Mareike Massow (Autor)

Linear Extension Graphs and Linear Extension Diameter

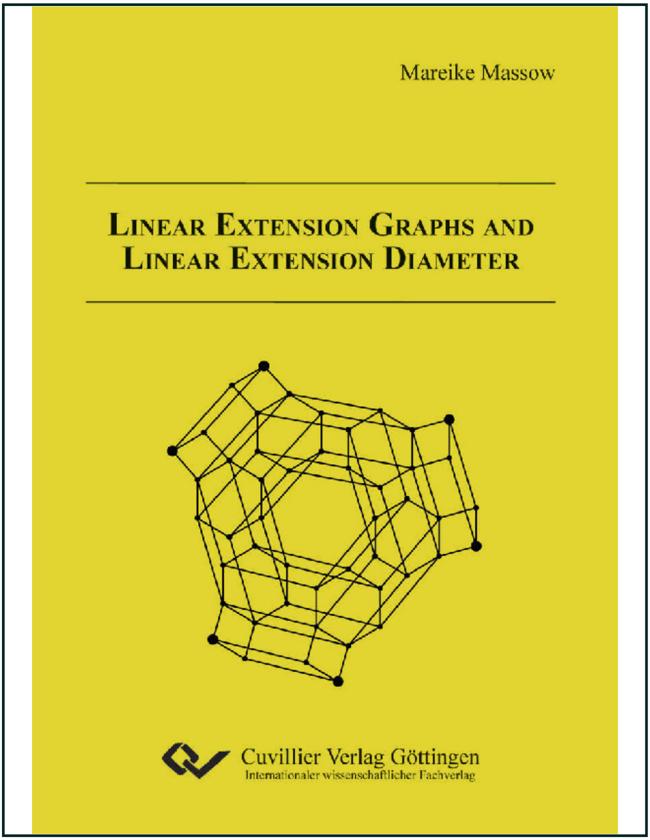

https://cuvillier.de/de/shop/publications/865

Chapter 2

Properties of Linear Extension Graphs

Linear extension graphs play a central role in this thesis. In this chapter we explore their beautiful structure and exhibit connections to the underlying posets.

A linear extension graph is a graph on the set of linear extensions of a poset, where two linear extensions are adjacent exactly if they differ in one adjacent swap of elements. This yields a coloring of the edges with the corresponding swaps. Figure 2.1 shows the Chevron with its linear extension graph.

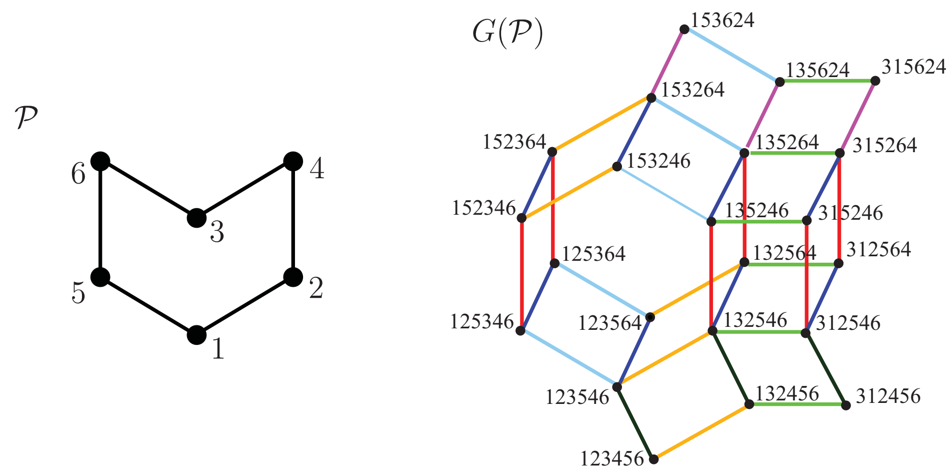

Figure 2.1: The Chevron and its linear extension graph with a swap coloring.

Linear extension graphs were originally defined by Pruesse and Ruskey in [48]. The first line of research on linear extension graphs was concerned with the existence of a Hamilton path; see also [52] and [61]. There has been subsequent research on other structural properties of $G(\mathcal{P})$, see [51], [50], [44], [45] and [22].

In Section 2.1 we present basic properties and previous results of linear extension graphs, and place them into a larger context. In Section 2.2 we concentrate on properties of the edge classes induced by the swap colors. This prepares Section 2.3, in which we will characterize which pairs of swap colors share an element. The last section of this chapter deals with the question which modifications of the poset $\mathcal{P}$ leave the linear extension graph $G(\mathcal{P})$ invariant.

2.1 Context and Previous Results

Before pointing out the connection of linear extension graphs to larger known graph classes, we want to make our notation precize and present some basic features of linear extension graphs.

Definition 2.1. *The* linear extension graph *$G(\mathcal{P}) = (V, E)$ of a poset $\mathcal{P}$ has as vertices the linear extensions of $\mathcal{P}$, with two of them being adjacent if they only differ in one adjacent transposition, or* swap, *of elements.*

Equivalently, for an edge $LL' \in E$ there is exactly one pair $xy \in Inc(\mathcal{P})$ with $x < y$ in L and $x > y$ in L'. This pair is the swap color *of LL'. The set of all edges with the same swap color forms a* color class *of $G(\mathcal{P})$.*

A graph G is a linear extension graph *if there is an underlying poset $\mathcal{P}$ such that $G = G(\mathcal{P})$.*

Recall that a reversal between L and L' is a pair of elements of $\mathcal{P}$ appearing in different orders in L and L'. Thus, each swap color corresponds to a reversal. By definition, two linear extensions of $\mathcal{P}$ are adjacent in $G(\mathcal{P})$ exactly if there is only one reversal between them, i.e., if the distance between them is 1. The very first observation about linear extension graphs we want to prove is that this generalizes to higher distances: The distance between two linear extensions L, L' of $\mathcal{P}$ equals the graph distance between the two corresponding vertices in $G(\mathcal{P})$.

Informally, this holds because changing L into L' is just the same as sorting the elements of $\mathcal{P}$ into the linear order L', starting with the linear order L. Thus we may, for example, use a selection sort algorithm (see e.g. [54]) to show that we never have to make a superfluous reversal in the sorting process.

For completeness, we give an explicit proof in the following lemma. A similar version of this lemma appeared in [44]. Note that we identify a vertex of $G(\mathcal{P})$ with the corresponding linear extension of $\mathcal{P}$.

Lemma 2.2. *Let L and L' be two linear extensions of a poset $\mathcal{P}$. Let T be a shortest L-L'-path in $G(\mathcal{P})$, and let S be the set of swap colors appearing on T. Then the swap colors in S are in bijection with the reversals between L and L', and each swap color in S appears only once on T.*

Proof. An L-L'-path T in $G(\mathcal{P})$ with swap color set S corresponds to a sequence of swaps modifying L into L'. To get from L to L' by swaps, we need to swap every reversal between L and L' at least once. We will describe a process modifying L into L' which swaps only reversals, and every reversal only once. This corresponds to a shortest L-L'-path in $G(\mathcal{P})$, and it follows that every shortest L-L'-path in $G(\mathcal{P})$ has the desired properties.

Let $L = u_1u_2 \dots u_n$ and $L' = v_1v_2 \dots v_n$. Start with finding the position of v_1 in L, say, $v_1 = u_j$. Then u_j is a minimal element of $\mathcal{P}$, so we can swap it with u_{j-1}, then u_{j-2}, and so on, until we obtain a linear extension of $\mathcal{P}$ which coincides with L' in the first element.

Now assume that we are given $L^i = u_1u_2 \dots u_n$ which coincides with L' in the first i elements. Find the position of v_{i+1} in L^i, say, $v_{i+1} = u_j$. It follows that $j > i$, and u_j is a minimal element of $\mathcal{P} - \{u_1, \dots, u_i\}$. We can thus swap u_j with u_{j-1}, then u_{j-2}, and so on, until we arrive at a linear extension L^{i+1} of $\mathcal{P}$ which coincides with L' in the first $i + 1$ elements. Inductively, we obtain $L^n = L'$.

Consider a pair of elements appearing in the same order in L and L', that is, a pair $x, y \in \mathcal{P}$ such that y has higher u-index and higher v-index than x. Then x and y are never swapped in our process, and hence we swap only reversals. Also, all of our swaps take an element which appears in L above some element with higher v-index, and swap it below that element. Therefore no pair of elements is swapped twice. This means that each reversal is swapped exactly once. $\square$

Let $xy \in \mathrm{Inc}(\mathcal{P})$ and set $G = G(\mathcal{P})$. We denote by W_{xy} the set of linear extensions of $\mathcal{P}$ in which $x < y$, and by W_{yx} the set of linear extensions of $\mathcal{P}$ in which $y < x$. Then the edges of swap color xy are exactly the edges connecting a linear extension in W_{xy} with a linear extension in W_{yx}. On the other hand, every path connecting a linear extension in W_{xy} with a linear extension in W_{yx} has to pass an edge with swap color xy. It follows that each color class is an edge cut of G.

Observe that the proof of the lemma above yields that each linear extension graph is connected. Now set $G_{xy} = G[W_{xy}]$ and $G_{yx} = G[W_{yx}]$.

Then G_{xy} is the linear extension graph of $\mathcal{P} \cup (x < y)$, and G_{yx} is the linear extension graph of $\mathcal{P} \cup (x > y)$. Therefore G_{xy} and G_{yx} are connected. Thus every color class cuts the graph G into exactly two components.

Lemma 2.3 ([51]). *Let $\mathcal{P}$ be a poset and $xy \in Inc(\mathcal{P})$. Then G_{xy} and G_{yx} are convex subgraphs of $G(\mathcal{P})$.*

Proof. Let $L, L' \in W_{xy}$. We need to show that an arbitrary shortest L-L'-path T in $G(\mathcal{P})$ is contained in W_{xy}. By Lemma 2.2, all swap colors appearing on T are reversals between L and L'. Hence xy does not appear as a swap color on T. Therefore T lies fully in W_{xy}. $\square$

The above lemma is the crucial tool embedding linear extension graphs into the much larger class of partial cubes.

Theorem 2.4. *Linear extension graphs are partial cubes, and the color classes equal the Djoković-Winkler classes.*

Proof. We use the characterization from Theorem 1.10. Let G be a linear extension graph, and let LL' be an edge in G. Suppose L and L' differ by the swap of the elements x and y, and $x < y$ in L. By Lemma 2.2, we have $W_{LL'} = W_{xy}$, that is, the set of linear extensions which are closer to L than to L' in G is exactly the set of linear extensions in which $x < y$. Hence it follows from Lemma 2.3 that G is a partial cube.

The definition of the Djoković-Winkler relation and $W_{LL'} = W_{xy}$ yield that the color classes of G equal the Djoković-Winkler classes of G. $\square$

For the following, we want to fix some more notation concerning the swap colors of linear extension graphs. In view of the above theorem we use the style of the partial cube notation.

Definition 2.5. *Let $G = G(\mathcal{P})$ be a linear extension graph. We denote the color classes of $G(\mathcal{P})$ by $\theta_1, \theta_2, \ldots, \theta_r$. The* swap partition $\Theta(G)$ *is the partition of $E(G)$ into the color classes $\theta_i, i = 1, \ldots, r$.*

A swap coloring *of G is a bijection c between $\Theta(G)$ and $Inc(\mathcal{P})$ which assigns to every color class the swap color of its edges. We usually denote the reverse mapping of c by θ. That is, if $\theta_i \in \Theta(G)$, and $c(\theta_i) = xy \in Inc(\mathcal{P})$, we write $\theta(xy) = \theta_i$. We extend this notation to the edges of G, i.e., if the swap color of e is xy, we write $c(e) = xy$, and if $e \in \theta_i$, we write $\theta(e) = \theta_i$.*

If a graph G is a linear extension graph and we need to specify which underlying poset a swap coloring of G refers to, we say that c is a swap coloring of G with respect to $\mathcal{P}$.

If G is a linear extension graph, and thus a partial cube, then it has a unique partition of its edges into Djoković-Winkler classes. Hence it follows from Theorem 2.4 that G has a unique swap partition $\Theta(G)$. Furthermore, as a partial cube, G has a Hamming labeling. By Lemma 1.11, this Hamming labeling is essentially unique, and partitioning the edges of G according to its coordinates again yields the swap partition $\Theta(G)$.

With Theorem 1.15, the following result now follows directly from Theorem 2.4.

Corollary 2.6 ([44]). *Let G be a linear extension graph with swap coloring c, and let C be a cycle in G. Then C is an isometric cycle exactly if for any two edges e, f on C, the following two conditions are equivalent:*

(i) e and f are opposite on C.

(ii) $c(e) = c(f)$. □

The fact that linear extension graphs are partial cubes is the background setting that we will use most in this thesis, but it is by far not the only larger context in which linear extension graphs can be viewed. In the remainder of this section, we want to point out some other connections. We only hint at the rich fields behind them. More connections and references can be found in Reuter's article [51].

The linear extension graph of the antichain, $G(\mathcal{A}_n)$, is the 1-skeleton of the *permutahedron* Π_{n-1}, see e.g. Ziegler's book on polytopes [66]. The permutahedron is a well-known polytope which is defined as the convex hull of all vectors that are obtained by permuting the coordinates of the vector $(1\ 2\ \dots\ n)^t$. In the following, we use the term permutahedron, and the notation Perm_n, to denote the 1-skeleton of Π_{n-1}.

The permutahedron Perm_n is also a hyperplane arrangement graph: The set of hyperplanes $H_{i,j} = \{x \in \mathbb{R}^n : x_i = x_j\}$ for $1 \leq i < j \leq n$ is called a *braid arrangement*. It is not difficult to see that Perm_n is isomorphic to the graph on all regions of the braid arrangements, where two regions are adjacent if they are separated by only one hyperplane. For an introduction to hyperplane arrangements, see e.g. Stanley's recent overview [56].

The braid arrangement is a Coxeter arrangement of type A_{n-1}. See the overview [24] by Fomin and Reading for an introduction to reflection groups and Coxeter arrangements. Hyperplane arrangement graphs are also the tope graphs of oriented matroids, see e.g. the book [4] by Björner et al..

A fact that we will use in the following is that hyperplane arrangement graphs are also partial cubes [46], thus, they interpolate between linear extension graphs and partial cubes.

In the last paragraphs we only considered $G(\mathcal{A}_n) = \mathrm{Perm}_n$. Now let $\mathcal{P}$ be an arbitrary poset on n elements, and consider the regions of the braid arrangement again. For each pair $i < j$ in $\mathcal{P}$, let us restrict the hyperplane arrangement to the *halfspace* induced by $H_{i,j}$ in which $x_i < x_j$. Then $G(\mathcal{P})$ is the region graph of the modified braid arrangement. If we restrict the modified braid arrangement to the box in which all coordinates of vectors are at most 1, we obtain Stanley's *order polytope*, see [55]. This additional restriction does not change the region graph. It also follows from Stanley's results that $G(\mathcal{P})$ is the dual graph of the canonical triangulation of the order polytope.

Let us consider what happens to Perm_n if we restrict the braid arrangement to halfspaces. Since Perm_n is a hyperplane arrangement graph, each edge e of Perm_n corresponds to a hyperplane $H_{i,j}$. Equivalently, e has swap color ij. Restricting Perm_n to the vertices induced by one halfspace of $H_{i,j}$ is the same as restricting it to one side of the edge cut $\theta(ij)$, thus to W_{ij} or W_{ji}.

This correspondence sheds new light on the notion of *convex* subgraphs in hyperplane arrangement graphs, since we can now see the close connection to geometric convexity. We formulate the following lemma for the larger class of partial cubes. For an adjacent pair x, y of vertices of a partial cube G, we call the subgraph $G[W_{xy}]$ a *halfspace* of G, cf. [45].

Lemma 2.7 ([45]). *Let G be a partial cube. A subgraph of G is convex if and only if it is an intersection of halfspaces.* □

With this lemma, it can now be seen that linear extension graphs are in bijection with convex subgraphs of the permutahedron: If we start with the full graph Perm_n and an n-element poset $\mathcal{P}$, we take each comparable pair in $\mathcal{P}$ and delete the halfspace of Perm_n where the corresponding relation is violated. This yields a convex subgraph of Perm_n. Conversely, starting with a convex subgraph we may consider it as intersection of halfspaces, and this yields the relations of the corresponding poset.

In fact, something even stronger is true. The convex subgraphs of Perm_n, ordered by inclusion, form a lattice $\mathrm{Conv}(\mathrm{Perm}_n)$. Now let us consider the set of all extensions of an n-element poset $\mathcal{P}$. If we order it by inclusion of the relations, it forms a poset $\mathrm{Ext}(\mathcal{P})$. Let us add an artificial global maximum to it, corresponding to the complete relation $\mathcal{P} \times \mathcal{P}$, and denote the result by $\mathrm{Ext}(\mathcal{P})^+$.

The following fundamental correspondance was first proved by Feldman in [19], and later rediscovered by Björner and Wachs [5] and by Reuter [51].

Theorem 2.8. *For an n-element poset $\mathcal{P}$, the lattice $Conv(Perm_n)$ is isomorphic to the dual of $Ext(\mathcal{P})^+$.* □