqrt{n}$ would imply that NP = ZPP [14]. (The complexity class ZPP contains problems for which there is a probabilistic Turing machine that runs in polynomial time and either returns the correct answer or says "Do Not Know." It is known that P $\subseteq$ ZPP, and many computer scientists believe that P = ZPP.)

2.1. Previous Models and Old Results.

Two models have received the bulk of the attention in studying the relationship between graphs and posets. The principal model is the so-called *incidence poset*, or the *vertex-edge poset*. This is a height two poset with maximal points corresponding to edges, minimal points corresponding to vertices, and a vertex v is less than an edge e if and only if v is an end of e. In [3] Babai and Duffus proved that when the incidence poset of a graph has dimension at most three, the graph is planar. But real interest in the incidence poset was sparked by a theorem of W. Schnyder when he proved the converse, and hence the following theorem.

Theorem 2. [18] *A graph is planar if and only if the dimension of its incidence poset is at most three.*

This work spurred further research, and a number of very nice results ensued. In particular the following two theorems stand out. Both results consider the *vertex-edge-face* poset, which is the height three poset defined analagously to the vertex-edge poset, with faces corresponding to maximal elements.

Theorem 3. [5] *If G be a 3-connected planar graph drawn in the plane without edge-crossings, then the dimension of the vertex-edge-face poset of G is four. Furthermore, if an element corresponding to either a vertex or a face is removed from the poset, then the dimension is three.*

Theorem 4. [6] *Let G be a planar multigraph drawn in the plane without crossings. Then the dimension of the vertex-edge-face poset of G is at most four.*

Closely related to the dimension of the incidence poset is the idea of the dimension of a graph. In fact, for a given graph these two parameters can differ by at most one. A result concerning the dimension of a graph follows. Recall that a graph is outerplanar if it can be embedded in the plane with all vertices incident to a single face.

Theorem 5. [12] *A graph is outerplanar if and only if the dimension of the graph is at most three. Furthermore, in the event that the dimension is three, it is possible to choose a realizer of size three such that two extensions are dual.*

Attempts have been made to extend this theory to surfaces of higher genus. However, evidence suggests that such extensions are impossible — it is known that the dimension of the incidence poset of a complete bipartite graph is at most four [2]. Hence, the incidence poset of the graph obtained from K_n by subdividing each edge exactly once, which is a subgraph of K_{2n}, has dimension at most four for all n. But K_n is a minor of this graph, and the dimension of the incidence poset of K_n goes to infinity as n goes to infinity [13].

More recently, a model called the *adjacency poset* has been examined. This is another height two poset, this time with the vertex set corresponding to both the minimal and maximal elements. The minimal

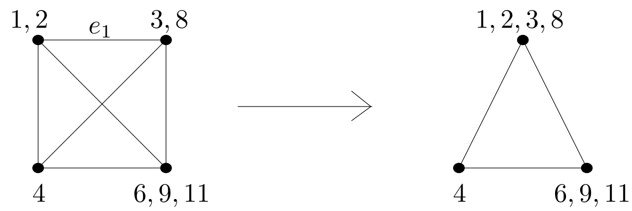

Figure 1: Contraction of the edge e_1.

element v' is less than maximal element u" if and only if $vu \in E(G)$. This poset has the nice property that dimension is related to chromatic number. As such, the Four Color Theorem can be used to prove the following theorem.

Theorem 6. [11] *If* $\mathbf{P}$ *is the adjacent poset of a planar graph* G*, then* $\dim(\mathbf{P}) \leq 8$*. If* G *is outerplanar, then* $\dim(\mathbf{P}) \leq 5$.

The best lower bounds for the statements in Theorem 6 are five and four, respectively. It would be of interest to improve the stated bounds or prove that no such improvement is possible. Efforts have been made to obtain results concerning higher genuses using adjacency posets. The primary result of this nature is the following theorem, whose proof is based on the existence of an upper bound on the acyclic chromatic number of graphs on a fixed surface.

Theorem 7. [11] *For every non-negative integer* g*, there exists an integer* d_g *so that if* $\mathbf{P} = (X, P)$ *is the adjacency poset of a graph* G *with genus* g*, then the dimension of* $\mathbf{P}$ *is at most* d_g.

2.2. **Terminology.** Let $G = (V, E)$ be a simple, undirected, unlabeled graph. The *neighborhood* of a vertex $v \in V(G)$, denoted $N(v)$, is the set of vertices adjacent to v. A *subgraph* of G is obtained by deleting edges and/or vertices (the removal of a vertex requires that all incident edges be removed as well). We will often abuse notation and write $G \backslash v$ and $G \backslash e$ to mean G with a vertex v deleted and G with the edge e deleted, respectively. Given an edge $e = \{u, v\} \in E(G)$, which will often be denoted e_{uv}, the graph G/e is obtained from G by *contracting* e; that is, we identify u and v and delete all resulting loops and parallel edges. A graph H is a *minor* of G, written $H \preceq G$, if H can be obtained by contracting edges in a subgraph of G. A minor of G is *proper* if it is not G itself.

In this paper we will be concerned with labeled graphs, so we need another definition of minor. Let G be a simple, undirected graph with vertex labels. Given $\{u, v\} \in E(G)$, the graph G/e is obtained by contracting e and labeling the new, identified vertex, with the union of the labels of u and v. For an example, see Figure 1. As in the previous case, a labeled graph H is a *labeled-minor* of G if H can be obtained by contracting edges in a subgraph of G. When it is clear that all graphs in question are labeled, we will simply refer to H as a minor of G and write $H \preceq G$. For convenience we will assume without loss of generality that the label function on G is the identity function. Thus, for any graph $H \preceq G$ we can denote by $\mathcal{L}_H \colon V(H) \longrightarrow 2^{V(G)}$ the label function on the vertices of H. Also, when it is clear from the context, we will write $\mathcal{L}_H(v)$ for some $v \in V(H)$ to mean the subgraph of G induced by the labels on v.

As we will be dealing with the relationship between various graph structures it will be helpful to define terminology and notation from graph theory. We first define several means of constructing new graphs from collections of other graphs. Given two graphs G_1 and G_2, the *k-clique sum* of G_1 and G_2, denoted $G_1 \oplus_k G_2$ is formed by identifying a clique of size k in G_1 with a clique of size k in G_2. Notice that $\oplus_0$ is precisely the *disjoint union of* G_1 *and* G_2, which we will denote $G_1 \cup G_2$. Given a subset of the vertices S, we denote by $G_1[S]$ the subgraph of G_1 induced by the vertices S.

For any two sets of vertices $U, V \subseteq V(G)$, we denote by $E(U, V)$ the set of edges $\{u, v\}$ with $u \in U$ and $v \in V$. We note that U and V do not need to be disjoint; in particular, for any set of vertices S, $E(S, S)$ is the set of edges in $G[S]$.

For any connected graph G, a set of edges whose deletion disconnects G is called an *edge cut* of G. If a single edge e is an edge cut, then e is called a *bridge*. (Similarly, a *vertex cut* is a set of vertices whose deletion disconnects G, and a single vertex that disconnects the graph is a *cutvertex*.) Observe that $E(U, V)$ is an edge cut for $G[U \cup V]$, as long as $G[U \cup V]$ is connected. We will also denote, for any set of vertices

S, the complement of S by $\overline{S}$. Notice, if $S \subset V(G)$ for any connected graph G, that $E\big(S, \overline{S}\big)$ is an edge cut of G.

We now turn to some important structural properties of graphs. A graph G is *2-connected* if it is connected, has at least three vertices, and does not have a cutvertex. Bridges (including their ends) and maximal 2-connected subgraphs of G are called *blocks*. Observe that two blocks have at most one vertex in common, vertices in more than one block are precisely the cutvertices, and the blocks partition $E(G)$ in the obvious way. For any unfamiliar terms from graph theory we refer the reader to [8].

We also require definitions from the theory of posets. We say that two posets $\mathbf{P} = (X, P)$ and $\mathbf{Q} = (Y, Q)$ are *isomorphic* if there exists a bijection $f : X \to Y$ such that $x_1 \leq_P x_2$ if and only if $f(x_1) \leq_Q f(x_2)$. An isomorphism from $\mathbf{P}$ to a subposet of $\mathbf{Q}$ is called in *embedding* of $\mathbf{P}$ in $\mathbf{Q}$. Now let $\mathbf{P} = (X, P)$ be a poset and let $x, y \in \mathrm{inc}(X, P)$, the set of incomparable pairs of $\mathbf{P}$. The ordered pair (x, y) is a *critical pair* in (X, P) if (1) $z <_P x$ implies that $z <_P y$, for all $z \in X \setminus \{x, y\}$, and (2) $y <_P w$ implies $x <_P w$, for all $w \in X \setminus \{x, y\}$. We say that a linear extention L *reverses* the critical pair (x, y) if $y <_L x$. With this in mind, we can state the following important fact.

Theorem 8. [15] *Let $\mathbf{P} = (X, P)$ be a poset and let $\mathcal{R}$ be a set of linear extensions of $\mathbf{P}$. Then $\mathcal{R}$ is a realizer of $\mathbf{P}$ if and only if for every critical pair (x, y) in (X, P) there is some linear extension $L \in \mathcal{R}$ such that $y < x$ in L.*

Another crucial definition is that of an alternating cycle. An *alternating cycle* in $\mathbf{P}$ is a sequence $\{(x_i, y_i) \mid 1 \leq i \leq k\}$ of ordered pairs from $\mathrm{inc}(X, P)$ with $y_i \leq x_{i+1}$ in P cyclically for all $i \in 1, 2, \ldots, k$. From this definition it is straightforward to verify the following remark.

Remark 1. Let $X = \{(x_1, y_1), (x_2, y_2), \ldots, (x_n, y_n)\}$ be a set of critical pairs. If either $x_1 = x_2 = \cdots = x_n$ or $y_1 = y_2 = \cdots = y_n$, then X does not contain an alternating cycle.

An alternating cycle is *strict* if $y_i \leq x_j$ implies $j = i + 1$. This leads us to the following theorem, which states that there is a linear extension of $\mathbf{P}$ that reverses each imcomparable pair in $S \subseteq \mathrm{inc}(X, P)$ if and only if S does not contain an alternating cycle, or, equivalently, a strict alternating cycle.

Theorem 9. [22] *Let $\mathbf{P} = (X, P)$ be a poset and $S \subseteq \mathrm{inc}(X, P)$. Then the following statements are equivalent.*

(1) *The transitive closure of $P \cup S$ is partial order on X.*
(2) *S does not contain an alternating cycle.*
(3) *S does not contain a strict alternating cycle.*

In referring to the poset structure, we will denote by $\mathrm{U}(x)$ as the set of elements y where $x < y$. Similarly, $\mathrm{D}(x)$ is the set of elements z where $z < x$. These sets are referred to as the *up-set* and *down-set*, respectively. We will also say that an element y is a *cover* of an element x if $x < y$ and there is no element z such that $x < z < y$. Lastly, we introduce the notation $\mathbf{2}^n$ to refer to the subset lattice on n elements; that is, the poset whose elements are all of the subsets of a k-element set, which are ordered by set-inclusion. It is well know that $\dim(\mathbf{2}^n) = n$. For any unfamiliar poset terminology we refer the reader to [20].

3. The Minor Poset

Motivated by the work cited in Section 2, we now introduce our model.

Let G be a graph with vertex labels and define $\mathrm{MP}(G) = (X, P)$, the *labeled-minor poset* of G, in the following natural way. The elements of X are the labeled-minors of G, and $G_2 <_P G_1$ if and only if G_2 is a labeled-minor of G_1. For example, the minor poset for a labeled path on three vertices, $\mathrm{MP}(P_3)$, is drawn in Figure 2. In an analogous way, we define the *unlabeled-minor poset* of a graph G and denote this as $\mathrm{MP}_u(G)$.

Notice that both $\mathrm{MP}(G)$ and $\mathrm{MP}_u(G)$ have the property that if G_2 is a labeled-minor of G_1, then $\mathrm{MP}(G_2)$ is a subposet of $\mathrm{MP}(G_1)$. Therefore, as dimension is monotonic over subposets, $\dim(\mathrm{MP}(G))$ is a minor-monotone graph invariant. This is precisely the property that failed to hold for incidence posets, and is what gives this model potential for greater success in terms of relating structural results from graph theory to this parameter. (For more on the importance of minor-monotone graph invariants, see [19].) Another important property, which is easy to check, is that if G_2 is obtained from G_1 by some sequence of operations σ, then any permutation of σ for which the operations remain valid also yields G_2.

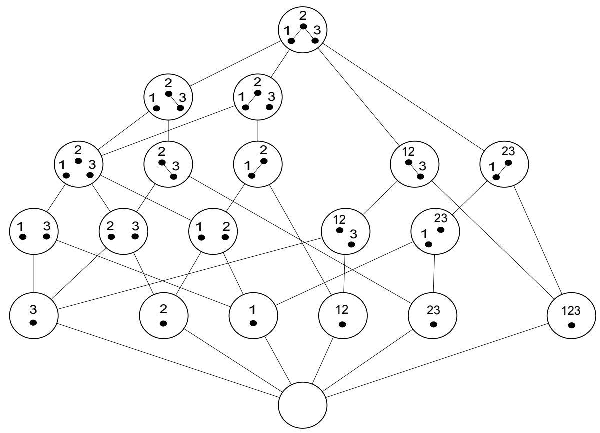

Figure 2: The labeled-minor poset for P_3.

3.1. **Labeled-Minor versus Unlabeled-Minor Posets.** In this paper we focus on $\mathrm{MP}(G)$ instead of $\mathrm{MP}_u(G)$. In this section we provide some justification for this choice.

The primary difference between the unlabeled-minor and labeled-minor posets is that the poset of labeled-minors seems to amplify small changes to the graph. For instance, consider the graph with no edges and k vertices. The unlabeled-minor poset for this graph is a k-element chain, which has dimension one, and the labeled-minor poset is isomorphic to $\mathbf{2}^k$, which has dimension k. In particular, any minor operation performed on this graph does not change the dimension of the unlabeled-minor poset, while the dimension of the labeled-minor poset will change under any minor operation. While there is no effective difference in the unlabeled-minor poset for this graph for any k, the labeled-minor poset amplifies the differences between graphs with different number of vertices.

Another way of viewing the difference between the unlabeled-minor and labeled-minor posets is that, in some sense, the labeled-minor poset has a full "memory" while the unlabeled-minor poset has only a partial memory. For instance, suppose $e = \{u, v\}$ is a cutedge in G and consider the minor poset of $G\backslash e$. In the unlabeled case, very little can be said about the relationship between the minor poset of $G\backslash e$ and the minor poset of G. However, in the labeled case, there is a clear delineation between the minors of $G\backslash e$ and G. Specifically, no minor of G that has a vertex label containing u and v can be a minor of $G\backslash e$.

We would like to prove that $\dim(\mathrm{MP}_u(G)) \leq \dim(\mathrm{MP}(G))$. Unfortunately this statement is false, as the following discussion illustrates.

Given that each element in the labeled-minor poset can be viewed as an element of the unlabeled-minor poset, once labels are removed, it seems reasonable to suspect that, for every labeled graph G, there is an embedding ϕ of $\mathrm{MP}_u(G)$ in $\mathrm{MP}(G)$. Further, we may suspect that this embedding would take a natural form, such as $\phi(G')$ is some labeling of the graph G' for every $G' \in \mathrm{MP}_u(G)$. However a quick comparison of the structures of the labeled-minor and unlabeled-minor posets for the cycle on four vertices, C_4, makes it clear such a embedding does not exist. For instance, consider the two chains $C_4 > K_3 > P_2 > K_2 \cup K_1$ and $C_4 > P_3 > P_1 \cup P_1 > K_2 \cup K_1 \cup K_1 > P_2 \cup K_1$, where $\cup$ denotes the disjoint union of graphs. The only way to create K_3 as a minor of C_4 is via contraction, and so $\phi(K_3)$ must have a vertex with two labels. Since the first chain has no vertex deletions, that means $\phi(K_2 \cup K_1)$ has a vertex with two labels. Now consider the second chain, where every operation is either an edge deletion or a vertex deletion. This implies that $\phi(K_2 \cup K_1)$ has no vertex with two labels. In particular, this means that there is no embedding of the unlabeled-minor poset in the labeled-minor poset that can be described via a natural labeling scheme. Nevertheless, as we can see from Figure 3, there is still an embedding of $\mathrm{MP}_u(C_4)$ in $\mathrm{MP}(C_4)$.

So, while no "natural" embedding exists, the question remains as to whether we can always find an embedding of $\mathrm{MP}_u(G)$ in $\mathrm{MP}(G)$, for every graph G. Unfortunately, there is no subposet of $\mathrm{MP}(C_5)$ that is isomorphic to $\mathrm{MP}_u(C_5)$. To see this, notice that the height of any graph G in $\mathrm{MP}(G)$ or $\mathrm{MP}_u(G)$ is $|V(G)| + |E(G)| + 1$. Therefore, any embedding of $\mathrm{MP}_u(C_5)$ in $\mathrm{MP}(C_5)$ maps elements of height h to other

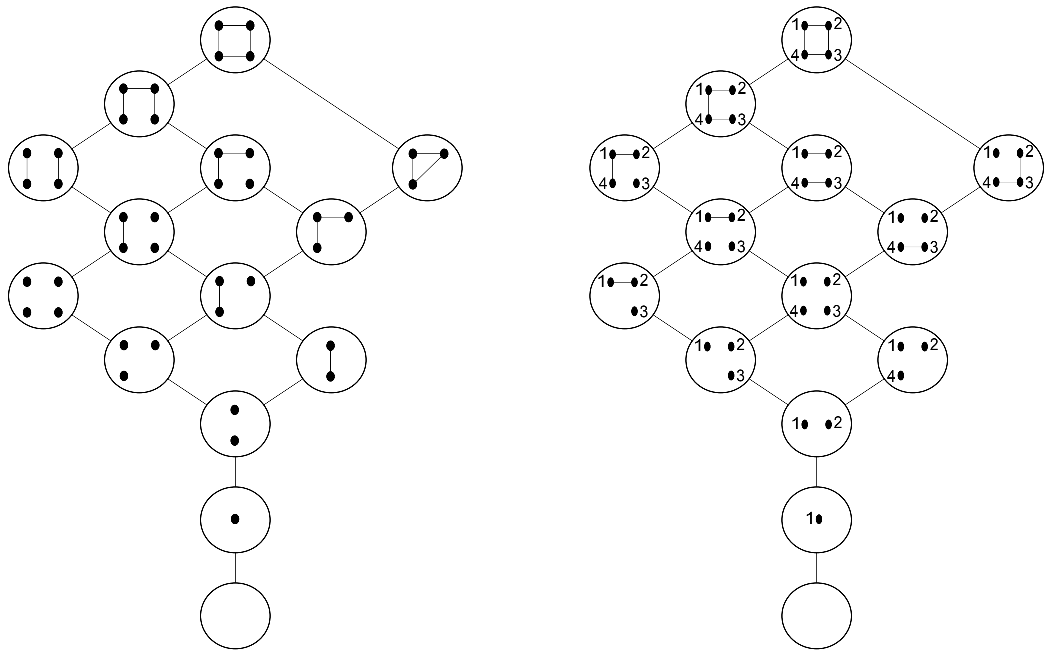

(a) $\mathrm{MP}_u(C_4)$ (b) Subposet of $\mathrm{MP}(C_4)$ isomorphic to $\mathrm{MP}_u(C_4)$

Figure 3: Minor posets for C_4

elements of height h. Consider the two elements of $\mathrm{MP}_u(C_5)$ of height eight that are below the element of height nine (see Figure 4). Neither of these can map to a labeled C_4, as C_4 is not below any element of height nine in $\mathrm{MP}(C_5)$. So they map to a labeled $P_3 \cup K_2$ or $P_4 \cup K_1$. But, given a fixed labeling of C_5, $\mathrm{MP}(C_5)$ does not have two elements of height eight, neither isomorphic C_4, that induce a $\mathbf{2}+\mathbf{2}$ with two elements of height seven. Thus, no such embedding exists. This is particularly surprising considering that $\mathrm{MP}_u(C_5)$ has 21 elements and $\mathrm{MP}(C_5)$ has 429.

It is worth noting that the labeled-minor poset can also be viewed as a forgetful version of the minor poset of a multigraph, where the contraction operation is modified to take into account multiple-edges, loops, and/or edge labels. Specifically, just as we may consider the unlabeled-minor poset as the labeled-minor poset where the vertex labels have been forgotten, we may consider the labeled-minor poset as the multigraph-minor poset where the number of edges are forgotten. Since we are focused on the topological consequences of the dimension of the minor poset, we will touch on these generalizations only briefly in Section 7.

3.2. **Preliminary Results.** According to Theorem 8, we can obtain a realizer of a poset by finding linear extensions that reverse all of the critical pairs. This idea is central to providing upper bounds on dimension for the minor poset, and leads us to the following lemma.

Lemma 10. *Let G be a graph. The ordered pair (G_1, G_2) is a critical pair of $\mathrm{MP}(G)$ if and only if G_1 and G_2 satisfy one of the following:*

- *G_1 is a single-vertex graph with label $\{v\}$ for some $v \in V(G)$, and G_2 is $G \backslash v$.*
- *G_1 is a single-vertex graph with label $U \cup V$, where $U \cap V = \varnothing$, $G[U]$ and $G[V]$ are connected, and $E(U,V) \neq \varnothing$. Further, $G_2 = G \backslash E(U,V)$.*
- *G_1 is a single-vertex graph with label Z and G_2 is G/e where $e \in E\big(Z, \overline{Z}\big)$.*

Proof. Let (G_1, G_2) be a critical pair in $\mathrm{MP}(G)$. Suppose G_1 has vertices $v_1, v_2, \ldots, v_t$ and $t \geq 2$. Since (G_1, G_2) is a critical pair, every proper minor, and in particular every proper subgraph, of G_1 is a minor of

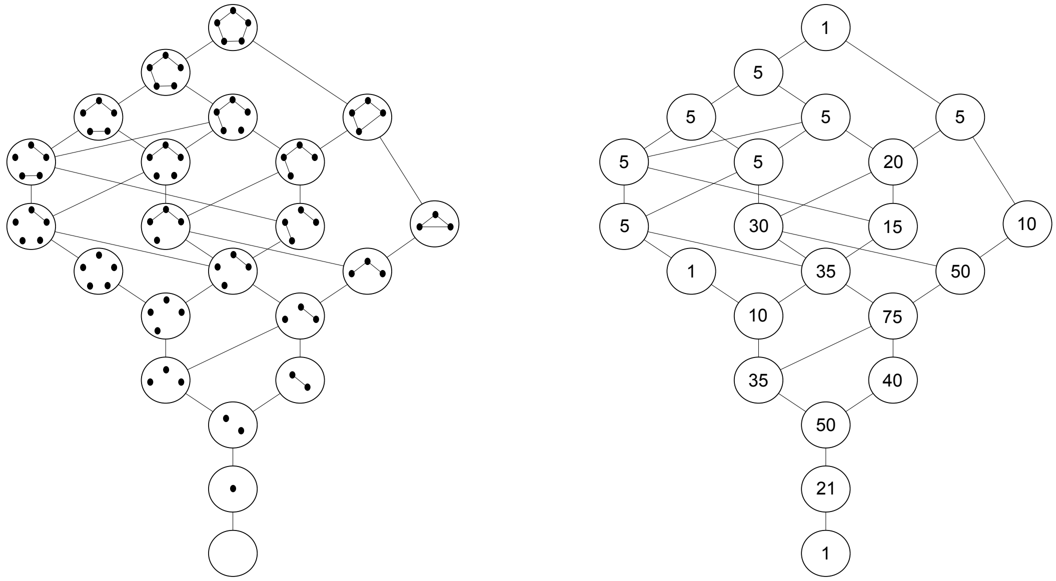

(a) $\mathrm{MP}_u(C_5)$ (b) Approximate structure of $\mathrm{MP}(C_5)$

Figure 4: Minor Posets for C_5. The numbers in 4b indicate the number of labeled-minors whose underlying graph is isomorphic to the corresponding graph in 4a.

G_2. Thus for every v_i there is a set $U_i \subseteq V(G_2)$ such that $G_2[U_i]$ is connected and $\mathcal{L}_{G_1}(v_i) = \cup_{u \in U_i} \mathcal{L}_{G_2}(u)$. Let G_2' be the minor of G_2 formed by deleting all the vertices in $\overline{\cup_{i=1}^{t} U_i}$ and contracting each set U_i to a single vertex. Note that there is a natural bijection from $V(G_1)$ to $V(G_2')$ by identifying vertices with the same set of labels. Since $G_1 \parallel G_2$, then $G_1 \parallel G_2'$ and in particular $G_1 \not\preceq G_2'$. Thus there is some edge $e_{v_i v_j} \in E(G_1)$ that is not present between the corresponding vertices in G_2'. More specifically, $E(U_i, U_j) = \varnothing$. However, by assumption $G_1/e_{v_i,v_j} \preceq G_2$ and thus $G_2[U_1 \cup U_2]$ is a connected, a contradiction.

So $t = 1$, G_1 has a single vertex v_1, and the only minor of G_1 is $\varnothing$. Now suppose $\mathcal{L}_{G_1}(v_1) = \{v\}$ for some $v \in V(G)$. Then, by the definition of critical pair, G_2 is either $G\backslash\{v\}$ or G/e for some edge adjacent to v in G. Thus (G_1, G_2) is either the first or third type of critical pair, respectively. So now assume $|\mathcal{L}_{G_1}(v_1)| \geq 2$ and define $U_{v_1} = \{u \in V(G_2) : \mathcal{L}_{G_1}(v_1) \cap \mathcal{L}_{G_2}(u) \neq \varnothing\}$. As $G_1 \parallel G_2$, either $G_2[U_{v_1}]$ is disconnected or $\mathcal{L}_{G_1}(v_1) \neq \cup_{u \in U_{v_1}} \mathcal{L}_{G_2}(u)$.

Suppose first that $G_2[U_{v_1}]$ is disconnected and $\mathcal{L}_{G_1}(v_1) = \cup_{u \in U_{v_1}} \mathcal{L}_{G_2}(u)$. Let (L_1, L_2) be the partition of $\mathcal{L}_{G_1}(v_1)$ induced by the cut (the defintion of critical pair ensures that this partition has exactly two parts). Since (G_1, G_2) is a critical pair, any graph G_2' such that $G_2 \preceq G_2'$ must be such that the set $\{u \in V(G_2') : \mathcal{L}_{G_1}(v_1) \cap \mathcal{L}_{G_2'}(u) \neq \varnothing\}$ induces a connected graph. Specifically, the minimal elements of the up-set of G_2 satisfy this and thus, if G_2' covers G_2, then $G_2 = G_2'\backslash\{e\}$ where the edge $e \in E_G(L_1, L_2)$. In particular, this implies that $G_2 = G\backslash E_G(L_1, L_2)$ and (G_1, G_2) is the second type of critical pair.

Now suppose that $\mathcal{L}_{G_1}(v_1) \neq \cup_{u \in U_{v_1}} \mathcal{L}_{G_2}(u)$. First consider the case where $\cup_{u \in U_{v_1}} \mathcal{L}_{G_2}(u) \subsetneq \mathcal{L}_{G_1}(v_1)$. Since (G_1, G_2) is a critical pair by assumption, every graph G_2' such that $G_2 \preceq G_2'$ must have $\mathcal{L}_{G_1}(v_1) \subseteq \cup_{u \in V(G_2')} \mathcal{L}_{G_2'}(u)$. However, since the only way to remove labels is via vertex deletion, this implies that $G_2 = G\backslash\{v\}$ for some $v \in \mathcal{L}_{G_1}(v_1)$. But, $G_2 = G\backslash\{v\} \preceq G\backslash E_G\left(\{v\}, \overline{\{v\}}\right)$ and, since $|\mathcal{L}_{G_1}(v_1)| \neq 1$, we have $G_1 \parallel G\backslash E_G\left(\{v\}, \overline{\{v\}}\right)$, contradicting the criticality of (G_1, G_2). Thus, we may assume that $\mathcal{L}_{G_1}(v_1) \subsetneq \cup_{u \in U_x} \mathcal{L}_{G_2}(u)$. Now the criticality of (G_1, G_2) implies that every graph G_2' such that $G_2 \preceq G_2'$ has a set of vertices U such that $\cup_{u \in U} \mathcal{L}_{G_2'}(u) = \mathcal{L}_{G_1}(v_1)$. However, the only way to combine labels is via contraction. This implies that there is some edge e in $E_G\left(\mathcal{L}_{G_1}(v_1), \overline{\mathcal{L}_{G_1}(v_1)}\right)$ such that $G_2 = G/e$. Thus (G_1, G_2) is the third type of critical pair. □

Perhaps the most interesting thing about Lemma 10 is that there is one type of critical pair corresponding to each minor operation; the first corresponding to vertex deletion, the second to edge deletion, and the third to edge contraction. For conciseness of notation, we will denote by $\mathcal{V}$ the critical pairs corresponding to vertex deletion and by $\mathcal{V}_v$ the particular critical pair corresponding to deleting the vertex v. Similarly, we will let $\mathcal{E}$ be the critical pairs corresponding to edge deletion and let $\mathcal{E}_{E(U,V)}$ denote of the critical pairs corresponding to the deletion of the edges $E(U, V)$. Finally, we will denote by $\mathcal{C}$ the critical pairs corresponding to contraction, with $\mathcal{C}_Z$ being those critical pairs where the label of the single-vertex graph is the set Z. Additionally, since many of these critical pairs deal with a single-vertex graph with a specified label, we write $\mathcal{G}_L$ to denote the single-vertex graph whose label is the set L. If $L = \{v\}$ we simply write $\mathcal{G}_v$.

Let S_k denote the *standard example* of size k; that is, the height two poset with minimal elements $a_1, a_2, \ldots, a_k$ and maximal elements $b_1, b_2, \ldots, b_k$ in which $a_i < b_j$ if and only if $i \neq j$. It is well known that $\dim(S_k) = k$. Given a poset $\mathbf{P}$, the size of the largest standard example that is a subposet of $\mathbf{P}$ is called the *breadth* of $\mathbf{P}$. Thus, $\dim(\mathbf{P})$ is lower-bounded by its breadth.

Proposition 11. *Let G be a graph with $|V(G)| = n$ and $|E(G)| = m$. The subposet induced by the critical pairs in $\mathcal{V}$ is isomorphic to S_n. The subposet induced by the critical pairs of type $\mathcal{E}_{\{\{u,v\}\}}$ for $\{u, v\} \in E(G)$ is isomorphic to S_m.*

Proof. We first observe that for every pair of distinct vertices u and v, $G\setminus\{v\} \parallel G\setminus\{u\}$. Furthermore, $\mathcal{G}_v \parallel \mathcal{G}_u$. Finally, $\mathcal{G}_v \preceq G\setminus\{u\}$ and $\mathcal{G}_u \preceq G\setminus\{v\}$, and thus the poset induced by $\mathcal{V}$ is isomorphic to S_n.

For any pair of distinct edges $e_1 = \{u_1, v_1\}$ and $e_2 = \{u_2, v_2\}$ we note that $G\setminus\{e_1\} \parallel G\setminus e_2$ and $\mathcal{G}_{\{u_1,v_1\}} \parallel \mathcal{G}_{\{u_3,v_2\}}$. Further, $\mathcal{G}_{\{u_1,v_1\}} \preceq G\setminus\{e_2\}$ and $\mathcal{G}_{\{u_2,v_2\}} \preceq G\setminus\{e_1\}$. Thus the critical pairs in $\mathcal{E}_{\{\{u,v\}\}}$ are isomorphic to S_m. □

Therefore the dimension of $\mathrm{MP}(G)$ is at least the maximum of $|V(G)|$ and $|E(G)|$. In the case of paths, we can construct a realizer whose size equals this maximum (see Theorem 22). Unfortunately, as we shall see, this lower bound is rarely tight.

Our next lemma is a decomposition result. Before we state it we require another definition concerning posets. Let $\mathbf{P} = (X, P)$ and $\mathbf{Q} = (Y, Q)$ be posets. The *Cartesian product*, $\mathbf{P} \times \mathbf{Q}$, is the poset on the elements $X \times Y$ where $(x_1, y_1) \leq (x_2, y_2)$ if and only if $x_1 \leq x_2$ and $y_1 \leq y_2$. Using this notation we have the following lemma.

Lemma 12. *Let G_1 and G_2 be graphs. The poset $\mathrm{MP}(G_1 \oplus_k G_2)$ is a subposet of $\mathrm{MP}(G_1) \times \mathrm{MP}(G_2)$ for $k \in \{0, 1\}$. Furthermore, if $k = 0$ then the two posets are isomorphic, whereas if $k = 1$ they are not.*

Proof. We will proceed first by defining a map ψ_k from $\mathrm{MP}(G_1 \oplus_k G_2)$ to the subgraphs of $G_1 \oplus_k G_2$. Specifically, given a $G \preceq G_1 \oplus_k G_2$, $\psi_k(G)$ will be a subgraph of $(G_1 \oplus_k G_2)[\cup_{v\in(G)}\mathcal{L}_G(v)]$, the subgraph induced by the labels of G, with only a selection of the edges. In order to determine which edges to keep we first fix an ordering σ of $E(G_1 \oplus_k G_2)$. Then for every $\{u, v\} \in E(G)$, keep the smallest edge in $E_{G_1\oplus_k G_2}(\mathcal{L}_G(u), \mathcal{L}_G(v))$ according to σ, and delete the rest. Similarly, for every vertex $v \in V(G)$, keep and mark the smallest spanning tree in $(G_1 \oplus_k G_2)[\mathcal{L}_G(v)]$ in accordance with σ (at least one spanning tree exists since $\mathcal{L}_G(v)$ must be connected in G). For an example of ψ_1, see Figure 5. We now define $\phi_k\colon \mathrm{MP}(G_1 \oplus_k G_2) \longrightarrow \mathrm{MP}(G_1) \times \mathrm{MP}(G_2)$ by $\phi_k(G) = (H_1, H_2)$, where H_i is $(\psi_k(G))[V(G_i)]$ with the marked edges contracted.

Now it is clear that ϕ_k is well defined and that it is an injection. So it suffices to show that $\mathrm{MP}(G_1 \oplus_k G_2)$ is a subposet of $\mathrm{MP}(G_1) \times \mathrm{MP}(G_2)$ by showing that ϕ_k preserves order relationships. It is worth observing that if $\phi_k(G) = (H_1, H_2)$ then H_1 may be thought of as the graph G with all references to vertices only in G_2 removed; similarly for H_2.

Suppose then that $A \preceq B \preceq G_1 \oplus_k G_2$. We first note that if A can be reached by vertex and edge deletion from B, then $\psi_k(A)$ is a subgraph of $\psi_k(B)$ and further, the marked edges are the same. Thus $\phi_k(A) < \phi_k(B)$. Now suppose that $A = B/e$ for some $e = \{u, v\} \in E(B)$. Then $\psi_k(A)$ is the same as $\psi_k(B)$ except over the vertices $\mathcal{L}_B(u) \cup \mathcal{L}_B(v)$. Furthermore, since $k \leq 1$, then at least one of $\mathcal{L}_B(u)$ or $\mathcal{L}_B(v)$ must consist of vertices entirely within G_1 or G_2. Without loss of generality, suppose then that $\mathcal{L}_B(u) \subseteq V(G_1)$. Then $\phi_k(A)$ and $\phi_k(B)$ agree in the second coordinate and so we restrict our attention to the first coordinate. However, it is clear that the first coordinate of $\phi_k(A)$ is formed by contracting the edge corresponding to e in the first coordinate of $\phi_k(B)$. Thus if $A \preceq B$ then $\phi_k(A) \leq \phi_k(B)$.

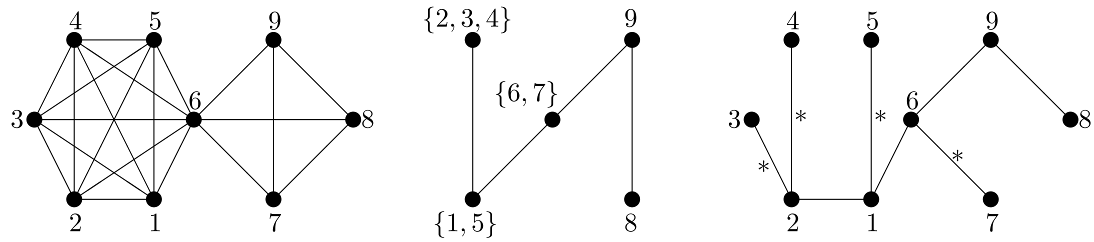

(a) $K_6 \oplus_1 K_4$ (b) $G' \preceq (K_6 \oplus_1 K_4)$ (c) $\psi_k(G')$

Figure 5: Example of ψ_k where σ is the lexicographic ordering.

Now suppose, by way of contradiction, that $A \parallel B$ in $\mathrm{MP}(G_1 \oplus_k G_2)$, $\phi_k(A) = (A_1, A_2)$, $\phi_k(B) = (B_1, B_2)$, and $A_1 \preceq B_1$ and $A_2 \preceq B_2$. Notice that each edge in any minor of $G_1 \oplus_k G_2$ can be uniquely associated with G_1 or G_2 and similarly every vertex can be uniquely associated with G_1 or G_2 (other than the vertex in the clique sum when $k = 1$). Therefore, we can obtain A by applying the minor operations which yield A_1 from B_1 and A_2 from B_2 to B, contradicting that $A \parallel B$. Thus ϕ_k is an embedding, as desired.

In order to show that $\mathrm{MP}(G_1 \oplus_0 G_2) = \mathrm{MP}(G_1) \times \mathrm{MP}(G_2)$ it suffices to show that ϕ_0 is invertible. However this follows trivially from the disjointness of G_1 and G_2 in $G_1 \oplus_0 G_2$. On the other hand, if $k = 1$, there exist vertices $v_1 \in V(G_1)$ and $v_2 \in V(G_2)$ that get identified in $G_1 \oplus_1 G_2$. As $(\mathcal{G}_{v_1}, \varnothing)$ is not in the range of ϕ_1, we see that ϕ_1 is not an isomorphism. □

For the next result we combine Lemma 12 with well-known results concerning the Cartesian product of posets. Specifically, for all posets $\mathbf{P}$ and $\mathbf{Q}$ we have $\dim(\mathbf{P} \times \mathbf{Q}) \leq \dim(\mathbf{P}) + \dim(\mathbf{Q})$ [10]. Further, equality holds for all non-trivial posets having a "1" and a "0"; that is, having unique maximal and minimal elements [4].

Lemma 13. *Let G_1 and G_2 be graphs. Then* $\dim(\mathrm{MP}(G_1 \oplus_1 G_2)) \leq \dim(\mathrm{MP}(G_1)) + \dim(\mathrm{MP}(G_2))$ *and* $\dim(\mathrm{MP}(G_1 \oplus_0 G_2)) = \dim(\mathrm{MP}(G_1)) + \dim(\mathrm{MP}(G_2))$.

Proof. By Lemma 12 we have that

$$\dim(\mathrm{MP}(G_1 \oplus_1 G_2)) \leq \dim(\mathrm{MP}(G_1) \times \mathrm{MP}(G_2)) \leq \dim(\mathrm{MP}(G_1)) + \dim(\mathrm{MP}(G_2)),$$

where the second inequality follows from [10]. Lemma 12 also yields

$$\dim(\mathrm{MP}(G_1 \oplus_0 G_2)) = \dim(\mathrm{MP}(G_1) \times \mathrm{MP}(G_2)) = \dim(\mathrm{MP}(G_1)) + \dim(\mathrm{MP}(G_2)),$$

where the second inequality follows from [4] since $\mathrm{MP}(G_i)$ has a "0" (the graph with no vertices) and has a "1" (the whole graph) which are distinct for $i \in \{1, 2\}$. □

The following theorem is an immediate consequence of Lemma 13.

Theorem 14. *Let G be a graph with connected components $C_1, C_2, \ldots, C_k$. Then*

$$\dim(\mathrm{MP}(G)) = \sum_{i=1}^{k} \dim(\mathrm{MP}(C_i)).$$

Let H be a connected graph. Let the blocks of H be enumerated $B_1, B_2, \ldots, B_k$. Then

$$\dim(\mathrm{MP}(H)) \leq \sum_{i=1}^{k} \dim(\mathrm{MP}(B_i)).$$

Remark 2. Amongst other things, Theorem 14 implies that $\dim(\mathrm{MP}(E_n)) = n$, where E_n is the graph with n vertices and zero edges. This should come as no surprise, since $\mathrm{MP}(E_n)$ is isomorphic to $\mathbf{2}^n$, as noted in section 3.1.

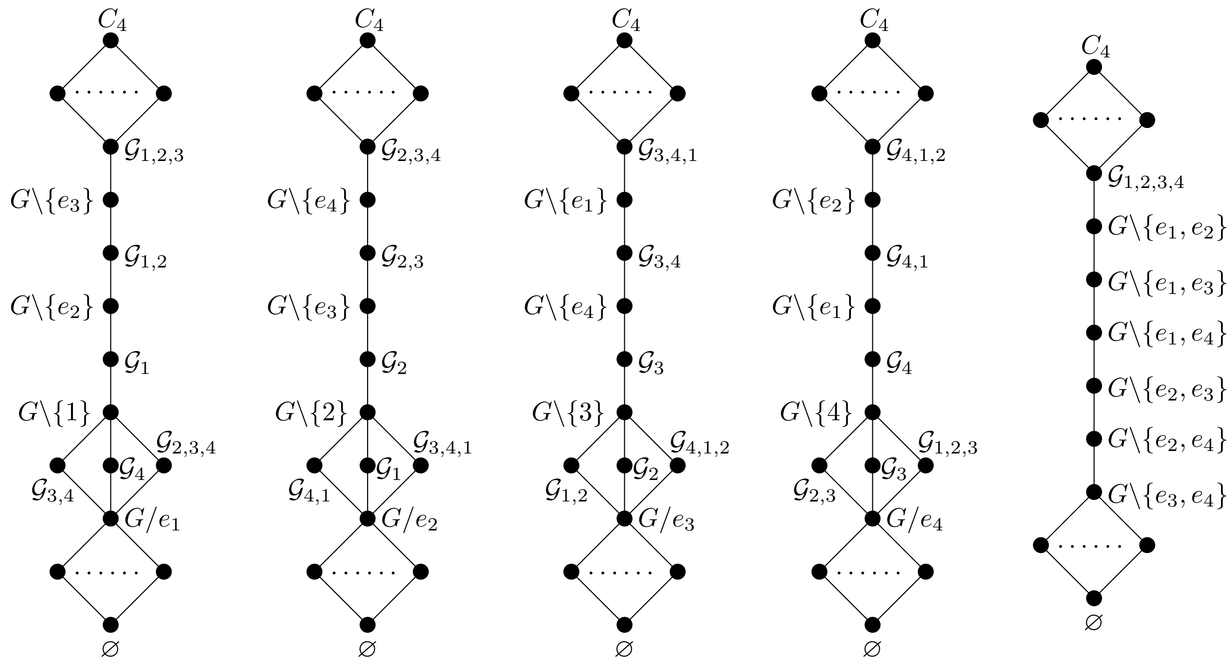

(a) L_1 (b) L_2 (c) L_3 (d) L_4 (e) L_5

Figure 6: Realizer of size five for $\mathrm{MP}(C_4)$, where $E(C_4) = \{\{1,2\},\{2,3\},\{3,4\},\{1,4\}\}$.

4. Dimension of the Minor Poset for Selected Classes of Graphs

Unsurprisingly, given the relationship between poset dimension and hypergraph coloring [20], the classification of the critical pairs in Lemma 10 quickly yields upper bounds on $\dim(\mathrm{MP}(G))$ for many graphs G. However, as we shall see in this section and the next, the calculation of lower bounds is non-trivial. This is especially true as the trivial lower bound of $\max\{|V(G)|,|E(G)|\}$ is rarely tight. Even in the following result, where $\dim(\mathrm{MP}(G))\,G = \max\{|V(G)|,|E(G)|\}+1$, a non-trivial pigeonhole argument is needed to lift the lower bound even one above the trivial bound.

Theorem 15. *Let G be a cycle with n vertices. Then $\dim(\mathrm{MP}(G)) = n+1$.*

Proof. Let the cycle have vertices $v_1, v_2, \ldots, v_n$ and edges $e_1, e_2, \ldots, e_n$ where $e_i = \{v_{i-1}, v_i\}$, where indicies are taken cyclically. We will employ the standard convention that $[i,j] = \{v_i, v_{i+1}, \ldots, v_j\}$ and $[n] = \{1,2,\ldots,n\}$.

To prove $\dim(\mathrm{MP}(G)) \le n+1$, we exhibit $n+1$ extensions that reverse all critical pairs. For each $i \in [n]$ let L_i be any linear extension of $\mathrm{MP}(G)$ that also has the following relations:

$$G/e_i < \left\{\mathcal{G}_{[i-1-j,i-1]}\right\}_{j=0}^{n-2} < G\backslash\{v_i\} < \mathcal{G}_{\{v_i\}} < G\backslash e_{i+1} < \mathcal{G}_{[i,i+1]} < \cdots < G\backslash\{e_{i+n-2}\} < \mathcal{G}_{[i,i+n-2]}.$$

We note that L_i reverses the critical pairs $\mathcal{V}_{v_i}$, those $\mathcal{E}_{E(U,V)}$ where $\{v_i, v_{i-1}\} \cap (U \cup V) = \{v_i\}$, and the $\mathcal{C}$ of the form $(\cdot, G/e_i)$. We note that all the critical pairs of the form $(\mathcal{G}_{V(G)}, \cdot)$ can be reversed in a single extension, L_{n+1}. For example, consider the realizer for C_4 in Figure 6. We observe that all of the critical pairs in $\mathcal{V}$ and $\mathcal{C}$ are reversed in $L_1, \ldots, L_n$. Now consider the critical pairs in $\mathcal{E}$. Since each of these corresponds to a minimal edge cut in a connected subgraph in G, there are two types of critical pairs in $\mathcal{E}$. The first are the single edge cuts, which are all reversed in one of $L_1, \ldots, L_n$, and the other are the two edge cuts, which are all reversed in L_{n+1}. Thus $\dim(\mathrm{MP}(G)) \le n+1$.

For the lower bound, suppose that $\dim(\mathrm{MP}(G)) \le n$. Let $\mathcal{R}$ be a realizer with n extensions, $L_1, L_2, \ldots, L_n$. At least one linear extension must reverse $\left(\mathcal{G}_{V(G)}, G\backslash\{e_1, e_n\}\right)$. Without loss of generality, this extension is L_n.

Now consider the set $C = \{\mathcal{C}_v : v \in V(G)\}$. As each vertex has degree two in G, $|C| = 2n$. Since for any vertex v_i, $\mathcal{G}_{\{v_i\}} \preceq G\setminus\{e_1, e_n\}$ and $\mathcal{G}_{V(G)} \preceq G/e_i, G/e_{i+1}$, there is an alternating cycle of length two formed between any critical pair in C and $\left(\mathcal{G}_{V(G)}, G\setminus\{e_1, e_n\}\right)$. Thus no critical pair in C is reversed in L_n and so, by the pigeonhole principle, there exists an extension in $\mathcal{R}$ which reverses at least 3 elements of C.

Since $|\mathcal{R}| = n$ and the critical pairs in $\mathcal{V}$ induce an S_n by Proposition 11, it must be the case that each L_i reverses a unique critical pair in $\mathcal{V}$. Now note that $\mathcal{G}_{\{v_i\}} \preceq G/e_j$ unless j is i or $i+1$. Further, if $j \neq i, i+1$, then $\mathcal{G}_{[j-1,j]} \preceq G\setminus\{v_i\}$, which forms an alternating cycle of length two. Thus the only elements of C that can be reversed with $\left(\mathcal{G}_{\{v_i\}}, G\setminus v_i\right)$ are

$$\left\{\left(\mathcal{G}_{\{v_i\}}, G\setminus\{e_i\}\right), \left(\mathcal{G}_{\{v_{i-1}\}}, G\setminus\{e_i\}\right), \left(\mathcal{G}_{\{v_i\}}, G\setminus\{e_{i+1}\}\right), \left(\mathcal{G}_{\{v_{i+1}\}}, G\setminus\{e_{i+1}\}\right)\right\}.$$

Thus there is some set of three consecutive vertices u, v and w and a linear extension $L \in \mathcal{R}$ such that L reverses $\left(\mathcal{G}_{\{v\}}, G\setminus\{v\}\right)$ and at least three of

$$\left(\mathcal{G}_{\{v\}}, G/\{u,v\}\right), \left(\mathcal{G}_{\{u\}}, G/\{u,v\}\right), \left(\mathcal{G}_{\{w\}}, G/\{w,v\}\right), \text{ and } \left(\mathcal{G}_{\{v\}}, G/\{w,v\}\right).$$

However, by Proposition 11, the critical pairs in $\mathcal{E}_{\{\{v_j,v_{j+1}\}\}}$ for $1 \leq j \leq n$ induce a copy of S_n and thus L must also reverse a critical pair of this form. Thus let $e = \{x, y\}$ be the edge so that $\left(\mathcal{G}_{\{x,y\}}, G\setminus\{e\}\right)$ is reversed in L. Notice that if $v \notin \{x, y\}$ then $\mathcal{G}_{\{v\}} \preceq G\setminus\{e\}$ and $\mathcal{G}_{\{x,y\}} \preceq G\setminus\{v\}$, forming an alternating cycle of length two. Thus $v \in \{x, y\}$. Suppose that $\{x, y\} = \{w, v\}$. Now we note that $\mathcal{G}_{\{w\}}, \mathcal{G}_{\{v\}} \preceq G\setminus\{\{w, v\}\}$ and $\mathcal{G}_{\{w,v\}} \preceq G/\{w, v\}$, and thus neither $\left(\mathcal{G}_{\{w\}}, G/\{w, v\}\right)$ nor $\left(\mathcal{G}_{\{v\}}, G/\{w, v\}\right)$ can be reversed in L. Similarly, if $\{x, y\} = \{u, v\}$. But this contradicts the fact that L reverses three critical pairs from C, and thus $\dim(\mathrm{MP}(G)) > n$, finishing the proof. □

In the next theorem we provide the best known bounds for the dimension of the minor poset of a complete graph. The exponential lower bound is far better than the quadratic bound implied by Theorem 11.

Theorem 16. *Let $n \geq 2$ be an integer. Then*

$$\frac{4^{n-1}}{\sqrt{n-1}} \leq \sum_{i=1}^{n-1}\left(\binom{2n-2i}{n-i} + \binom{2n-2i+1}{n-i}\right) \leq \dim(K_{2n}) \leq 4^n.$$

Proof. We first show that $\dim(K_{2n}) \leq 4^n$. To this end, notice that for every $X \subseteq V(G)$, all of the critical pairs of the form $(\mathcal{G}_X, \cdot)$ can be reversed in a single linear extension by Remark 1. Since every critical pair is of this form at most $2^{2n} = 4^n$ linear extensions suffice.

Next we turn to the lower bound. For this we exhibit a standard example of the appropriate size, made entirely of critical pairs from $\mathcal{E}$. Let $V(K_{2n}) = \{1, 2, \ldots, 2n\}$ and let $\binom{[i,j]}{k}$ for $i \leq j$ represent all subsets of $\{i, i+1, \ldots, j\}$ of size k. Then, for $1 \leq i \leq 2n-2$, define

$$X_i = \left\{\mathcal{E}_{E(\{i\},T)} : T \in \binom{[i+1, 2n]}{n - \lceil \frac{i}{2} \rceil}\right\}$$

Consider two such critical pairs $\left(\mathcal{G}_{\{i\}\cup T_1}, K_{2n}\setminus E(\{i\}, T_1)\right)$ and $\left(\mathcal{G}_{\{j\}\cup T_2}, K_{2n}\setminus E(\{j\}, T_2)\right)$ which are distinct. Note that $\mathcal{G}_{\{i\}\cup T_1} \preceq K_{2n}\setminus E(\{j\}, T_2)$ unless $\{i\}\cup T_1 \subseteq \{j\}\cup T_2$ and $j \in \{i\}\cup T_1$. Now since $\{i\}\cup T_1 \subseteq [i, 2n]$ and $\{j\}\cup T_2 \subseteq [j, 2n]$, this implies that the only way these two critical pairs do not form an alternating cycle is if $i = j$ and $T_1 = T_2$, contradicting that they are distinct. Thus every pair of distinct critical pairs in $\cup_{i=1}^{2n-2} X_i$ induces an alternating cycle of length two; hence it induces a standard example. Furthermore, since $E(\{i\}, T_1)$ is distinct from $E(\{j\}, T_2)$ unless $i = j$ and $T_1 = T_2$, no minor appears in more than one critical pair in $\cup_{i=1}^{2n-2} X_i$. Thus the standard example has size

$$\left|\bigcup_{i=1}^{2n-2} X_i\right| = \sum_{i=1}^{n-1}\left(|X_{2i-1}| + |X_{2i}|\right) = \sum_{i=1}^{n-1}\left(\binom{2n-2i+1}{n-k} + \binom{2n-2i}{n-i}\right).$$

This quantity can be bounded below as follows:

$$\sum_{i=1}^{n-1}\left(\binom{2n-2i+1}{n-i}+\binom{2n-2i}{n-i}\right) \geq \sum_{i=1}^{n-1}\left(\frac{2n-2i+1}{n-i+1}\binom{2(n-i)}{n-i}+\binom{2(n-i)}{n-i}\right)$$
$$\geq \sum_{i=1}^{n-1}\left(3-\frac{1}{n-i+1}\right)\frac{4^i}{2\sqrt{i}}$$
$$\geq \sum_{i=1}^{n-1}\frac{5}{2}\cdot\frac{4^i}{2\sqrt{i}}$$
$$\geq \sum_{i=1}^{n-1}\frac{4^i}{\sqrt{i}}.$$

This last summation is clearly dominated by the final term in the summation, yielding the simplified lower bound that is stated in the result. □

For emphasis, we state the following corollary.

Corollary 17. *For all $n \geq 4$, $\log(\dim(\mathrm{MP}(K_n))) = (1+o(1))n$.*

5. Trees

Let T be a tree. We begin this section by defining a collection of *fundamental* critical pairs for $\mathrm{MP}(T)$. For each vertex $v \in V(T)$, the critical pair $(\mathcal{G}_v, T\backslash v) \in \mathcal{V}$ is fundamental. For each edge $e = \{u, v\} \in E(T)$ the critical pair $(\mathcal{G}_{u,v}, T\backslash e) \in \mathcal{E}$ is fundamental. Additionally, the critical pairs $(\mathcal{G}_u, T/e), (\mathcal{G}_v, T/e) \in \mathcal{C}$ are fundamental.

The following lemma states which fundamental critical pairs can be reversed in a single linear extention of $\mathrm{MP}(T)$. Figure 7 demonstrates these sets pictorially, where "ν" refers to the fundamental critical pair in $\mathcal{V}$ associated with the nearest vertex, "ε" refers to the fundamental critical pair in $\mathcal{E}$ associated with the nearest edge, and "c" refers to the fundamental critical pair in $\mathcal{C}$ associated with the nearest vertex and edge.

Lemma 18. *Let T be a tree and let L be a linear extension of $\mathrm{MP}(T)$. The fundamental critical pairs reversed in L are a subset of either*

(i) $\{(\mathcal{G}_v, T\backslash v)\} \cup \{(\mathcal{G}_u, T/e_{uv})\} \cup \{(\mathcal{G}_v, T/e_{vw}) \mid w \in N(v))\}$, *or*
(ii) $\{(\mathcal{G}_v, T\backslash v)\} \cup \{(\mathcal{G}_{uv}, T\backslash e_{uv})\} \cup \{(\mathcal{G}_v, T/e_{vw}) \mid w \in N(v)\backslash u\} \cup \{(\mathcal{G}_{u'}, T/e_{u'v})\}$ *where $e_{u'v}$ is distinct from e_{uv},*

for some $v \in V(T)$ and some $u \in N(v)$.

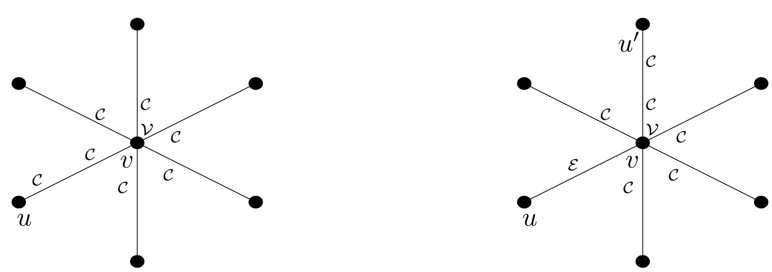

(a) A (i)-subset of fundamental critical pairs.

(b) A (ii)-subset of fundamental critical pairs.

Figure 7

Proof. By Proposition 11, two distinct critical pairs in $\mathcal{V}$ cannot be simultaneously reversed. Similarly for $\mathcal{E}$. Now consider two fundamental critical pairs in $\mathcal{C}$, $(\mathcal{G}_u, T/e_{uv})$ and $(\mathcal{G}_x, T/e_{xy})$. If $|\{u,v\} \cap \{x,y\}| = 0$, then these pairs induce an alternating cycle. Thus $|\{u,v\} \cap \{x,y\}| \geq 1$. Suppose then that e_{uv}, e_{xy}, and $e_{u'v'}$ are three distinct edges with each pair having exactly one endpoint in common. Then since T is a tree, there exists a vertex, say v, that is in all three edges. Thus the fundamental critical pairs in $\mathcal{C}$ must be a subset of $\{(\mathcal{G}_v, T/e_{uv})\} \cup \{(\mathcal{G}_u, T/e_{uv})\} \cup \{(\mathcal{G}_v, T/e_{wv}) \mid w \in N(v)\backslash u\}$. In sum, the fundamental critical pairs in $\mathcal{V}$ and $\mathcal{E}$ must be reversed in individual extensions and the fundamental critical pairs in $\mathcal{C}$ can be reversed in a "star" around a single vertex with at most one edge having both ends reversed in critical pairs in $\mathcal{C}$.

We now turn to the interrelations between the sets of fundamental critical pairs. Consider the fundamental critical pair $(\mathcal{G}_v, T\backslash v) \in \mathcal{V}$ and the fundamental critical pair $(\mathcal{G}_{xy}, T\backslash e_{xy}) \in \mathcal{E}$. Note that $\mathcal{G}_v \preceq T\backslash e_{xy}$ and $\mathcal{G}_{xy} < T\backslash v$ unless $v \in \{x,y\}$. Thus the critical pairs in $\mathcal{V}$ and $\mathcal{E}$ must have a vertex in common. Now instead of $(\mathcal{G}_{xy}, T\backslash e_{xy})$ consider the fundamental critical pair $(\mathcal{G}_x, T/e_{xy}) \in \mathcal{C}$. Clearly if $x = v$ there is no alternating cycle between the two critical pairs, thus suppose $x \neq v$. Then $\mathcal{G}_x \preceq T\backslash v$ and then unless $v \in \{x,y\}$, there is an alternating cycle. Thus if there is a reversed critical pair in $\mathcal{V}$, the reversed critical pairs in $\mathcal{C}$ must be on edges incident with a common vertex.

Now consider fundamental critical pairs $(\mathcal{G}_{xy}, T\backslash e_{xy}) \in \mathcal{E}$ and $(\mathcal{G}_v, T/e_{uv}) \in \mathcal{C}$. First note that $\mathcal{G}_v \preceq T\backslash e_{xy}$. Further if $\{u,v\} \cap \{x,y\} = \varnothing$ or $\{u,v\} = \{x,y\}$, then $\mathcal{G}_{xy} \preceq T/e_{uv}$ forming an alternating cycle. Thus the fundamental critical pairs in $\mathcal{E}$ and $\mathcal{C}$ must share a common vertex but not involve the same edge. Therefore the set of fundamental critical pairs reversed by a linear extension L must be a subset of either (i) or (ii). □

For convenience of notation, we denote a linear extension reversing a maximal number of fundamental critical pairs by L_{uv} if it is of type (i) and as $L_{uvu'}$ if it is of type (ii). Note that $L_{uv\varnothing}$ and $L_{\varnothing vu'}$ are maximal type (ii) linear extensions when v is a leaf.

In order to bound the dimension of the minor poset for a tree, T, we introduce a few terms. Let o_T denote the number of vertices of odd degree in T (recall that every graph has an even number of vertices of odd degree). Define a *path decomposition* of T as a three-tuple $(\mathcal{P}, \mathcal{O}, \mathcal{S})$ where $\mathcal{P}$ is a partition of the edges of T into $o_T/2$ paths, $\mathcal{O}$ is an orientation of the edges of T (we do not require that the orientation of the edges be consistent along any path in $\mathcal{P}$), and $\mathcal{S} \colon \mathcal{P} \to 2^{V(T)}$ is a function that assigns a subset of $V(T)$ to each path. For $P \in \mathcal{P}$, the set $\mathcal{S}(P)$ is referred to as the *stars* of P. Additionally, we require the following properties of a path decomposition:

(i) for every path $P = p_0, p_1, \ldots, p_k$ in $\mathcal{P}$, we have $\mathcal{S}(P) \subseteq \{p_i\}_{i=0}^{k}$,

(ii) $\left(\bigcup_{P \in \mathcal{P}} \mathcal{S}(P)\right) \cup \{v \in V(T) \mid \text{ there exists a } u \in V(T) \text{ such that } (u,v) \in \mathcal{O}\} = V(T)$,

(iii) if $\mathcal{S}(P) = \varnothing$ (such paths are *special*) then P can be decomposed into subpaths $S_1, S_2, \ldots, S_\ell$ where each subpath S_j satisfies one of the following:

  (a) $S_j = s_0, s_1$ and $s_0, s_1 \in \mathcal{I}(P)$,

  (b) $S_j = s_0, s_1, s_2$ and $(s_0, s_1), (s_2, s_1) \in \mathcal{O}$ and $s_0, s_2 \in \mathcal{I}(P)$,

  (c) $S_j = s_0, s_1, s_2, \ldots s_m$ and $(s_i, s_{i+1}) \subseteq \mathcal{O}$ and $s_0, s_1 \in \mathcal{I}(P)$,

  and $\mathcal{I}(P) = \{v \in V(P) \mid \exists u \notin V(P), (u,v) \in \mathcal{O}\} \cup \{v \in V(P) \mid \exists P' \in \mathcal{P}, P' \neq P, v \in \mathcal{S}(P')\}$. Less formally, $\mathcal{I}(P)$ is the set of vertices of P that have a non-P edge oriented towards them or that are a star of some $P' \neq P$.

Figure 8 provides an example of a path decomposition. In that example, $\mathcal{P} = \{P_1, P_2, P_3\}$ where $P_1 = v_1, v_3, v_2$, $P_2 = v_3, v_4, v_5$, and $P_3 = v_6, v_5, v_7$. The orientation $\mathcal{O}$ is defined by the arrows on the edges. The stars of each path are $\mathcal{S}(P_1) = \{v_1\}$, $\mathcal{S}(P_2) = \varnothing$, and $\mathcal{S}(P_3) = \{v_7\}$. Thus P_2 is special. Notice that P_2 satisfies condition (b) of property (iii).

Remark 3. It is not hard to see that every tree T has at least one path decomposition. To this end, pick two odd-degree vertices $u, v \in V(T)$ and let P be the unique path with ends u and v. Put P into $\mathcal{P}$. Notice that the degrees of u and v in $T \setminus P$ are even, whereas all other vertices in $V(T)$ have maintained the parity of their degree. Repeat this deletion procedure inductively on the components of $T \setminus P$. This procedure terminates when all vertices have degree zero, at which point each odd-degree vertex is the end of exactly one path in $\mathcal{P}$. Since the ends of each path are odd-degree vertices in T, there are precisely $o_T/2$ paths in $\mathcal{P}$. Adding an arbitrary orientation of the edges and an appropriate function $\mathcal{S}$ yields a path decomposition of T.

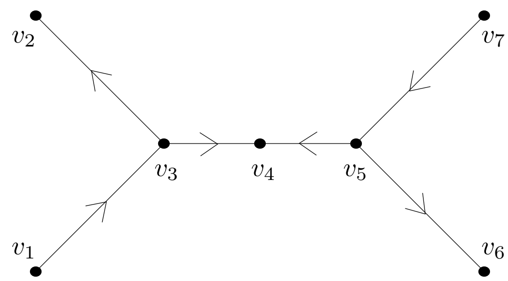

Figure 8: This orientation, $\mathcal{O}$, together with $\mathcal{P}$ and $\mathcal{S}$, as described in the text, form a path decomposition of the tree above.

Define the *size* of a path decomposition $\mathcal{D} = (\mathcal{P}, \mathcal{O}, \mathcal{S})$ as $|\mathcal{D}| = \sum_{P \in \mathcal{P}} (|E(P)| + |\mathcal{S}(P)|)$. That is, the size of a path decomposition is the number of edges in T plus the number of stars in the decomposition. We will denote by $\mathfrak{D}$ the set of path decompositions for a fixed tree T.

For a path $P = p_0, p_1, \ldots, p_k$ we will denote by $P_{p_i,p_{i+1}}$ the subpath $p_i, p_{i+1}, \ldots, p_k$. Similarly $P_{\overline{p_i,p_{i+1}}}$ will denote the complementary subpath $p_i, p_{i-1}, \ldots, p_0$ (we use this notation instead of $P_{p_i,p_{i-1}}$ in the case that we have limited information about the vertex p_{i-1}). Note that $P_{p_i,p_{i+1}} \cup P_{\overline{p_i,p_{i+1}}} = P$ and $P_{p_i,p_{i+1}} \cap P_{\overline{p_i,p_{i+1}}} = \{p_i\}$.

Theorem 19. *Let T be a tree with $|E(T)| \geq 1$. The fundamental critical pairs can be reversed with $\min_{\mathcal{D} \in \mathfrak{D}} |\mathcal{D}|$ linear extensions of* $\mathrm{MP}(T)$*, and this is a best possible.*

Proof. We first show that we can reverse the fundamental critical pairs with $\min_{\mathcal{D} \in \mathfrak{D}} |\mathcal{D}|$ linear extensions. To this end let $\mathcal{D} = (\mathcal{P}, \mathcal{O}, \mathcal{S})$ be a path decomposition of T and let $P = p_0, p_1, \ldots, p_k$ be a non-special path in the decomposition. Let $\mathcal{O}'$ be an orientation of the edges of T that agrees with $\mathcal{O}$ on $E(T) \backslash E(P)$ and has the oriented edge (p_i, p_{i+1}) for $0 \leq i \leq k-1$. Further, let $\mathcal{S}' \colon \mathcal{P} \to 2^{V(T)}$ be a function that satisfies $\mathcal{S}'(P) = \{p_0\}$ and $\mathcal{S}'(P') = \mathcal{S}(P')$ for every path $P' \in \mathcal{P}$ with $P' \neq P$. Now $\mathcal{D}' = (\mathcal{P}, \mathcal{O}', \mathcal{S}')$ is a valid path decomposition for T as every vertex in $V(P)$ has indegree at least one or has an associated star, and moreover $|\mathcal{D}'| \leq |\mathcal{D}|$. Thus, without loss of generality, we may consider a path decomposition of T in which every non-special path has exactly one star, that star is at one end of the path, and the edges of the non-special paths are oriented away from the star.

Now consider the following set of linear extensions of $\mathrm{MP}(T)$: for each edge $(u, v) \in \mathcal{O}'$ with u' as the next vertex on the path containing edge $\{u, v\}$ (note that u' may not exist if v is a leaf of T), consider the extension $L_{uvu'}$. For every star on v with u as the adjacent vertex on the path associated with the star, consider the extension L_{uv}. Then we claim that the union of these extensions reverses all fundamental critical pairs in $\mathrm{MP}(T)$.

We first note that, since every vertex has an indegree or is a star, all of the critical pairs in $\mathcal{V}$ are reversed. Since every edge $\{u, v\}$ has an orientation, say (u, v), every critical pair in $\mathcal{E}$ is reversed by some extension $L_{uvu'}$. Thus it suffices to show that the $2|E(T)|$ critical pairs in $\mathcal{C}$ are reversed. Suppose that $p_0, p_1, \ldots, p_k$ is a non-special path in the path decomposition, with edges $(p_i, p_{i+1}) \in \mathcal{O}'$ for each $i \in \{0, 1, \ldots, k-1\}$. Then, for $1 \leq i \leq k-1$, the two critical pairs in $\mathcal{C}$ that are associated with the edge $\{p_i, p_{i+1}\}$ are reversed in $L_{p_{i-1}p_ip_{i+1}}$. Since there is a star at p_0, the two critical pairs in $\mathcal{C}$ that are associated with the edge $\{p_0, p_1\}$ are reversed by $L_{p_0p_1}$. Thus we only need to consider the critical pairs in $\mathcal{C}$ that are on edges of special paths. Further, it suffices to show that the critical pairs in $\mathcal{C}$ that correspond to edges in a given subpath are reversed. But this is clear, as the prescribed orientations and inclusions in $\mathcal{I}(P)$ (conditions (a), (b), and (c) in the definition of special path) are precisely the conditions needed to ensure that all of the critical pairs in $\mathcal{C}$ are reversed for each type of subpath.

In order to show that we require at least $\min_{\mathcal{D} \in \mathfrak{D}} |\mathcal{D}|$ linear extensions of $\mathrm{MP}(T)$ to reverse all fundamental critical pairs, we will need some new terminology. Call a set of linear extensions of $\mathrm{MP}(T)$ that reverses all of the fundamental critical pairs an *$\mathcal{F}$-realizer*. An $\mathcal{F}$-realizer is *optimal* there does not exist an $\mathcal{F}$-realizer with fewer linear extensions.

Now consider optimal $\mathcal{F}$-realizers in which each linear extension reverses a maximal set of critical pairs (no other critical pair can be reversed without introducing an alternating cycle) and in which no critical pair in $\mathcal{E}$ is reversed in more than one extension. Any $\mathcal{F}$-realizer can be transformed into this form by first saturating all of the linear extensions (reversing additional critical pairs until the addition of any new reversed critical pair introduces an alternating cycle). Second, for any edge $\{u, v\}$ whose associated fundamental critical pair in $\mathcal{E}$ is reversed in more than one linear extension, we can replace all but one such extension by $L_{u'v}$ for the appropriate u' (see Figure 7). We will refer to such an $\mathcal{F}$-realizer as *saturated*. Further, we will show that there is a saturated optimal $\mathcal{F}$-realizer from which we can build a path decomposition of T of the same size.

Let $\mathcal{R}$ be a saturated optimal $\mathcal{F}$-realizer for $\mathrm{MP}(T)$ and let $\mathcal{P}$ be a partition of the edges of T into $o_T/2$ paths. We can uniquely associate each linear extension in $\mathcal{R}$ to a path in $\mathcal{P}$ by associating the linear extensions of the form L_{uv} and $L_{uvu'}$ with the path containing the edge $\{u, v\}$. Since there is precisely one linear extension reversing the critical pair in $\mathcal{E}$ associated with $\{u, v\}$, orient each edge as (u, v) if $L_{uvu'}$ is in $\mathcal{R}$ or as (v, u) if $L_{vuu'}$ is in $\mathcal{R}$, for some vertex u'. Additionally, associate extensions of the form L_{uv} to stars at vertex v on the path containing the edge $\{u, v\}$.

We would like the set of paths associated with the $\mathcal{F}$-realizer to be a path decomposition of T. In order to find such a path decomposition, we define two quantities to measure the relationship between a saturated optimal $\mathcal{F}$-realizer and a path decomposition. Given an optimal saturated $\mathcal{F}$-realizer $\mathcal{R}$ and a partition $\mathcal{P}$ of T into paths we define the *consistency*, $C(\mathcal{R}, \mathcal{P})$, of the pair as the number of linear extensions in $\mathcal{R}$ that are of the form $L_{uvu'}$ such that either u, v, u' is a subpath of a path in $\mathcal{P}$ or v is the end of the path in $\mathcal{P}$ containing edge $\{u, v\}$. We also define the *switch points* of $(\mathcal{R}, \mathcal{P})$ as the set of vertices v in which the associated critical pair in $\mathcal{V}_v$ is reversed by two different linear extensions in $\mathcal{R}$ that are associated to the same path in $\mathcal{P}$. We denote the number of switch points as $S(\mathcal{R}, \mathcal{P})$.

From the collection of ordered pairs $(\mathcal{R}, \mathcal{P})$, where $\mathcal{R}$ is an optimal $\mathcal{F}$-realizer and $\mathcal{P}$ is a partition of the edges of T into $o_T/2$ paths, consider a pair which maximizes $C(\mathcal{R}, \mathcal{P})$ and subject to that minimizes $S(\mathcal{R}, \mathcal{P})$. Let $\mathcal{O}$ be the orientations derived from $\mathcal{R}$ and let $\mathcal{S}$ be the star associations. If $(\mathcal{P}, \mathcal{O}, \mathcal{S})$ is a path decomposition of T, we are done. If not, it must be the case that a path with no stars does not fit the requirements of a special path, as all conditions are easily verified. So consider a path $P^* \in \mathcal{P}$ with $\mathcal{S}(P^*) = \varnothing$. We will either show that this path is special or that $(\mathcal{P}, \mathcal{O}, \mathcal{S})$ can be altered to obtain a valid path decomposition of T of the same size.

Partition the edges of P^* into its maximal consistently-oriented subpaths, where "consistently-oriented" means that no vertex in the subpath has indegree more than one within P^*. Label the subpaths $S_1, S_2, \ldots, S_l$, where $|V(S_i) \cap V(S_{i+1})| = 1$ for $i \in \{1, 2, \ldots, l-1\}$. Define a *subpath-pair* to be a pair of consecutive subpaths $\{S_i, S_{i+1}\}$ of P^* that are oriented toward each other; that is, $S_i = v_1, v_2, \ldots, v_n$ and $S_{i+1} = u_1, u_2, \ldots, u_m$ is a subpath-pair if $V(S_i) \cap V(S_{i+1}) = \{v_n\} = \{u_m\}$ and $\{(v_j, v_{j+1})\}_{j=1}^{n-1}, \{(u_j, u_{j+1})\}_{j=1}^{m-1} \in \mathcal{O}$. Notice that, with the possible exception of S_1 and/or S_l, each subpath S_i is a member of a unique subpath-pair. In the case that S_1 is not in a subpath-pair, artifically pair it with the empty-path. Do the same for S_l. In this way the subpath-pairs partition the subpaths of P^*.

Consider a subpath-pair $\{S_i, S_{i+1}\}$. Let s_0 be the root vertex of S_i (the vertex with zero indegree in S_i) and let the subsequent vertex in S_i be s_1 (note that each subpath has at least one edge). If there are edges not in P^* directed into s_0 and s_1 or if there are stars at s_0 and s_1 (which by the choice of P^* are not associated with P^*), then S_i satisfies condition (a) or (c) in the definition of special path. So suppose not. Note that we must have $s_0 \in \mathcal{I}(P^*)$ since no extension associated with P^* reverses the critical pair in $\mathcal{V}_{s_0}$. Thus, since each extension in $\mathcal{R}$ is saturated, the critical pair $(\mathcal{G}_{s_0}, T/e_{s_0 s_1}) \in \mathcal{C}$ is reversed in the same extension. Now consider the ways in which the critical pair $(\mathcal{G}_{s_1}, T/e_{s_0 s_1}) \in \mathcal{C}$ could be reversed with $s_1 \notin \mathcal{I}(P^*)$. There is either some $v \notin V(P^*)$ such that $L_{v s_0 s_1} \in \mathcal{R}$ or there is some $s' \in S_{i+1}$ such that $L_{s' s_1 v'} \in \mathcal{R}$ for some vertex v'. Notice that in the latter case $S_i = s_0, s_1$.

Suppose that $L_{v s_0 s_1} \in \mathcal{R}$ and let P' be the path containing the edge $\{v, s_0\}$ (e.g. see Figure 9a). Consider a new parition of $E(T)$ into $o_T/2$ paths, $\mathcal{P}'$, which is identical to $\mathcal{P}$ except that it contains paths $P^*_{s_0 s_1} \cup P'_{s_0 v}$ and $P^*_{\overline{s_0 s_1}} \cup P'_{\overline{s_0 v}}$ instead of P^* and P'. We observe that any changes in consistency and/or switch points must occur surrounding the vertex s_0. First consider $C(\mathcal{R}, \mathcal{P})$ relative to $C(\mathcal{R}, \mathcal{P}')$. Since the extension $L_{v s_0 s_1}$ is consistent with the paths in $\mathcal{P}'$, we have $C(\mathcal{R}, \mathcal{P}') > C(\mathcal{R}, \mathcal{P})$ unless there is a $v' \in P'$ such that $L_{v' s_0 v} \in \mathcal{R}$. However, in this case $S(\mathcal{R}, \mathcal{P}') < S(\mathcal{R}, \mathcal{P})$, where we have used the fact that s_0 is the root of its subpath. Thus, by the choice of $\mathcal{P}$, there is no v not on P^* such that $L_{v s_0 s_1} \in \mathcal{R}$.

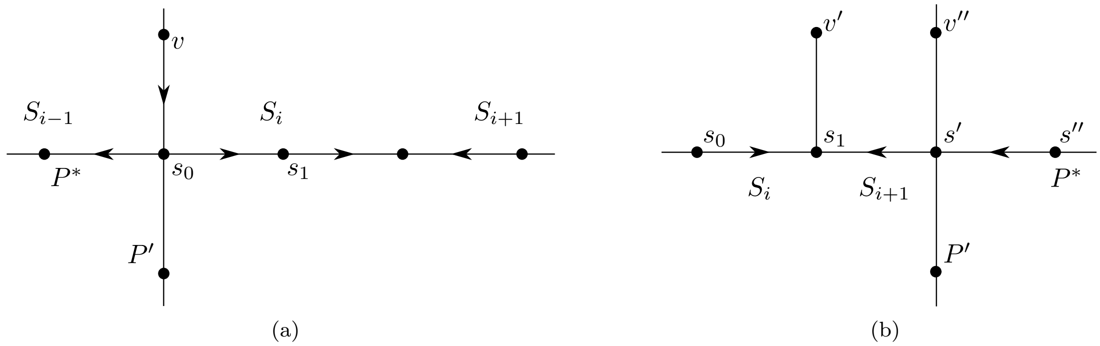

Figure 9: In 9b it may be the case that $v' = s_0$ or $v'' = s_1$.

So we may assume that there exists some $s' \in S_{i+1}$ such that $L_{s's_1v'} \in \mathcal{R}$, and, in particular $(s', s_1) \in \mathcal{O}$ is an edge of S_{i+1}. If $s' \in \mathcal{I}(P^*)$ then $S_i \cup \{s'\}$ satisfies condition (b) in the defintion of special path. Hence, we can remove s' from S_{i+1} and repeat the arguments above to $S_{i+1} - \{s'\}$ (unless it is now empty). Notice that the orientation of the edge $\{s', s_1\}$ prevents this argument from cycling, and hence S_{i+1} will satisfy either condition (a) or (c). Thus, we have established that the edges comprising the subpath-pair $\{S_i, S_{i+1}\}$ can be partitioned into two disjoint subpaths, each of which is one of (a), (b), or (c). Therefore, by way of contradiction, we may assume that $s' \notin \mathcal{I}(P^*)$. Since the critical pair in $\mathcal{V}_{s'}$ must be reversed, there is a vertex $s'' \in S_{i+1}$ such that $L_{s''s'v''} \in \mathcal{R}$ for some vertex v'' (e.g. see Figure 9b). Assume for the moment that $v'' \neq s_1$ and let $P' \in \mathcal{P}$ be the path containing the edge $\{s', v''\}$. Then, since $s' \notin \mathcal{I}(P^*)$, we can replace P^* and P' by $P^*_{s's''} \cup P'_{s'v''}$ and $P^*_{\overline{s's''}} \cup P'_{\overline{s'v''}}$ to obtain a partition of $E(T)$ that contradicts the choice of $\mathcal{P}$, as in the previous case.

Thus we may assume that $v'' = s_1$, and in particular, the extension $L_{s''s's_1}$ reverses both of the critical pairs in $\mathcal{C}$ associated with the edge $\{s_1, s'\}$. We may also assume that $L_{s_0s_1s'} \in \mathcal{R}$ as otherwise we can apply an analogous argument to that in the previous paragraph to contradict the choice of $\mathcal{P}$. Thus, the only critical pair that is reversed in $L_{s_0s_1s'}$ that is not reversed elsewhere in $\mathcal{R}$ is the critical pair $(\mathcal{G}_{s_0s_1}, T\backslash e_{s_0s_1})$. Therefore, we may replace $L_{s_0s_1s'}$ in $\mathcal{R}$ with $L_{s_1s_0s_{-1}}$, where s_{-1} is the vertex adjacent to s_0 in P^* that isn't s_1 (if no such vertex exists, which may be the case for S_1 or S_l, let $s_{-1} = \varnothing$). The resulting realizer/path partition pair has at least the same consistency and one less switch point (since s_0 is the root of S_i) in comparison to the original, contradicting the choice of $\mathcal{P}$. In this case, we have found S_i to satisfy condition (a) or (c) of the definition of special path. Next, we apply the analogous arguments to S_{i+1}. Again, we find that the edges comprising the subpath-pair $\{S_i, S_{i+1}\}$ can be partitioned into two disjoint subpaths, each of which is one of (a), (b), or (c).

Since the subpath-pairs partition the subpaths of P^*, and since there is no point in the argument in which we swap edges between subpath-pairs, we can apply all of the arguments above to each pair separately. After doing so, we find that the edges of P^* can be partitioned into disjoint subpaths, each of which is one of (a), (b), or (c). Hence P^* is special, which completes the proof. $\square$

Theorem 19 proves that $\dim(\mathrm{MP}(T)) \geq \min_{\mathcal{D}\in\mathfrak{D}} |\mathcal{D}|$, since we require that number of extensions just to reverse the fundamental critical pairs. The next theorem yields a similar upper bound.

Theorem 20. *Let T be a tree. Then* $\dim(\mathrm{MP}(T)) \leq |E| + \frac{o_T}{2}$.

Proof. Remark 3 guarantees the existence of a path decomposition of T into exactly $o_T/2$ paths. Let $\mathcal{P}$ be one such decomposition, and let $P = v_1, v_2, \ldots, v_k$ be an arbitrary path in $\mathcal{P}$. We aim to show that the critical pairs $\{\mathcal{V}_{v_i}\}_{i=1}^{k}$, $\{(\,\cdot\,, T\backslash e_{v_iv_{i+1}})\}_{i=1}^{k-1} \subseteq \mathcal{E}$, and $\{(\,\cdot\,, T/e_{v_iv_{i+1}})\}_{i=1}^{k-1} \subseteq \mathcal{C}$ can be reversed in k linear extensions of $\mathrm{MP}(T)$. Notice that this suffices in proving the result; if we do this for any path in $\mathcal{P}$ then all critical pairs in $\mathcal{V}$, $\mathcal{C}$, and $\mathcal{E}$ will be reversed (all critical pairs in $\mathcal{E}$ are of the form $T\backslash e$ for some $e \in E(T)$ since all minimal edgecuts in trees have exactly one edge), and $\sum_{P\in\mathcal{P}} |V(P)| = |E(T)| + o_T/2$.

We note that every single-vertex graph which contains a vertex in P participates in some critical pair of the form $(\cdot , T\backslash e_{v_iv_{i+1}})$ or $(\cdot , T/e_{v_iv_{i+1}})$. We introduce some notation to specify a limited class of such single-vertex graphs. Let $P[i,j]$ be the set of all $S \subseteq V(T)$ such that $\{v_i, v_j\} \subseteq S$ and S induces a connected subgraph of T. Let $P^+[i,j] \subseteq P[i,j]$ be the subset consisting of sets S where $v_{i+1} \notin S$. Similarly, let $P^-[i,j] \subseteq P[i,j]$ consist of the subsets S where $v_{i-1} \notin S$. Lastly, let $P^{\pm}[i,j] = P^+[i,j] \cap P^-[i,j]$. If $i = j$ then we may simply write $P[i]$, $P^+[i]$, $P^-[i]$, and $P^{\pm}[i]$, respectively.

Next, we define k sets of critical pairs with respect to P. Let

$$J_i = \begin{cases} \mathcal{V}_{v_1} \cup \{(\cdot , T/e_{v_1v_2})\} & i = 1 \\ \mathcal{V}_{v_i} \cup \left\{(\cdot , T/e_{v_iv_{i+1}})\right\} \cup \left(\bigcup_{j=1}^{i-1} \left\{(\mathcal{G}_S, T\backslash e_{v_jv_{j+1}}) \mid S \in P^+[j,i]\right\}\right) & 2 \leq i \leq k-1 \\ \mathcal{V}_{v_k} \cup \left(\bigcup_{j=1}^{k-1} \left\{(\mathcal{G}_S, T\backslash e_{v_jv_{j+1}}) \mid S \in P[j,k]\right\}\right) & i = k \end{cases}.$$

To finish the proof we must show two things; first, that

$$\bigcup_{i=1}^{k} J_i = \{\mathcal{V}_{v_i}\}_{i=1}^{k} \cup \left\{(\cdot , T\backslash e_{v_iv_{i+1}})\right\}_{i=1}^{k-1} \cup \left\{(\cdot , T/e_{v_iv_{i+1}})\right\}_{i=1}^{k-1},$$

and second, that for each i there is a linear extension L_i which reverses the critical pairs in J_i.

It is clear that $\cup_{i=1}^{k} J_i$ contains all critical pairs in $\{\mathcal{V}_{v_i}\}_{i=1}^{k} \cup \left\{(\cdot , T/e_{v_i,v_{i+1}})\right\}_{i=1}^{k-1}$ since each J_i for $1 \leq i \leq k-1$ contains $\mathcal{V}_{v_i}$ and $(\cdot , T/e_{v_i,v_{i+1}})$, whereas J_k contains $\mathcal{V}_{v_k}$. It remains to verify that all critical pairs in $\left\{(\cdot , T\backslash e_{v_i,v_{i+1}})\right\}_{i=1}^{k-1}$ are contained in $\cup_{i=1}^{k} J_i$. To see that this is indeed the case, consider the edge $e_{v_jv_{j+1}}$ and some connected set S such that $(\mathcal{G}_S, T\backslash e_{v_jv_{j+1}}) \in \mathcal{E}$. Let i be the largest index such that $v_i \in S$. Note then that $S \in P^+[j,i]$ and thus $(\mathcal{G}_S, T\backslash e_{v_jv_{j+1}})$ is in J_i.

In order to complete we must show that, for each $i \in [k]$, there is a linear extension L_i reversing the critical pairs in J_i. By Theorem 9 it suffices to show that the elements in J_i do not induce an alternating cycle.

First consider J_1. As $\mathcal{G}_{v_1} \parallel T/e_{v_1v_2}$ and all single-vertex graphs are incomparable, the only comparability induced by J_1 is $\mathcal{G}_S \preceq T\backslash v_1$ where $S \in P^-[2]$. Hence there is no alternating cycle induced by J_1.

Now we consider J_k. As $T\backslash v_k \parallel \mathcal{G}_{v_k}$ and $T\backslash v_k \parallel \mathcal{G}_S$ for all $S \in \cup_{j=1}^{k-1} P[j,k]$, we have $T\backslash v_k$ does not participate in an alternating cycle induced by the critical pairs in J_k. We note as well that $\mathcal{G}_{v_k} \preceq T\backslash e_{v_iv_{i+1}}$ for all $i \in [k-1]$, and so $\mathcal{G}_{v_k}$ does not participate in any alternating cycles. Observing that among the remaining graphs the only potential comparabilities are between $\mathcal{G}_S$, where $S \in P[j,k]$ for some j, and $T\backslash e_{v_iv_{i+1}}$ for some i, we can define a bipartite graph which encapsulates the potential for alternating cycles. Let $a_j = \{\mathcal{G}_S \mid S \in P^-[k-j,k]\}$ for $1 \leq j \leq k-1$. The set $\{a_1, \ldots, a_{k-1}\}$ forms one side of the bipartition, and the set $\left\{T\backslash e_{v_{k-i}v_{k-i+1}}\right\}_{i=1}^{k-1}$ forms the other side of the bipartition. Then $\left\{a_j, T\backslash e_{v_{k-i}v_{k-i+1}}\right\}$ is an edge if and only if $j \leq i+1$. Notice that this graph entirely encapsulates the minor relation as $\mathcal{G}_S \preceq T\backslash e_{v_{k-i}v_{k-i+1}}$ if and only if $\mathcal{G}_S \in a_j$ for some $j \leq i+1$. Thus, since any altnerating cycle here is forced to increase along the indices, there cannot be an alternating cycle induced by J_k. See Figure 10a.

Lastly, we consider J_i for $2 \leq i \leq k-1$. Notice first that if $\mathcal{G}_S$ is a single-vertex graph in some critical pair in J_i then $S \in P^+[i] \cup P^-[i+1]$, and thus $\mathcal{G}_S \parallel T/e_{v_iv_{i+1}}$. Thus $T/e_{v_iv_{i+1}}$ participates in no alternating cycle within J_i. Furthermore, since for any $S \in P^-[i+1]$ we have that $\mathcal{G}_S \preceq T\backslash e_{v_jv_{j+1}}$ for $1 \leq j \leq i-1$ and $\mathcal{G}_S \preceq T\backslash v_i$, the single-vertex graph $\mathcal{G}_S$ does not participate in any alternating cycle. Similarly, as $T\backslash v_i$ is incomparable to the remaining graphs and $\mathcal{G}_{v_i}$ is comparable to all of the remaining graphs except $T\backslash v_i$, neither participates in an alternating cycle. Thus, if there is an alternating cycle induced by J_i it must be among $\mathcal{G}_S$, where $S \in \cup_{j=1}^{i} P^+[j,i]$, and $T\backslash e_{v_jv_{j+1}}$, where $1 \leq j \leq i-1$. As the only possible comparabilities are between single-vertex graphs and graphs of the form $T\backslash e_{v_jv_{j+1}}$, we will again define a bipartite graph which encapsulates the minor relations. To this end, let $a_j = \{\mathcal{G}_S \mid S \in P^{\pm}[i-j+1,i]\}$ for $1 \leq j \leq i$, and define the bipartition by $\{a_1, \ldots, a_i\}$ and $\left\{T\backslash e_{v_{i-\ell}v_{i-\ell+1}}\right\}_{\ell=1}^{i-1}$. Then $\left\{a_j, T\backslash e_{v_{i-\ell}v_{i-\ell+1}}\right\}$ is an edge if and only if $j \leq \ell$. Furthermore, $\mathcal{G}_S \preceq T\backslash e_{v_{i-\ell}v_{i-\ell+1}}$ if and only if $\mathcal{G}_S \in a_j$ for some $j \leq \ell$. Thus, since any alternating cycle would have to strictly increase amongst the indices, there are no alternating cycles induced by J_i. See Figure 10b. $\square$

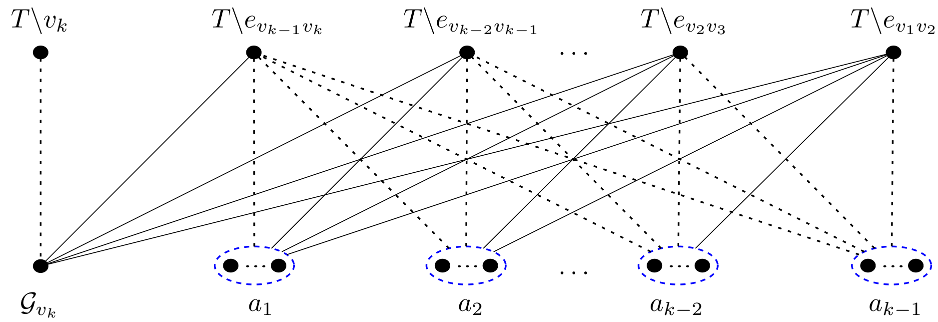

(a) The extension L_k has no alternating cycles.

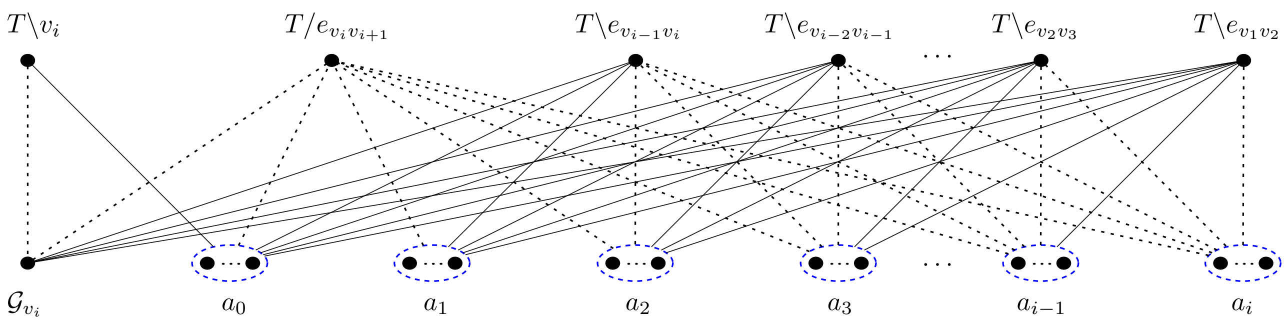

(b) The extension L_i has no alternating cycles. Here $a_0 = \{\mathcal{G}_S \mid S \in P^-[i+1]\}$.

Figure 10: Straight edges are comparabilities. Dashed edges are critical pairs.

Notice that Theorems 19 and 20 allow one to bound dimension of $\mathrm{MP}(T)$ with no notion of a realizer, or even of the poset itself, using only the topology of the tree. In fact, in the cases in which the lower and upper bounds match, the dimension itself can be computed in no time at all.

Corollary 21. *Let T be a tree with at least one edge. Then $|E(T)| + 1 \le \dim(\mathrm{MP}(T)) \le \lceil \frac{3}{2} \cdot |E(T)| \rceil$. Moreover, the bounds are tight.*

Proof. The lower bound follows from Theorem 11, as $|V(T)| = |E(T)| + 1$. For the upper bound, notice that Theorem 20 implies that

$$\dim(\mathrm{MP}(T)) \le |E(T)| + \frac{o_T}{2} \le |E(T)| + \frac{|V(T)|}{2} \le \left\lceil \frac{3}{2} \cdot |E(T)| \right\rceil.$$

Let P_n is the path with n vertices. Since P_n has exactly two vertices of odd degree we have $o_{P_n} = 1$, and hence $\dim(\mathrm{MP}(P_n)) = |E(T)| + 1$, showing that the lower bound is tight. Now let STAR_n be the star with n vertices. It is easy to verify that the minimal path decomposition has $\lceil |E(T)|/2 \rceil$ paths, none of which are special. Hence $\dim(\mathrm{MP}(\mathrm{STAR}_n)) \ge \lceil \frac{3}{2} |E(T)| \rceil$, showing that the upper bound is tight. □

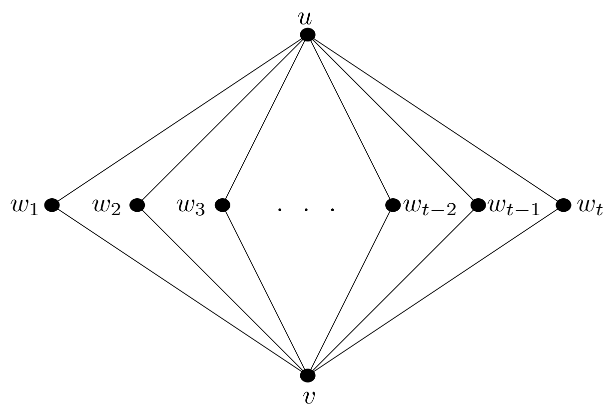

Figure 11: $K_{2,t}$.

As alluded to in Section 3.2, the following theorem characterizes the graphs whose minor posets have dimension equal to its breadth in the case that its breadth is the size the vertex set.

Theorem 22. *Let G be a graph with $|V(G)| = n$. Then $\dim(\mathrm{MP}(G)) = n$ if and only if G is a subgraph of P_n, the path on n vertices.*

Proof. If G is a subgraph of P_n then Theorem 14 and the proof of Corollary 21 imply that $\dim(\mathrm{MP}(G)) = n$. For the forward direction, notice that Proposition 11 yields $\dim(\mathrm{MP}(G)) \geq n$ for any G. Now suppose $\dim(\mathrm{MP}(G)) \leq n$. We wish to show that G is a subgraph of P_n. We may assume that G is connected, as if any component of G has dimension greater than the size of its vertex set, then Theorem 14 implies that $\dim(\mathrm{MP}(G)) \geq n + 1$. We may further assume that G has no cycles, as otherwise we could pick a cycle C, delete all edges with one end in C and the other in $G \setminus C$, and conclude from Theorem 15 and Theorem 14 that $\dim(\mathrm{MP}(G)) \geq n + 1$. So G is a tree. If G is not a path then it has at least one vertex of degree at least three, say v. The graph induced by $v \cup N(v)$, say G', is a star with at least four vertices. Thus, by the proof of Corollary 21, $\dim(\mathrm{MP}(G')) \geq |V(G')| + 1$. So if we delete all edges from $G \setminus G'$ to G' and apply Theorem 14 as before, we see that $\dim(\mathrm{MP}(G)) \geq n + 1$. Thus G is a subgraph of P_n, as desired. □

6. Subexponential Dimension — Excluding $K_{2,t}$-Minors

The parameter *treewidth* plays a significant role in the graph minors theorem (for the definition of treewidth and its basic properties, see [8]). Because of this, and the fact that a connected graph with at least two vertices has treewidth one if and only if it is a tree, it is tempting to think that Corollary 21 can be generalized to any class of graphs with bounded treewidth. For instance, it might be the case that there are constant c_1 and c_2, depending only on k, such that the minor posets of all graphs G with treewidth at most k have dimension at most $c_1 |E(G)|^{c_2}$. However, such a result is false for any $k \geq 2$, as the following example illustrates.

Example 1. Consider the complete bipartite graph $K_{2,t}$. It is clear that $K_{2,t}$ has no subdivision of K_4 since it only has two vertices with degree bigger than two. Furthermore, since the maximum degree of K_4 is three, having a subdivision of K_4 is equivalent to having a K_4-minor (see [8], Proposition 1.7.2.). Thus $K_{2,t}$ is series-parallel, or equivalently, $K_{2,t}$ has treewidth two. The aim of this example is to show that $\dim(\mathrm{MP}(K_{2,t}))$ is exponential in t, and hence exponential in $|E(K_{2,t})|$.

Let u and v be the two vertices of degree t in $K_{2,t}$ and let $w_1, w_2, \ldots, w_t$ be the t vertices of degree two, as in Figure 11. Let S, S' be any two distinct subsets of $\{w_1, \ldots, w_t\}$ of size $\lfloor \frac{t}{2} \rfloor$. Then we have $\mathcal{G}_{S \cup \{u,v\}} \preceq K_{2,t} \backslash E(v, S')$, whereas $(\mathcal{G}_{S \cup \{u,v\}}, K_{2,t} \backslash E(v, S)) \in \mathcal{E}$. Thus $\mathrm{MP}(K_{2,t})$ contains a standard example of size $\binom{t}{\lfloor \frac{t}{2} \rfloor}$. Therefore $\dim(\mathrm{MP}(K_{2,t})) \geq 2^{\lfloor t/2 \rfloor}$, which is exponential in $|E(K_{2,t})| = 2t$.

Example 1 suggests the following problem: characterize those graphs for which the dimension of its minor poset is subexponential is the size of its edge set. As we have seen, the dimension for trees and cycles is linear, whereas the dimension for graphs containing a $K_{2,t}$-minor, which of course includes large complete graphs, is exponential. With this in mind, we make the following conjecture.

Conjecture 23. *Let G be a connected graph. Then there is an increasing function $f(\cdot)$ and constant a C, depending only on t, such that if G is $K_{2,t}$-minor free then* $\dim(\mathrm{MP}(G)) \leq C\,|E(G)|^{f(t)}$.

Roughly this says that the only reason for the dimension of the minor poset to be large, with respect to the number of edges, is the presence of a large $K_{2,t}$-minor. Example 1 shows that the converse is true, assuming that t is sufficiently large with respect to the size of the edge set. Although we can not prove this conjecture in full generality, the following results will show that if G is $K_{2,t}$-minor free for $t \leq 4$ then there is some constant c such that $\dim(\mathrm{MP}(G)) \leq |E(G)|^c$. We start our discussion with outerplanar graphs, motivated by the fact that outerplanar graphs avoid $K_{2,3}$-minors. (Outerplanar graphs can be characterized as the class of graphs whose members are both $K_{2,3}$-minor free and K_4-minor free.)

Recall that an outerplanar graph is a graph that has an embedding in the plane in which all vertices are incident to the same face (without loss of generality, this is the infinite face). A graph is *maximally outerplanar* if it is impossible to add any edges and maintain outerplanarity. A *near triangulation* of the plane is a graph in which all but one face is bounded by exactly three edges. A graph is *Hamiltonian* if it has a cycle spanning all of the vertices. In the next theorem, we will need the following fact about maximally outerplanar graphs, which is easily verified.

Fact 24. All maximally outerplanar graphs are Hamiltonian, 2-connected, near-triangulations of the plane.

In the course of proving the following theorem, in order to understand the structure of the graph it is necessary to pass to the *weak plane dual* of a planar embedding of the outerplanar graph. Roughly speaking, this graph is formed by associating to each finite faces a vertex and having two vertices adjacent if and only if their faces have a common edge. For a more formal definition, we refer interested readers to [8].

Theorem 25. *If G is a connected outerplanar graph, then the dimension of the minor poset is polynomial in $|E(G)|$; specifically,* $\dim(\mathrm{MP}(G)) = \mathcal{O}\Big(|E(G)|^4\Big)$.

Proof. We will prove that $\dim(\mathrm{MP}(G)) = \mathcal{O}\big(n^4\big)$, where $|V(G)| = n$. This is sufficient since $n-1 \leq |E(G)| \leq 2n-3$ for connected outerplanar graphs. We may assume that G is maximally outerplanar since otherwise we can add edges and be sure not to decrease the dimension of $\mathrm{MP}(G)$ (for any additional edge e, $G \preceq G+e$). Therefore, by Fact 24, we may assume that G is a 2-connected, Hamiltonian, near-triangulation of the plane with $|E(G)| = 2n-3$.

The critical pairs in $\mathcal{V}$ can be reversed with n linear extensions and Remark 1 implies that the critical pairs in $\mathcal{C}$ can be reversed with $|E(G)| = 2n-3$ linear extensions. Thus, we can reverse all but the critical pairs in $\mathcal{E}$ with $\mathcal{O}(n)$ extensions. So it is sufficient to prove that all of the critical pairs in $\mathcal{E}$ can be reversed with $\mathcal{O}\big(n^4\big)$ extensions. Recall that each critical pair in $\mathcal{E}$ is of the form $(\mathcal{G}_{U\cup V}, G\backslash E(U,V))$ where $G[U]$ and $G[V]$ are connected, $U \cap V = \varnothing$, and $E(U,V) \neq \varnothing$. Therefore, it is sufficient to show that there are at most $\mathcal{O}\big(n^4\big)$ valid choices for $E(U,V)$, as all such critical pairs can be reversed in a single extension by Remark 1.

Consider an embedding of G in the plane such that every vertex is incident with the infinite face. By Fact 24 the boundary of the infinite face is a Hamiltonian cycle, say D. Let $\mathfrak{C}$ be the set of chords of D and let F be the bounded face of $G \setminus \mathfrak{C}$; that is, F is the complement of the infinite face of G.

Now let U and V be disjoint sets of vertices so that $G[U]$ and $G[V]$ are connected and $E(U,V)$ is non-empty. We will show that there are at most $\binom{n-2}{2}+1$ possible choices for subsets of edges in $E(U,V)\cap\mathfrak{C}$. Further, we will show that there are at most two edges in $E(U,V)\cap D$, and thus, since $|D| = n$, there are at most $\binom{n}{2}+n$ choices for $E(U,V)\cap D$. In total, since $E(G) = \mathfrak{C}\cup E(D)$, we will have that there are most $\left(\binom{n}{2}+n\right)\left(\binom{n-2}{2}+1\right) = \mathcal{O}\big(n^4\big)$ choices for $E(U,V)$, which will complete the proof.

In order to analyze the edges in $E(U,V)\cap\mathfrak{C}$, we need to consider the weak plane dual H of the embedding of G. Since the infinite face is ignored, every edge of H corresponds to an edge in $\mathfrak{C}$. Furthermore, since G is maximally outerplanar, H is a tree with maximum degree three. Let $H_{U,V}$ be the subgraph of H containing all edges corresponding to edges of $E(U,V)\cap\mathfrak{C}$ and their ends. We now will proceed to show that the $H_{U,V}$ is a path.

First we show that there are no vertices of degree three in $H_{U,V}$. Suppose not, and let $f \in V(H_{U,V})$ be such that $\deg_{H_{U,V}}(f) = 3$. Then $\deg_H(f) = 3$ as well, and so f corresponds to a triangular face of the embedding of G such that all three edges along the boundary have one edge in U and one in V. This yields a two-coloring of a triangle, a contradiction.

Now consider two edges in $H_{U,V}$ which correspond to the edges $\{u_1, v_1\}$ and $\{u_2, v_2\}$ in G (with $u_1, u_2 \in U$ and $v_1, v_2 \in V$). If $\{u_1, u_2, v_1, v_2\}$ are all distinct then proceeding around D we find that the elements of U and V either alternate, as in Figure 12a, or are consecutive, as in Figure 12b. If they are not distinct, we may assume without loss of generality that $u_1 = u_2$ as in Figure 12c.

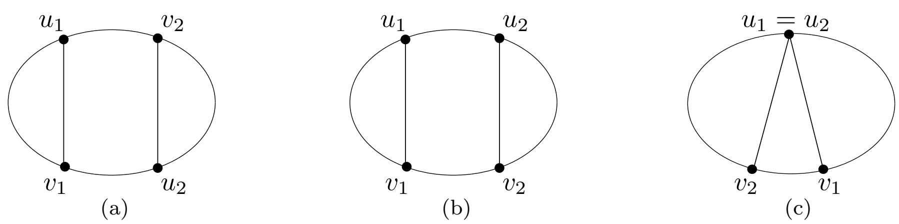

Figure 12: Arrangement of $\{u_1, u_2, v_1, v_2\}$ on D

Suppose in the first two cases that there is some chord $\{w_1, w_2\} \in \mathfrak{C} \backslash E(U, V)$ which topologically separates $\{u_1, v_1\}$ and $\{u_2, v_2\}$. Without loss of generality we may assume that $w_1, w_2 \notin V$. But then, connecting w_1 and w_2 by a second edge through the infinite face, we find a planar embedding in which the vertices in V disjoint from and on both sides of a separating cycle. This contradicts that $G[V]$ is connected, thus any separating chord must be in $E(U, V)$.

Similarly, in the third case, suppose there is a chord $\{u_1, w\}$ which topologically separates $\{u_1, v_1\}$ and $\{u_2, v_2\}$ with $w \notin V$. By adding another edge $\{u_1, w\}$ through the infinite face there is a separating cycle that is disjoint from V which separates V, contradicting that $G[V]$ is connected.

Thus given any two edges in $E(U, V) \cap \mathfrak{C}$, any chord which topologically separates them must also be in $E(U, V) \cap \mathfrak{C}$, and thus $H_{U,V}$ is connected. Since $H_{U,V}$ is a connected subgraph of a tree with no vertex of degree more than two, it follows that $H_{U,V}$ is a path. Since any path in a tree is uniquely defined by its endpoints, and since there are $(n-3)$ edges in H, there are $\binom{n-2}{2}$ such paths. So, including the empty subgraph, there are at most $\binom{n-2}{2} + 1$ choices for $E(U, V) \cap \mathfrak{C}$.

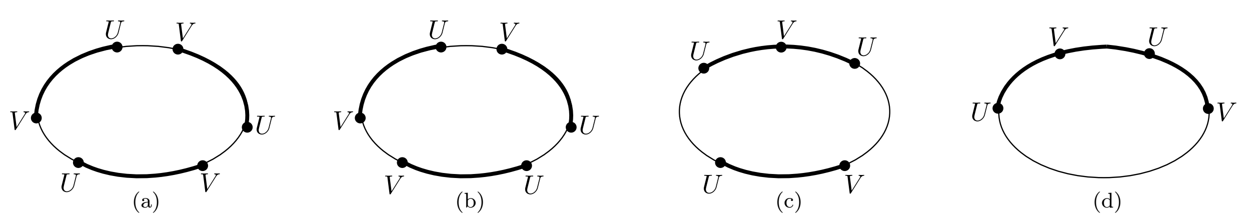

Figure 13: The four configurations of edges in $E(U, V) \cap D$ up to symmetry.

Next, we claim that $|E(U, V) \cap D| \leq 2$. Suppose not. Then there are at least three edges in $E(U, V) \cap D$. There are four cases, up to symmetry, for any three of these edges. These cases are depicted in Figure 13. In every case there are distinct vertices $u_1, u_2 \in U$ and $v_1, v_2 \in V$ such that, when one proceeds clockwise around D, one encounters u_1, then v_1, then u_2, then v_2. Since $G[U]$ is connected there is a path from u_1 to u_2 in G. Further, since these four vertices alternate along D, this path must use a chord which topologically separates v_1 and v_2. Then, by connecting u_1 and u_2 by an edge (possibly a second $\{u_1, u_2\}$ edge) through the infinite face, v_1 and v_2 are on opposite sides of a separating cycle in which all the vertices belong to U. Thus $G[V]$ is disconnected, a contradiction. Hence there are at most $\binom{n}{2} + n$ possible choices for $E(U, V) \cap D$, completing the proof. $\square$

Outerplanar graphs avoid $K_{2,3}$-minors. In this paper, we are more interested in the class of graphs that only include $K_{2,t}$-minors, for some t. For this we will need to employ a recent theorem of Dieng and Gavoille [7].

Theorem 26. [7] *Every 2-connected $K_{2,r+2}$-minor free graph has r vertices whose removal leaves the graph outerplanar, for each $r \in \{0, 1, 2\}$.*

With this theorem in hand, we can prove Theorem 27.

Theorem 27. *If G be a connected graph with no $K_{2,3}$-minor, then the dimension of the minor poset of G is polynomial in $|E(G)|$.*

In particular, a simplification of the proof of Theorem 30 shows $\dim(\mathrm{MP}(G)) = \mathcal{O}\left(|E(G)|^{60}\right)$. However, we have omitted the proof since Theorem 27 can be viewed as a corollary of Theorem 30. Before proving Theorem 30 we prove the following bound on the number of subtrees of a tree with limited branching.

Lemma 28. *Let T be a tree with n vertices and at most three leaves. Then there are $\mathcal{O}(n^7)$ ways of choosing three disjoint subtrees of T.*

Proof. As T is a tree, it must have at least two leaves. If T has exactly two leaves then it is a path, and hence so is each subtree. We can count the number of disjoint subpaths exactly by specifying the six vertices which are the endpoints of the paths. Clearly there are at most n^6 ways to do this.

Now suppose T has three leaves, v_1, v_2, v_3. As such, there is precisely one vertex, say v, whose degree is greater than two; in fact, $\deg(v) = 3$. Let n_i be the number of vertices on the path in T from v to v_i, for $i = 1, 2, 3$. If one subtree S contains v, then there are $n_1 n_2 n_3 = \mathcal{O}(n^3)$ ways to build this tree. The other two subtrees have three choices for which component of $T \setminus S$ they reside, at most n_i choices in length, and at most n_i choices of position in that component, for whichever i it chooses. As each $n_i \leq n$, there are at most $(3n_i^2)^2 = \mathcal{O}(n^4)$ choices for the other two subtrees. So in total there are $\mathcal{O}(n^7)$ choices of three subtrees. If none of the subtrees contain v, then, by similar logic, there are at most $((3n_i^2)^3) = \mathcal{O}(n^6)$ choices of three subtrees. $\square$

Define $\mathfrak{B}(G)$, the *block graph* of G, to be the bipartite graph with bipartition

$$(\{\tau_B \mid B \text{ is a block of } G\}, \{\tau_v \mid v \text{ is a cutvertex in } G\}).$$

The block graph has an edge $\{\tau_B, \tau_v\}$ if and only if the cutvertex v is a vertex in the block B. It is well known that $\mathfrak{B}(G)$ is a tree. Clearly the leaves of $\mathfrak{B}(G)$ are blocks and not cutvertices. Hence, we call the leaves of $\mathfrak{B}(G)$ *leaf blocks*. The next lemma concerns the block graph and will be used in the proof of Theorem 30.

Lemma 29. *Let $B_1, B_2, \ldots, B_k$ be the blocks of a graph G with $|V(G)| = n$. Let $|V(B_i)| = n_i$ for all $i \in [k]$. Then $\sum_{i=1}^{k} n_i \leq 2n - 1$.*

Proof. Let C be the set of cutvertices of G. In the quantity $\sum_{i=1}^{k} n_i$, each non-cutvertex is counted once, whereas each cutvertex $c \in C$ is counted $\deg_{\mathfrak{B}(G)}(\tau_c)$ times. Let $|V(\mathfrak{B}(G))| = n'$, C_1 be the set of leaves $\mathfrak{B}(G)$, C_2 be the set of vertices of degree two in $\mathfrak{B}(G)$, and $C_3 = V(\mathfrak{B}(G)) \backslash (C_1 \cup C_2)$. Since $\mathfrak{B}(G)$ has a bipartition corresponding to cutvertices and non-cutvertices, we have the following:

$$\begin{aligned}
\sum_{c \in C} \deg_{\mathfrak{B}(G)}(\tau_c) &= \frac{1}{2} \cdot \left(\sum_{\tau_v \in V(\mathfrak{B}(G))} \deg_{\mathfrak{B}(G)}(\tau_v) \right) \\
&= \frac{1}{2} \left(\sum_{v \in C_3} \deg_{\mathfrak{B}(G)}(\tau_v) + \sum_{v \in C_2} \deg_{\mathfrak{B}(G)}(\tau_v) + \sum_{v \in C_1} \deg_{\mathfrak{B}(G)}(\tau_v) \right) \\
&= \frac{1}{2} \left((|C_1| + 2|C_3| - 2) + 2|C_2| + |C_1| \right) \\
&= |C_3| + |C_2| + |C_1| - 1 \\
&= n' - 1,
\end{aligned}$$

where we have used the fact that $|C_1| = (\sum_{v\in C_3} \deg_{\mathfrak{B}(G)}(\tau_v) - 2) + 2$, an equation that holds for all trees. Now let N be the set of non-cutvertices in G and let B be the set of blocks. Then

$$\sum_{i=1}^{k} n_i = \left(\sum_{c\in C} \deg_{\mathfrak{B}(G)}(\tau_c)\right) + |N| = n' - 1 + |N| = |C| + |B| - 1 + |N| \leq |C| + 2\,|N| - 1.$$

But it's clear that $|C| + 2\,|N| \leq 2n$, which finishes the proof. ☐

Several times in the following proof we will forbid certain substructures by showing that they result in a $K_{2,4}$-minor. In particular, it is often the case that we will find distinct vertices x_1, x_2 and y_1, y_2, y_3, y_4 such that x_1 and x_2 are internal vertices of two internally disjoint subtrees, each of whose leaves include $\{y_1, y_2, y_3, y_4\}$. Clearly such a case results in a $K_{2,4}$-minor, and we will say that a $K_{2,4}$-minor can be constructed from $\{x_1, x_2\}$ and $\{y_1, y_2, y_3, y_4\}$ rather than explicitly designating the minor.

Theorem 30. *Let G be a connected graph with no $K_{2,4}$-minor. Then the dimension of the minor poset of G is polynomial in $|E(G)|$.*

Proof. Let $|V(G)| = n$ and $|E(G)| = m$. We note that by Theorem 26 and the fact that outerplanar graphs has at most twice as many edges as vertices, every two connected block of G has at most four times as many edges as vertices. This inequality is clearly true for those blocks that are not 2-connected as well — the bridges. Thus Lemma 29 implies that $m \leq 8n$. Hence it suffices to prove that the dimension of $\mathrm{MP}(G)$ is polynomial in n. In particular, we will prove that $\dim(\mathrm{MP}(G)) = \mathcal{O}(n^{1614})$.

Suppose that we can show that if G' is 2-connected, then $\dim(\mathrm{MP}(G')) \leq c\,|V(G')|^{\alpha}$ for some fixed $c > 0$ and $\alpha \geq 1$. Now consider the case where G is not 2-connected. Then G can be decomposed into blocks B_i where each block is either 2-connected or is a bridge. Then, by Theorem 14 (and Theorem 19 to show that $\dim(\mathrm{MP}(K_2)) = 2$),

$$\dim(\mathrm{MP}(G)) \leq \sum_{i=1} \dim(\mathrm{MP}(B_i)) \leq \sum_{i=1} \max\{2, c\,|V(B_i)|^{\alpha}\} \leq (c+1)\left(\sum_{i=1}^{k} |V(B_i)|\right)^{\alpha} \leq (c+1)(2n)^{\alpha}.$$

Hence, we may assume that G is 2-connected, and as such, Theorem 26 guarantees the existence of two vertices s and t such that $G' = G \setminus \{s, t\}$ is outerplanar.

Like the proof of Theorem 25, we will bound the dimension of $\mathrm{MP}(G)$ by finding an upper bound on the number of subsets $C \subseteq E(G)$ for which there exist disjoint sets U and V such that $G[U]$ and $G[V]$ are connected and $C = E(U, V)$. By Theorem 25, if $(U \cup V) \cap \{s, t\} = \varnothing$, then there are $\mathcal{O}(n^4)$ such sets. So we will only be concerned with the case that at least one of s and t is in $U \cup V$. Without loss of generality, $s \in U$.

Since G is not necessarily 3-connected, it may be the case that G' is disconnected. Suppose G' has k components. Each component must have a vertex adjacent to s and a vertex adjacent to t, as otherwise either s or t is a cutvertex of G, contrary to the fact that G is 2-connected. Therefore G has a $K_{2,k}$-minor. Hence, as G is $K_{2,4}$-minor free, $k \leq 3$.

At this point the proof will proceed as follows: for each component Q of G' we will partition the edges of Q and the edges incident with vertices in Q into five sets, say $E_1, E_2, \ldots, E_5$. For each set E_i we will show that there is a constant c_{E_i} such that the number of subsets of edges from E_i that could be in $E(U, V)$ is bounded by $\mathcal{O}(|V(Q)|^{c_{E_i}})$. Once we have shown this, the number of possible sets for $E(U, V)$ is $\mathcal{O}(n^{c_{E_1}+c_{E_2}+\cdots+c_{E_5}})$. Furthermore, we will show that $c_{E_1} + c_{E_2} + \cdots + c_{E_5} = 1614$.

Fix some component Q of G'. Recall from above that $\mathfrak{B}(Q)$, the block graph of Q, is a tree. Let T_s be the minimal subtree of $\mathfrak{B}(Q)$ such that $\tau_B \in V(T_s)$ whenever B contains a non-cutvertex neighbor of s, and such that $\tau_v \in V(T_s)$ whenever v is a cutvertex of G adjacent to s. See Figure 14. We can make the following observations about T_s.

(1) If there are at least four vertices in T_s, say $\tau_{B_1}, \tau_{B_2}, \tau_{B_3}, \tau_{B_4}$, such that B_i has a non-cutvertex u_i that is adjacent to s, then we can construct a $K_{2,4}$-minor in G; the 2-side is s and the vertex obtained by contracting $V(Q) \setminus \{u_1, u_2, u_3, u_4\}$, and the 4-side is $\{u_1, u_2, u_3, u_4\}$. (Note that $Q \setminus \{u_1, u_2, u_3, u_4\}$ is connected since each block is 2-connected.) Hence, T_s has at most three such vertices. See Figure 15a.

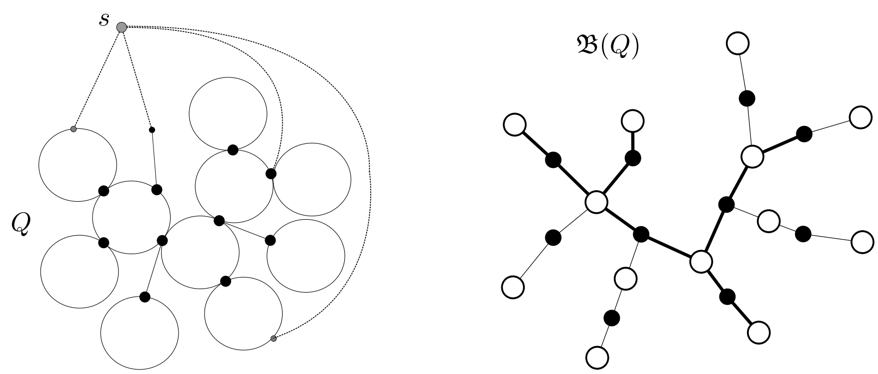

(a) (b)

Figure 14: Figure 14a is a component Q of G' with s and four edges incident to both s and Q. Figure 14b is $\mathfrak{B}(Q)$ with edges of T_s thickened.

(2) If there are at least two vertices in T_s, say τ_{B_1}, τ_{B_2}, such that B_i has at least two non-cutvertices $u_{i,1}$ and $u_{i,2}$, both adjacent to s, then we can construct a $K_{2,4}$-minor in G; the 2-side can be constructed from s and a cutvertex v between B_1 and B_2, and the 4-side is $\{u_{1,1}, u_{1,2}, u_{2,1}, u_{2,2}\}$. Hence, there is at most one such vertex of T_s. See Figure 15b.

(3) Let T'_s be the subtree of T_s formed by removing leaves that correspond to cutvertices of Q. If T'_s has at least four leaves, say $\tau_{B_1}, \tau_{B_2}, \tau_{B_3}, \tau_{B_4}$, then there are four vertices u_1, u_2, u_3, u_4 such that u_i is a witness for $\tau_{B_i} \in V(T'_s)$. That is, either u_i is a cutvertex in block B_i and adjacent to s and τ_{u_i} is a leaf in T_s or u_i is non-cutvertex in B_i and is adjacent to s. Thus we can construct at $K_{2,4}$-minor in G; the 2-side consists of s and the vertex obtained by contracting $V(Q) \setminus \{u_1, u_2, u_3, u_4\}$, and the 4-side consists of $\{u_1, u_2, u_3, u_4\}$. Thus T'_s has at most three leaves which we will denote $\tau_{B_1}, \tau_{B_2}, \tau_{B_3}$. See Figure 15c.

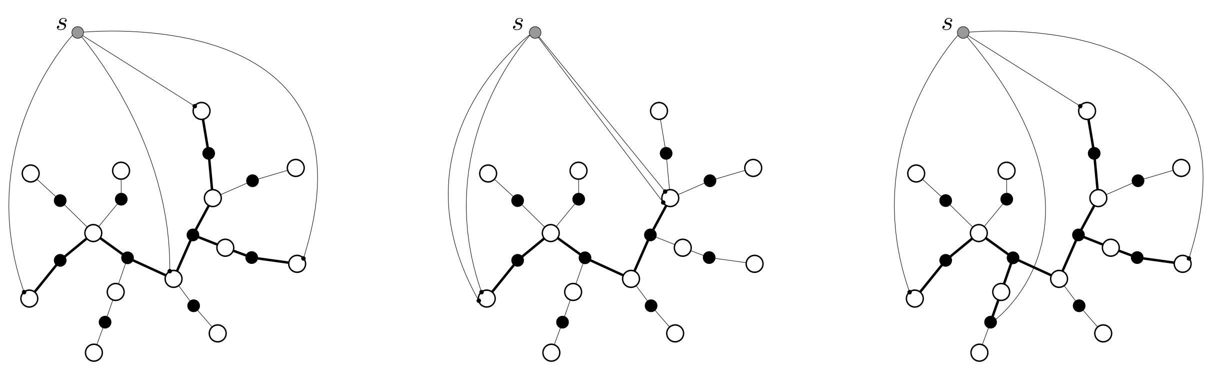

(a) (b) (c)

Figure 15

Let B be a block of Q with the maximum number of non-cutvertices adjacent to s. By observations (1) and (2), there are at most three other non-cutvertices in Q that are adjacent to s. Let E_1 be the set of edges incident with s and a cutvertex not in $V(B) \cup V(B_1) \cup V(B_2) \cup V(B_3)$. Let E_2 be the set of edges incident with s and a non-cutvertex not in B, and let E_3 be the set of edges incident with s and a vertex in B, B_1, B_2, or B_3. Notice that the sets E_1, E_2, E_3 cover all of the edges incident with s and Q, and that $|E_2| \leq 3$.

We first bound the number of subsets of E_1 that can be in $E(U,V)$. As we are looking for an upper bound, we may assume that s is adjacent to each cutvertex represented in T'_s. Consider the subgraph of T'_s induced by $\{\tau_v \cup N(\tau_v) \mid v \in V\}$. If this subgraph is disconnected then, since $G[V]$ is connected, we have $t \in V$ and for each component there is a corresponding vertex in $V(Q)$ that is adjacent to t. Thus, if this subgraph has at least four components, G has a $K_{2,4}$-minor; the 2-side is s and t, and the 4-side consists of the contracted vertices representing each component. Therefore, we can find an upper bound on the number of subsets of E_1 that can be in $E(U,V)$ by counting the number of ways we can choose at most three disjoint subtrees of T'_s. By observation (3), T'_s has at most three leaves. Furthermore, $|V(T_s)| \leq 2\,|V(Q)|$. Hence, by Lemma 28, there are $\mathcal{O}\big((2\,|V(Q)|)^7\big) = \mathcal{O}\Big(|V(Q)|^7\Big)$ subsets of E_1 that can be in $E(U,V)$.

We will now bound the number of subsets of $E_3 \cap E(V(B),\{s\})$ that can be in $E(U,V)$. Since B is 2-connected and outerplanar, it has a Hamiltonian cycle $\mathcal{D}$. Fix an outerplanar embedding of B in which the infinite face is bounded by $\mathcal{D}$. We will call two neighbors of s consecutive with respect to this embedding if there are no other neighbors of s between them in $\mathcal{D}$ (either clockwise or counter-clockwise). We then define $V_1,\ldots,V_\ell$ as the maximal sets of consecutive neighbors of s in V.

We note that $\ell \leq 3$. Suppose not, then there are four vertices u_1,u_2,u_3,u_4 which are neighbors of s and separate V_i and V_j along $\mathcal{D}$ for all $i \neq j$. But then since $G[V]$ is connected, there is a $K_{2,4}$-minor formed by $\{u_1,u_2,u_3,u_4\}$ together with s and the vertex formed by contracting the vertices in V. Thus the number of potential subsets for $E_3 \cap E(V(B),\{s\})$ is bounded by the number of ways of breaking all the neighbors of s in B into at most 6 intervals. There are clearly at most $2\,|V(B)|^6$ ways of doing this, which is $\mathcal{O}\Big(|V(Q)|^6\Big)$. Notice that this argument also holds for those edges in $E_3 \cap E(V(B_i),\{s\})$ for each $i \in \{1,2,3\}$.

Returning to the set $E_1 \cup E_2 \cup E_3$, we find that there are at most $\mathcal{O}\Big(|V(Q)|^7\Big) \cdot 2^3 \cdot \mathcal{O}\Big(|V(Q)|^6\Big)^4 = \mathcal{O}\Big(|V(Q)|^{31}\Big)$ ways of choosing subsets of edges incident with s and Q which can be in $E(U,V)$. Repeating this argument for edges incident with both t and Q yields that there are at most $\mathcal{O}\Big(|V(Q)|^{62}\Big)$ ways of choosing subsets of edges incident with either s or t and a vertex in Q that can be in $E(U,V)$.

We now turn to edges with both ends in Q. As each block of Q with at least three vertices is 2-connected and outerplanar, each has a Hamiltonian cycle. Fix this Hamilton cycle in each block. Let E_4 be the set of edges in Q that are in the fixed Hamiltonian cycle of some block together with all bridges of Q. Let E_5 be the remaining edges — the chords of Hamilton cycles.

We begin with E_4. Consider an arbitrary block B' of $\mathfrak{B}(Q)$. If B' is a bridge then it has one edge, and so there are a constant number of subsets of $E_4 \cap E(B')$ which are in $E(U,V)$. Otherwise, B' has a Hamiltonian cycle $\mathcal{D}$. The edges of $\mathcal{D}$ in $E(U,V)$ are those with one end in U and the other in V. We first consider the case where every vertex is in either U or V. This naturally partitions the vertices of B' into maximal intervals of U-vertices, $U_1,\ldots,U_k$, and maximal intervals of V-vertices, $V_1,\ldots,V_k$. We first note that there are at most three intervals of U-vertices such that there is a path from a vertex in the interval to s that is external to B', as we can construct a $K_{2,4}$-minor in G. Namely, since the outerplanarity of Q implies that all the paths are disjoint, s and the vertex formed by contracting V form the 2-side and the 4-side is formed by the four intervals. Now suppose there are four pairs (V_i,V_j) such that each pair is connected by a chord. Consider the planar graph formed by removing all other chords and contracting the intervals of U- and V-vertices. This graph has $2k$ vertices and $2k+4$ edges, and so by Euler's formula has six faces, one of which is the infinite face. Now at most one finite face does not border the infinite face. Thus there are at least four finite faces bordering the infinite face and further, the boundary of each of these faces contains a distinct vertex representing an interval of U-vertices. As these U-vertices must be connected (and can not be connected within B'), there are disjoint paths from these intervals of U-vertices to s, a contradiction. Thus there are at most three pairs (V_i,V_j) such that there is a chord between them. In fact, by a similar argument if there are three chords they must form a triangle to create a finite face that does not border the infinite face, thus there are either at most two chords or three chords which form a triangle. Now, as above, there are at most three intervals of V-vertices which have a path external to B', and thus, since the V-vertices are connected there are at most five intervals of V-vertices. Hence, there are at most $2\,|V(B')|^{10}$ subsets of $E_4 \cap E(B)$ which can be in $E(U,V)$ if every vertex in B' is in either U or V.

Now suppose that not every vertex in B' belongs to either U or V. We label the unlabeled vertices according to which labeled vertex is nearest to it in a counterclockwise traversal of $\mathcal{D}$; if a U-vertex is closest, label the unlabeled vertex U', and otherwise label it V'. This new labeling maintains the property that $G[U \cup U']$ and $G[V \cup V']$ are connected and further $E(U, V) \subseteq E(U \cup U', V \cup V')$. However, by the above work $E(U \cup U', V \cup V') \cap E(B') \leq 10$. Thus, overall there are at most $2^{11} |V(B')|^{10}$ subsets of $E_4 \cap E(B')$ which can be in $E(U, V)$.

We now consider how many blocks B' can have non-empty intersection with $E_4 \cap E(U, V)$. Suppose two distinct blocks B'_1 and B'_2 have an edge in $E_4 \cap E(U, V)$, say e_1 and e_2, respectively. That is, one end of e_i is in U and the other in V for $i \in 1, 2$. Because $\mathfrak{B}(Q)$ is a tree and Q is outerplanar, only one set of ends, either the U-set or the V-set, is connected in Q; the other must use either s or t. Thus if $t \in U$, there are at most four blocks with non-empty intersection with $E_4 \cap E(U, V)$. Suppose then that $t \in V$. Consider the complete graph whose vertices are blocks with non-empty intersection with $E_4 \cap E(U, V)$. Color each edge with the color U if the U-vertices in the blocks represented by the endpoints are not connected in Q, otherwise color the edge with color V (if an edge can receive both colors, then choose one arbitrarily). Notice that if this graph has a monochromatic K_4 then we can construct a $K_{2,4}$-minor in G, since both $G[U]$ and $G[V]$ are connected. See Figure 16. Thus there are at most $R(4, 4) - 1 = 17$ vertices [16]. (Note: $R(4, 4)$ is the Ramsey number for K_4 — the minimum number of vertices such that any two coloring of the edges of the complete graph on that number of vertices has a monochromatic K_4). Now if the block sizes are $n_1, n_2, \ldots, n_k$, this implies that there are

$$\sum_{S \in \binom{[k]}{17}} \prod_{i \in S} 2^{11} n_i^{10} \leq 2^{187} \binom{k}{17} \left(\frac{2n}{k}\right)^{170} = \mathcal{O}\big(n^{170}\big)$$

choices of $E_4 \cap E(U, V)$.

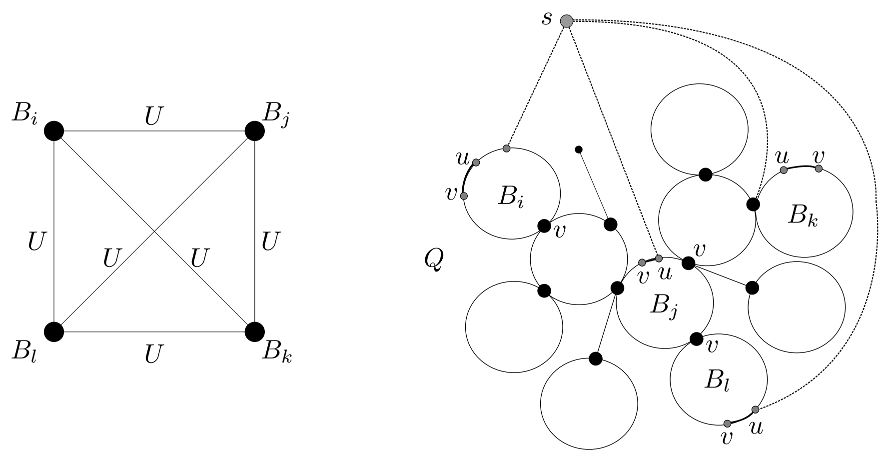

Figure 16: A monochromatic K_4, here with color U, allows the construction for a $K_{2,4}$-minor in G. The vertices labeled u are in U, the vertices labeled v are in V.

Finally, we consider E_5. Consider a block B'' and let $U_1, \ldots, U_k$ be the maximal sets of vertices in U such that $G'[U_i \cap B'']$ is connected. That is, the U_i are the maximal sets of U-vertices that are connected within B''. Define $V_1, \ldots, V_\ell$ analogously. Now, we may assume without loss of generality that $E(U_i, V) \neq \varnothing$ for each i. Thus, since each U_i must be connected through s, we have that $k \leq 3$ as above. Similarly, either $t \in V$ or $\ell = 1$. But then $\ell \leq 3$. Thus there are at most nine pairs (U_i, V_j). Now fix some maximal outerplanar graph B''' such that B'' is a subgraph of B'''. Since B'' is a subgraph of B''', we have that $E_{B''}(U_i, V_j) = E_{B'''}(U_i, V_j) \cap E(B'')$ and thus, by the arguments of Theorem 25, there are at most $|V(B'')|^2$ choices for $E_{B''}(U_i, V_j)$. Now since there are at most nine such pairs, there are at most $|V(B'')|^{18}$ choices for $E_5 \cap E(B'')$. Applying the same argument as in the case of E_4 there are at most 17 blocks which

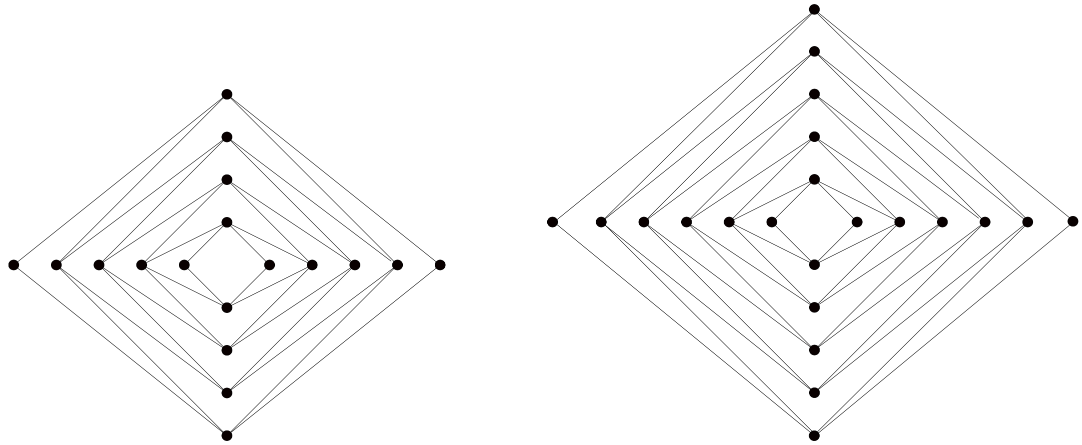

(a) G_4 (b) G_5

Figure 17: No graph in the family $\{G_n\}$ has a $K_{2,5}$-minor. Furthermore, there is no constant c such that c vertices can be removed from any member of $\{G_n\}$ to yield an outerplanar graph.

have a non-empty intersection with $E_5 \cap E(U,V)$, and thus there are at most $\mathcal{O}(n^{306})$ choices of subset for $E_5 \cap E(U,V)$.

All edges incident with a vertex in Q have been covered by $E_1 \cup E_2 \cup E_3 \cup E_4 \cup E_5$. Thus, the number of subsets of such edges that can be in $E(U,V)$ is $\mathcal{O}\Big(|V(Q)|^{62}\Big)\,\mathcal{O}\Big(|V(G)|^{170}\Big)\,\mathcal{O}\Big(|V(G)|^{306}\Big) = \mathcal{O}\Big(|V(Q)|^{538}\Big)$. Cubing this to account for the at most three components of G', multiplying by 2 to account for e_{st} (if it exists), and adding $\mathcal{O}(n^4)$ for the case where $\{s,t\} \cap (U \cup V) = \varnothing$, we find that there are $\mathcal{O}(n^{1614})$ choices for $E(U,V)$. Therefore, using Remark 1, we see that the type $\mathcal{E}$ critical pairs in $\mathrm{MP}(G)$ can be reversed with $\mathcal{O}(n^{1614})$ linear extensions. The type $\mathcal{V}$ critical pairs require n linear extensions, and since $m \leq 8n$, the type $\mathcal{C}$ critical pairs require at most $8n$ linear extensions. In total, we can reverse all critical pairs of $\mathrm{MP}(G)$ with $\mathcal{O}(n^{1614})$ linear extensions, as desired. □

Remark 4. The proof of Theorem 30 generalizes to $K_{2,t}$-minor free graphs as long as there exists a function $f(t)$, not depending on the size of the graph, such that, for all graphs without a $K_{2,t}$-minor, we can delete $f(t)$ vertices and obtain an outerplanar graph. Unfortunately this theorem is false for $t \geq 5$, as illustrated by the family of graphs $\{G_n\}_{n=4}^{\infty}$ in Figure 17. In [7], G_4 was used to show that Theorem 26 is false for $r = 3$.

The aim of the work in this section has been to characterize graphs without $K_{2,t}$-minors as those graphs whose minor poset has dimension subexponential in the size of its edge set. Remark 4 shows that a direct generalization of the techniques in the proof of Theorem 30 is not possible. However, in [9], Ding provides a characterization of graphs that exclude $K_{2,t}$-minors. Specifically, he shows that, for every natural number t, every sufficiently large 3-connected graph with minimum degree at least six, every 4-connected graph with a vertex of sufficiently large degree, and every sufficiently large 5-connected graph must have a $K_{2,t}$-minor. This characterization would seem to indicate that any counterexample to Conjecture 23 must occur in a graph with low connectivity.

7. Multigraphs

The first six sections of this paper has been focused on simple graphs. In order to obtain deeper results, it may be beneficial to consider a more general, multigraph model. Two such models are described in this section.

Given a ground set $[n]$, define a *pseudograph* to be a pair (V,E). The elements of the vertex set V are the equivalence classes of some equivalence relation on $[n]$, and $E : V \times V \mapsto \mathbb{Z}_+$ is a function that defines the number of edges between $v_1, v_2 \in V$. When $v_1 \neq v_2$, $E(v_1,v_2)$ is the number of parallel edges between v_1 and v_2, and when $v_1 = v_2$, $E(v_1,v_2)$ is the number of loops on v_1. For the sake of notation, we will write

v^l when that $E(v,v)=l$. We will also write $G[G_1]$ to mean the subgraph of G induced by those vertices whose equivalence class is contained in the ground set of G_1.

The minor operations are as follows. Deleting a vertex v removes the equivalence class v from both the ground set (and hence from V as well). Edge deletion corresponds to choosing $v_1, v_2 \in V$ such that $v_1 \neq v_2$ and $E(v_1, v_2) \geq 1$ and reducing $E(v_1, v_2)$ by one. For two vertices $v_1^{l_1}, v_2^{l_2} \in V$ such that $E(v_1, v_2) = e \geq 1$, the contraction of an edge incident to v_1 and v_2 results in a new vertex $v^{l_1+l_2+e-1}$ in which an element of the ground set x is in the equivalence class represented by v if and only if $x \in v_1 \cup v_2$.

Given this model, we have the following result, very much akin to Lemma 10. In fact, the proofs are analgous, and as such we omit the proof of the next theorem.

Theorem 31. *Let G be a pseudograph. The ordered pair (G_1, G_2) is a critical pair of* $\mathrm{MP}(G)$ *if and only if G_1 and G_2 satisfy one of the following:*

- *G_1 is a single-vertex pseudograph with no loops and $V(G_1) \in V(G)$, and $G_2 = G \backslash G_1$,*
- *G_1 is a single-vertex pseudograph with l loops and G_2 is $G \backslash E_1$ where $E_1 \subseteq E(G)$ is the union of s loops and t parallel edges on $G[V(G_1)]$ such that $m - s - t = l - 1$, where m is the largest integer such that G_1^m is a minor of G,*
- *G_1 is a single-vertex pseudograph with no loops and ground set $\{u_1 u_2 \ldots u_i v_1 v_2 \ldots v_j\}$, $G[U] = G[\{u_1, \ldots, u_i\}]$ and $G[V] = G[\{v_1, \ldots, v_j\}]$ are connected subgraphs of G, G_2 is $G \backslash C$ where C is the set of edges with one end in $G[U]$ and the other in $G[V]$ and $C \neq \varnothing$,*
- *G_1 is a single-vertex pseudograph with no loops and ground set $\{z_1 z_2 \ldots z_k\}$ and G_2 is G/e where $e \in E(G)$ has one end in $G[Z] = G[\{z_1, \ldots, z_k\}]$ and the other end in $V(G) \backslash G[Z]$.*

As before, there are critical pairs for each minor operation; the first to vertex deletion, the second and third to edge deletion, and the fourth to edge contraction.

The final model is the most general. It uses pseudographs with the additional property that the edges have labels. We call these graphs *edge-labeled-pseudographs*, or simply *ELP*. Formally, an ELP is a triple (V, E, L), where V and E are as defined in the pseudograph model, and L is a function that assigns labels to the elements of E, with $L(e_1) \neq L(e_2)$ for all $e_1 \neq e_2$. The notation v^{E_1} denotes a vertex v whose loops are precisely those edges in $E_1 \subseteq E$. The minor operations for ELPs are the same as those for pseudographs with the condition that one must pick which edge to delete or contract.

Once again, we can determine the critical pairs in this model. Again, we omit the proof.

Theorem 32. *Let G be an ELP graph. The ordered pair (G_1, G_2) is a critical pair of* $\mathrm{MP}(G)$ *if and only if G_1 and G_2 satisfy one of the following:*

- *G_1 is a single-vertex ELP with no loops and $V(G_1) \in V(G)$, and $G_2 = G \backslash G_1$,*
- *G_1 is a single-vertex ELP with loop e and G_2 is $G \backslash e$, where e is not a cut-edge in $G[G_1]$,*
- *G_1 is a single-vertex ELP with no loops and ground set $\{u_1 u_2 \ldots u_i v_1 v_2 \ldots v_j\}$ where $G[U] = G[\{u_1, \ldots, u_i\}]$ and $G[V] = G[\{v_1, \ldots, v_j\}]$ are connected subgraphs of G, G_2 is $G \backslash C$ where C is the set of edges with one end in $G[U]$ and the other in $G[V]$, and $C \neq \varnothing$,*
- *G_1 is a single-vertex ELP with no loops and ground set $\{z_1 z_2 \ldots z_k\}$ and G_2 is G/e where $e \in E(G)$ has one end in $G[Z] = G[\{z_1, \ldots, z_k\}]$ and the other end in $V(G) \backslash G[Z]$.*

As in the previous case, there are critical pairs for each minor operation; the first to vertex deletion, the second and third to edge deletion, and the fourth to edge contraction.

In our final theorem, we note that the main results for trees in Section 5 hold for ELPs as well.

Theorem 33. *Let T be an ELP whose underlying simple graph T' is a tree. Then*

$$\min_{\mathcal{D} \in \mathfrak{D}} |\mathcal{D}| \leq \dim(\mathrm{MP}(T)) \leq |E(T)| + \frac{o_T}{2},$$

where all definitions are taken from Theorem 19.

Proof. Since T' is a tree, all edges are bridges. Thus, the critical pairs in $\mathrm{MP}(T)$ are of the the first, third, or fourth kind of critical pair for ELPs; not the second. As these are the same critical pairs in the simple graph model, we find that $\dim(\mathrm{MP}(T)) = \dim(\mathrm{MP}(T'))$. Theorem 19 finishes the proof. □

8. Conclusion and Future Directions

The poset obtained by using minor operations on a graph is quite natural, however, to the authors' knowledge, it has never been studied. The intent of this paper is to lay the combinatorial foundation for studying this class of posets (depending on the various manners of defining the minor operation). Before concluding, we would like to make a few remarks concering the future of this line of research.

Asymptotically, the dimension for simple graphs is driven by the critical pairs in $\mathcal{E}$. As such, it seems worthwhile to study the structure of minimal edge-cuts of subgraphs of a given graph. Unfortunately, there is nothing in the current literature that addresses this issue. Perhaps there is an interesting class of graphs for which this cut structure is particularly nice.

The primary problem we would like to see resolved is Conjecture 23. As noted in Remark 4, the proof of Theorem 30 will not generalize, although perhaps the work of Ding will be useful in answering this conjecture.

Extremal problems often arise in this area of combinatorics (e.g. see [1], [2]). For example, given $n < t \in \mathbb{N}$, what is the maximum number of edges a graph on n vertices can have while maintaining dimension at most t? Solutions to extremal questions may also help to shed light on an answer to Conjecture 23.

Finally, there is an algorithmic aspect to this theory. In general, computing dimension is NP-Complete [23]. However, we have seen that in some instances the dimension can be computed easily as it is closely related to a corresponding graph parameter (e.g., the dimension of the minor poset of a path). Are there certain classes of graphs, perhaps minor-free classes, for which dimension is efficiently computed? This question remains open even for trees, not only because the upper and lower bounds do not always match, but also because we do not know how to identify those graphs for which the bounds do match. It is an open question as to the complexity of computing the minimum size of a path decomposition for a given tree.

In the graph minors theorem [17], the parameter treewidth is used to describe the structural complexity of a graph; the smaller the treewidth, the more tree-like the graph, and hence the less complex the structure. Dimension is used similarly for posets, as the smaller the dimension, the more chain-like the poset, and hence the less structural complexity. It would be natural to hope that graphs of small treewidth would yield minor posets of small dimension. Unfortunately, the two notions of complexity do not mesh as nicely as could be wished. However, we still have hope that there is a common notion of complexity that will yield significant results concerning the structure of graphs and posets.

Acknowledgements. The authors would like to thank Tom Trotter, Graham Brightwell, and Guoli Ding for their insightful conversations in the development of this work.

References

[1] G. Agnarsson, *Extremal graphs of order dimension 4*, Math. Scand., 90 (2002), pp. 5–12.
[2] G. Agnarsson, S. Felsner, and W. T. Trotter, *The maximum number of edges in a graph of bounded dimension, with applications to ring theory*, Discrete Math., 201 (1999), pp. 5–19.
[3] L. Babai and D. Duffus, *Dimension and automorphism groups of lattices*, Algebra Universalis, 12 (1981), pp. 279–289.
[4] K. A. Baker, *Dimension, join-independence and breadth in partially ordered sets.* honors thesis (unpublished), 1961.
[5] G. Brightwell and W. T. Trotter, *The order dimension of convex polytopes*, SIAM J. Discrete Math., 6 (1993), pp. 230–245.
[6] G. R. Brightwell and W. T. Trotter, *The order dimension of planar maps*, SIAM J. Discrete Math., 10 (1997), pp. 515–528.
[7] Y. Dieng and C. Gavoille, *La structure des graphes sans mineur $K_{2,4}$*, in 10èmes Journées Graphes et Algorithmes, November 2008.
[8] R. Diestel, *Graph theory*, vol. 173 of Graduate Texts in Mathematics, Springer-Verlag, Berlin, third ed., 2005.
[9] G. Ding, *Graphs without large $K_{2,n}$-minors.* In preperation, August 2001.
[10] B. Dushnik and E. W. Miller, *Partially ordered sets*, Amer. J. Math., 63 (1941), pp. 600–610.
[11] S. Felsner, C. M. Li, and W. T. Trotter, *Adjacency posets of planar graphs*, Discrete Math., 310 (2010), pp. 1097–1104.
[12] S. Felsner and W. T. Trotter, *Posets and planar graphs*, J. Graph Theory, 49 (2005), pp. 273–284.
[13] Z. Füredi, P. Hajnal, V. Rödl, and W. T. Trotter, *Interval orders and shift graphs*, in Sets, graphs and numbers (Budapest, 1991), vol. 60 of Colloq. Math. Soc. János Bolyai, North-Holland, Amsterdam, 1992, pp. 297–313.
[14] R. Hegde and K. Jain, *The hardness of approximating poset dimension*, Electronic Notes in Discrete Mathematics, 29 (2007), pp. 435–443. European Conference on Combinatorics, Graph Theory and Applications, European Conference on Combinatorics, Graph Theory and Applications.
[15] I. Rabinovitch and I. Rival, *The rank of a distributive lattice*, Discrete Math., 25 (1979), pp. 275–279.
[16] S. P. Radziszowski, *Small Ramsey numbers*, Electron. J. Combin., 1 (1994), pp. Dynamic Survey 1, 30 pp. (electronic).

[17] N. Robertson and P. D. Seymour, *Graph minors. XX. Wagner's conjecture*, J. Combin. Theory Ser. B, 92 (2004), pp. 325–357.
[18] W. Schnyder, *Planar graphs and poset dimension*, Order, 5 (1989), pp. 323–343.
[19] A. Schrijver, *Minor-monotone graph invariants*, in Surveys in combinatorics, 1997 (London), vol. 241 of London Math. Soc. Lecture Note Ser., Cambridge Univ. Press, Cambridge, 1997, pp. 163–196.
[20] W. T. Trotter, *Combinatorics and partially ordered sets*, Johns Hopkins Series in the Mathematical Sciences, Johns Hopkins University Press, Baltimore, MD, 1992. Dimension theory.
[21] W. T. Trotter, Jr. and J. I. Moore, Jr., *Some theorems on graphs and posets*, Discrete Math., 15 (1976), pp. 79–84.
[22] ———, *The dimension of planar posets*, J. Combinatorial Theory Ser. B, 22 (1977), pp. 54–67.
[23] M. Yannakakis, *The complexity of the partial order dimension problem*, SIAM J. Algebraic Discrete Methods, 3 (1982), pp. 351–358.

School of Mathematics, Georgia Institute of Technology, Atlanta, GA 30332-0160 U.S.A.
E-mail address: `nstreib3@math.gatech.edu`

Department of Mathematics, University of California, San Diego (UCSD), La Jolla, CA 92093-0112 U.S.A
E-mail address: `s7young@math.ucsd.edu`