THE PARALLEL RIGIDITY INDEX OF A GRAPH

ALEXANDER POSTNIKOV AND ALEXEY SPIRIDONOV

Contents

Note. This is a preprint, which needs to be rewritten to be clearer and more concise. I apologize for the quality of the writing. — Alexey Spiridonov

1. Introduction

The classical graph rigidity problem asks: given a graph G, and a length l_e for every edge, how can we draw this graph in d dimensions with the specified edge lengths? A graph is rigid if such a *valid drawing* cannot be continuously transformed into another one, while staying valid at every instant. Parallel rigidity is analogous, but it fixes edge *directions*, leaving the lengths to vary freely. The notion of parallel rigidity has been investigated before by Whiteley [2] and Develin, Martin, and Reiner [1]. In two dimensions, parallel rigidity is equivalent to classical rigidity (and both are very well understood).

The usual approach to graph rigidity (both classical and parallel) involves studying a *rigidity matroid*. The matroid's ground set are the edges of K_n. An edge set is rigid iff it contains a basis of the matroid. Therefore, bases are minimally rigid sets of edges that cover all n vertices. Removing edges from a basis results in a non-rigid subgraph; adding edges keeps the graph rigid. Both Whiteley [2] and Develin-Martin-Reiner [1] work with the generic parallel rigidity matroid $P_{d,n}$.

We define a new notion – $RI_d(G)$, the *(generic) parallel rigidity index* of a graph G in d dimensions. It is the dimension of the variety consisting of drawings of G in $\mathbb{R}^d$ with fixed generic edge directions. The key difference from the rigidity matroid above is: the edge directions, not the vertex positions are generic.

We can restate what $RI_d(G)$ means in terms of the rigidity matroid. In our model, every rigidity circuit collapses into a point, and does not contribute to $RI_d(G)$. What remains is an independent set I, and its parallel rigidity index is a function of its rank. If the independent set has full rank, it's rigid and $RI_d(G) = 1$; if one needs to add k edges to I to make it full rank, $RI_d(G) = k + 1$. In other words,

$$RI_d(G) = \operatorname{rank} P_d - \operatorname{rank} I + 1.$$

So, one can, somewhat inelegantly, compute $RI_d(G)$ from any other description of the parallel rigidity matroid. Of course, 2.5 is itself one such description; by combining it with the requirement that $RI_d(G) = 1$, we get:

Corollary 1.1. *A graph G is a basis of $P_{d,|G|}$ if and only if for every edge e, there are some d spanning trees T_i such that $\bigcap_{i=1}^{d} T_i = \{e\}$, and no d such trees have empty intersection.*

We list below three results from the literature that appear closely related to ours.

Specializing Theorem 3.6 from [1] to graphs, we obtain a result similar to 5.1:

Corollary 1.2 (Develin, Martin, Reiner)**.** *Take a graph G, and convert it to a multigraph G' with $d-1$ parallel edges replacing each edge of G. Then, choose $e \in G'$, and double it to make G'' (so there are now d parallel edges between e's vertices). G'' can be partitioned into d disjoint trees if and only if G is an independent set of $P_{d,|G|}$.*

Theorem 8.2.2 from [2] is similar to both 5.1 and the Develin-Martin-Reiner result:

Theorem 1.3 (Whiteley, d-Parallel Graphs Theorem)**.** *If D is a subset of edges of K_n, the following are equivalent:*

1. *D is a basis of $P_{d,n}$.*
2. *$|D| = dn-(d+1)$ and for all nonempty subsets $D' \subset D$, $|D'| \le d|V(D')| - (d+1)$, where $V(D')$ denotes the set of vertices incident with D'.*
3. *D can be partitioned into $d+1$ edge-disjoint trees, exactly d incident with each vertex, but not d non-empty subtreest span the same subset of vertices.*

Several combinatorial descriptions of the independent sets and bases of this matroid are proved in [1, 2].

This quantity can be calculated, though somewhat laboriously, from the rigidity matroid descriptions of [1, 2]. We give an elegant combinatorial characterization of $RI_d(G)$, with a simple linear-algebric proof. As a corollary, we get another description of the bases of the rigidity matroid (it appears closely related to, but not immediately equivalent to previous results).

In Section 3, we formalize the degree of parallel flexibility $RI_d(G)$ as the dimension of the kernel of a matrix H_G derived from a cographic matroid. Then, in Section 4, we prove some linear algebra results about matrices of this type. Finally, in Section 5, we deduce the above rigidity results. We conclude with a discussion of our approach from the point of view of the parallel rigidity matroid used by [1, 2], and some simple examples.

2. Main Results

Let G be a connected graph on $[n]$, with the edge set $E \subseteq \{\langle i,j\rangle \mid i,j \in [n], i<j\}$. Fix a dimension d, and associate a vector $v_e \in \mathbb{R}^d, e \in E$ to each edge of the graph. Define a *valid drawing* to be a choice of positions $p_i, i \in [n]$ of the vertices of G, such that $p_i - p_j = \lambda_{\langle i,j\rangle} e_{\langle i,j\rangle}$ for all edges $\langle i,j\rangle$ and some $\lambda_{\langle i,j\rangle} \in \mathbb{R}$. Fix $p_1 = 0$ to eliminate drawings equivalent up to translation. If $\{p_i\}$ and $\{p'_i\}$ are valid drawings for given $G, \{v_e\}, d$, then so is $\{ap_i + bp'_i\}$; indeed:

$$\begin{aligned} ap_i + bp'_i - (ap_j + bp'_j) &= (ap_i - ap_j) + (bp'_i - bp'_j) = \\ &= a(p_i - p_j) + b(p'_i - p'_j) = \lambda_{\langle i,j\rangle} e_{\langle i,j\rangle} - \lambda'_{\langle i,j\rangle} e_{\langle i,j\rangle} = (\lambda_{\langle i,j\rangle} - \lambda'_{\langle i,j\rangle}) e_{\langle i,j\rangle}. \end{aligned}$$

So, valid drawings form a linear subspace $V_{d,\{v_e\}}(G)$ of $\mathbb{R}^{dn}$. The dimension of this subspace depends on the choice of $\{v_e\}$.

Proposition 2.1. *If we fix d and G, the minimum of $\dim V_{d,\{v_e\}}(G)$ occurs for a generic choice of edge directions.*

We prove this claim in 3.7. On the other hand, obtaining a large space of drawings is so simple that we can do it right away:

Claim 2.2. The maximum possible value of $\dim V_{d,\{v_e\}}(G)$ is $n-1$; it occurs, for instance, when all $\{v_e\}$ are collinear.

Proof. The graph is connected, and thus it contains a spanning tree. Delete all edges not in the tree; the vertex 1 is fixed. Its neighbors contribute one degree of freedom each (each such neighbor i is confined to the line given by $v_{\langle 1,i\rangle}$). In the same way, neighbors of neighbors contribute one degree of freedom, and so on, for a total of $n-1$. Adding back edges that are not in the spanning tree can only decrease $\dim V_{d,\{v_e\}}(G)$, so $n-1$ is the maximum. When the $\{v_e\}$ are collinear, every vertex except 1 can be placed anywhere on that line, which gives an $(n-1)$-dimensional subspace. □

One can obtain values of $\dim V_{d,\{v_e\}}(G)$ between the minimum and the maximum by making the $\{v_e\}$ collinear on an induced subgraph, or by introducing some other linear dependences among the edge directions. So, the key parameter is

$$D_d(G) = \min_{\{v_e\}} \dim V_{d,\{v_e\}}(G),$$

which we will call the *degree of parallel flexibility* of G in d dimensions. We will calculate it it by choosing edge directions $\{v_e\}$ with algebraically independent coordinates (recall 2.1). In this paper, we prove several results about $D_d(G)$:

Theorem 2.3. *$D_d(G)$ is equal to the minimal size of the intersection of d spanning trees of G.*

It's clear from this theorem is that as d grows, $D_d(G)$ decreases: $D_1(G) \geq D_2(G) \geq D_3(G) \geq \ldots$. But, we can prove something stronger:

Proposition 2.4. *Let $\lambda_i = D_i(G) - D_{i+1}(G)$; then, $\lambda = (\lambda_1, \lambda_2, \lambda_3, \ldots)$ is a partition. That is, $\lambda_i \geq 0$ and $\lambda_1 \geq \lambda_2 \geq \lambda_3 \ldots$.*

For d large enough, $D_d(G)$ stabilizes at the number of bridges of G (a bridge is an edge that doesn't belong to any cycle).

Observe that the only way to draw a triangle with generic edge directions in $\mathbb{R}^3$ is by setting all edge lengths to zero. Then, it's natural to ask: which G can be drawn in $\mathbb{R}^d$ with generic edge directions, so that none of the vertices coincide? Here is a different way of asking the same question. Every choice of vertex positions is a point in $\mathbb{R}^{dn}$; let $\Delta \subseteq \mathbb{R}^{dn}$ be the set of vertex positions with some coinciding vertices. A point from $\mathcal{P} = \mathbb{R}^{dn} \backslash \Delta$ uniquely determines the edge directions $\{v_e\}$ of the resulting drawing; these $\{v_e\}$ correspond to a point in $\mathcal{E} = \left(\mathbb{RP}\{\}d-1\right)^{|E|}$. These correspondences give a map $\phi : \mathcal{P} \to \mathcal{E}$. If a graph has a drawing with noncoincident points and generic edge directions, then most of the points of $\left(\mathbb{RP}\{\}d-1\right)^{|E|}$ will be in the image of ϕ. Formally, we would like to know when $\overline{\phi(\mathcal{E})}$, the algebraic closure of the image, equals $\mathcal{E}$. This is equivalent to the original condition (is there a drawing with non-coinciding vertices for generic edge directions). Indeed, if a drawing exitsts for one set of algebraically independent edge directions, one exists for all other choices of algebraically independent edge directions. Therefore closure is equal to $\mathcal{E}$. Conversely, if $\overline{\phi(\mathcal{E})} = \mathcal{E}$, then what?

Corollary 2.5. *A graph does not impose algebraic dependences on its edge directions in d dimensions if and only if for every edge $e \in E$ some minimal intersection of d spanning trees contains e.*

In 5.4, we list several results from the literature that are closely related to this corollary.

3. A Parallel Rigidity Matrix

First, we set about the business of defining a parallel rigidity matrix. This matrix is not the same as the one defined by Develin, Martin and Reiner, nor is it the same as Whiteley's variant – both of those arise from the graphic, rather than cographic, matroid of G.

3.1. **Graph orientation.** An edge vector $l_{\langle i,j\rangle} v_{\langle i,j\rangle}$ can be interpreted either as pointing from i to j, or from j to i. We choose this orientation according to the following rule.

Definition 3.1 (Orientation of G)**.** If $i < j$, orient the edge $\langle i, j\rangle$ from i to j. Also define

$$o(i,j) = \begin{cases} 1 & i < j \\ -1 & i > j \end{cases}$$

for an adjacent pair of vertices i and j, and call the edge between them $e(i, j)$.

For example, $K_4 \hookrightarrow \mathbb{R}^2$ would be oriented as follows

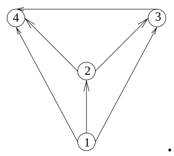

.

3.2. **The drawing in terms of edge lengths.** Let

$$l = (l_e \in \mathbb{R} \mid e \in E) \in \mathbb{R}^{|E|}$$

be the vector of edge lengths for a given drawing. Note that the lengths don't have to be positive: negative lengths simply serve to reverse the orientation of the vector.

Claim 3.2. The length vector l uniquely determines a drawing D_l (when it exists).

Proof. We start with the node 1 (position fixed at the origin); for every neighbor i of 1, we have the length and the direction of the edge $\langle 1, i\rangle$. Thus, we have the position p_i for every i adjacent to 1. Since the graph is connected, we proceed inductively to get the positions of all the vertices. □

3.3. **Cycles constrain the edge lengths.** For which l does D_l exist? In the construction above, all vertex positions are completely determined by (any) spanning tree of G. The other edges *must* have the proper direction and length for the drawing to exist. More formally: pick any spanning tree T, and consider $\langle i, j\rangle \in E \backslash T$. Then, the drawing exists iff $l_{\langle i,j\rangle}$ and $v_{\langle i,j\rangle}$ have the unique values that follow from the (already determined) positions of i and j.

Restating the preceding condition in another way, D_l exists iff every cycle C in the graph "closes up", with

$$\sum_{(i,j)\in C} o(i,j) v_{e(i,j)} l_{e(i,j)} = 0, \tag{3.1}$$

where (i, j) ranges over all pairs of consecutive vertices in the cycle (reading around the cycle in one direction). Multiplying by $o(i, j)$ adjusts the orientation so that the vector $v_{e(i,j)} l_{e(i,j)}$ points from i to j.

For every cycle of G, fix one of the two possible reading orders (e.g. the cycle (123) can also be read (132)). Let the *cycle matrix* M^G have rows indexed by the cycles, and columns indexed by the edges of G. The entry for cycle C and edge e, denoted $M^G_{C,e}$, is 0 if the cycle does not contain the edge. If $e \in C$, let (i, j) be the pair of vertices of e listed in the reading order of the cycle. Then, $M_{C,e} = o(i, j)$. For our K_4 example,

$M^G =$

	12	13	14	23	24	34
(123)	1	−1		1		
(124)	1		−1		1	
(134)		1	−1			1
(234)				1	−1	1
(1234)	1		−1	1		1
(1243)	1	−1			1	−1
(1324)		1	−1	−1	1	
(1342)	−1	1			−1	1
(1423)		−1	1	1	−1	
(1432)	−1		1	−1		−1

.

3.4. **Only generating cycles are needed.** Let $E_{\mathbb{Z}}$ be a vector space over $\mathbb{Z}$ with the edges E forming a basis. The vector $ne, e \in E, n \in \mathbb{Z}$ is interpreted as $|n|$ copies of e. If n is positive, e is oriented as defined in 3.1; if n is negative, e is taken with the reverse orientation. In the above K_4 example, let

$$C_{(123)} = e_{12} + e_{23} - e_{13} \in E_{\mathbb{Z}}.$$

This is an oriented cycle. Let $C_{\mathbb{Z}} \subset E_{\mathbb{Z}}$ be the subspace spanned by all oriented cycles. Starting with the whole set of cycles, we can omit some to get a cycle basis for $C_{\mathbb{Z}}$, which we will call a set of *generating cycles*. For instance, this linear dependence between 3 cycles in K_4 allows us to omit one of them:

$$C_{(123)} - C_{(234)} = C_{(1243)}.$$

Note that the rows of M^G (defined in the previous section) are the cycles of G written out in the edge basis of $E_{\mathbb{Z}}$. Let $\tilde{M}^G$ be a minor with some generating cycles as the rows. Here's one possible choice for K_4:

$\tilde{M}^G =$

	12	13	14	23	24	34
(123)	1	−1		1		
(124)	1		−1		1	
(134)		1	−1			1

The final observation is that if the cycles of $\tilde{M}^G$ close up, then so do the cycles they generate. Indeed, if the cycle C is a linear combination of generating cycles C_i, then the corresponding linear combination of (3.1) for the C_i proves that C closes up.

3.5. **$\tilde{M}^G$ is a representation of the cographic matroid.** This fact will not be used directly until 5.1, but it should be helpful in understanding the intervening sections. The bases of the cographic matroid are the complements of the graphic matroid's bases, which are spanning trees. So, we will prove:

Lemma 3.3. *A set of columns of $\tilde{M}^G$ is maximal independent iff its complement is a spanning tree.*

Proof. The rows of $\tilde{M}^G$ are linearly independent over $\mathbb{Z}$ by construction. Since all the entries of $\tilde{M}^G$ are also in $\mathbb{Z}$, the rows are independent over $\mathbb{R}$ as well.

Claim 3.4. The number of generating cycles is $|C_G| = |E| - (n-1)$.

Proof. We'll start with spanning tree of G (no cycles), and add edges one-by-one until we get G. Since the graph is connected at every stage of the game, adding a new e adds at least one cycle. None of the previous generating cycles contain this edge, so we add at least one *generating* cycle. But, suppose that adding e introduced two generating cycles. Orient them with e pointing in opposite directions. Adding the two cycles yields a third cycle C (possibly a disjoint union of cycles; that's okay) that doesn't e. So, C is also generated by some of the previous generating cycles, and we get a linear dependence. Adding each edge adds one generating cycle, and we have to add exactly $|E| - (n-1)$ edges. □

It follows that $\tilde{M}_G$ has full rank, namely $|C_G|$, and therefore the number of columns $|E| \geq |C_G|$. Thus, a maximal independent set of columns $I \in E$ has size $|C_G|$. Let M_I be the minor induced by this column set; it has full rank, so no nonzero linear combination of the rows is zero. Every cycle of G is some nonzero linear combination of the rows of $\tilde{M}_G$. Thus, every cycle in G contains an edge from I. Therefore, the complement $\overline{I}$ has no cycles, and by 3.4, it has $n-1$ edges. So, $\overline{I}$ is a spanning tree. That is exactly what we needed to show.

□

3.6. Drawings are described by the parallel rigidity matrix. We now return to the original question: for which l does there exist a drawing D_l of a fixed graph G with a fixed set of edge directions $\{v_e\}$? In the preceding paragraphs, we established that D_l exists iff

$$\sum_{e \in C} \left(\tilde{M}^G_{C,e} v_e \right) l_e = 0$$

for all C in the set of generating cycles C_g; that's d linear equations for every cycle. So, the answer is given by a system of $d \cdot |C_g|$ linear equations in $|E|$ variables. Suppose that $\{v_e\}$ for our K_4 example is given by the following matrix:

$$(v_{e,i}) = \begin{pmatrix} x_1 & x_2 & x_3 & x_4 & x_5 & x_6 \\ y_1 & y_2 & y_3 & y_4 & y_6 & y_6 \end{pmatrix};$$

then, the linear system in question has the matrix:

$$H_G = \begin{array}{cc|cccccc} \text{dim} & C & 12 & 13 & 14 & 23 & 24 & 34 \\ \hline 1 & (123) & x_1 & -x_2 & & x_4 & & \\ & (124) & x_1 & & -x_3 & & x_5 & \\ & (134) & & x_2 & -x_3 & & & x_6 \\ 2 & (123) & y_1 & -y_2 & & y_4 & & \\ & (124) & y_1 & & -y_3 & & y_5 & \\ & (134) & & y_2 & -y_3 & & & y_6 \end{array}.$$

In the general case, that's d copies of $\tilde{M}^G$ stacked on top of each other, with every column of every A_G copy multiplied by the corresponding $v_{e,i}$. We call this the *parallel rigidity matrix.* In the terminology of Section 4, H_G is the column-wise tensor $\tilde{M}^G \otimes_c (v_{e,i})$. A drawing exists for every $l \in \mathbb{R}^{|E|}$ such that $H_G l = 0$, so the object of interest is the kernel of H_G.

In the introduction, we said $D_d(G)$ is the number of degrees of freedom of G's family of drawings (with generic edges, more on this in the next subsection). We now have an equivalent, more explicit, definition: for a graph G, the *degree of parallel flexibility* is $D_d(G) = \dim\ker H_G$. Degree 0 means that the graph can only be drawn with the specified edge directions by collapsing all the vertices into one point. Degree 1 – there exists a drawing of the graph, such that not all the vertices are coincident, and the drawing is unique up to scaling (this is called *rigid*). Higher degrees give additional degrees of freedom to the drawing.

3.7. Generic edge directions minimize $D_d(G)$. As noted in Section 2, the above description of $D_d(G)$ in terms of H_G depends on the choice of $\{v_e\}$, and not just on the graph. For instance, if we set $x_i = y_i = 1$ in the K_4 example of 3.6, the resulting H_G has kernel dimension 4. On the other hand, picking generic values for x_i and y_j gives $\dim\ker H_G = 0$.

In order to obtain a combinatorial problem, we will choose $\{v_e\}$ that minimize $\dim\ker H_G$, or, equivalently, maximize rank H_G. The rank of H_G is the size of it largest nonsingular square minor M. Let $v_{e,i}$ be the ith component of v_e. Take $v_{e,i}, e \in E, i \in [d]$ to be indeterminates; then, the determinant of every minor is a polynomial in those indeterminates. Take M to be the largest minor of H_G such that the determinant polynomial is not identically zero. Obviously, there is no larger minor that is nonsingular for any fixed numeric choice of $(v_{e,i})$. Moreover, since the polynomial is not identically zero, there is a dense set of $(v_{e,i}) \in \mathbb{R}^{d \cdot |E|}$ that make the polynomial nonzero. So, a generic choice of $\{v_e\}$ minimizes $\dim\ker H_G$.

To make some of the subsequent arguments easier, we will assume specifically that the $(v_{e,i})$ are algebraically independent; this is equivalent to treating all the $v_{e,i}$ as indeterminates. So, the final definition of $D_d(G)$ depends only on combinatorial data:

Definition 3.5. The *degree of parallel flexibility* $D_d(G)$ of a graph G in d dimensions is $\dim\ker H_G$, with H_G constructed using edge directions $\{v_e\}$ such that the coordinates $(v_{e,i})$ are algebraically independent.

4. Column Tensor Product by a Generic Matroid

All of the rigidity results of Section 5 are derived by applying linear-algebraic statements from this section. Though the results are about matrices, it will be convenient to use the language of matroids. In this section, all matroids are explicitly represented by the columns of a matrix. We will use interchangeably the matroids and the corresponding matrices.

Both results in this section use the following construction.

Definition 4.1. If A and B are matroids on the same ground set E, then the *column tensor product* $A \otimes_c B$ is a matroid on E with $e \in E$ represented by the Kronecker product of the eth columns of A and B.

Here is an example:

$$A = \begin{pmatrix} a & d \\ b & e \\ c & f \end{pmatrix}, B = \begin{pmatrix} x & u \\ y & v \end{pmatrix} \Rightarrow A \otimes_c B = \begin{pmatrix} xa & ud \\ xb & ue \\ xc & uf \\ ya & vd \\ yb & ve \\ yc & vf \end{pmatrix}.$$

Note that that the column tensor product is commutative on matroids (and on matrices, up to a rearrangement of rows). We will actually use a less symmetric formulation: for every row r of B, take a copy of A, and multiply the ith column by r_i. Then, $A \otimes_c B$ is the union of all the rows of all the modified copies of A.

In this section, we address the scenario where a given matroid is column-tensored with a "generic" representable matroid of a certain rank. Here's what generic means.

Definition 4.2. Given a matroid A, its *d-generic matroid* $G_d(A)$ has the same ground set, and is defined by a matrix with d rows such that all its entries are algebraically independent.

4.1. Rank of a column tensor of a matroid with a d-generic matroid.

Proposition 4.3. *Let A be a represented matroid of positive rank. For $d \in \mathbb{N}$, define $H = A \otimes_c G_d(A)$. Then,* $\operatorname{rank} H$ *is the maximum size of the union of d bases of A.*

Proof. For most of this proof A and H will be treated as matrices. Recall that the column tensor product H can be decomposed into d modified copies of A; call these $H^i, i = 1, \ldots, d$. The columns in H^i are just those of A, each multiplied by some nonzero constant. This does not change the rank, so a minor of of H^i is nonsingular iff the corresponding minor of A is nonsingular. To take advantage of this, we will express the rank of H in terms of the ranks of H^i.

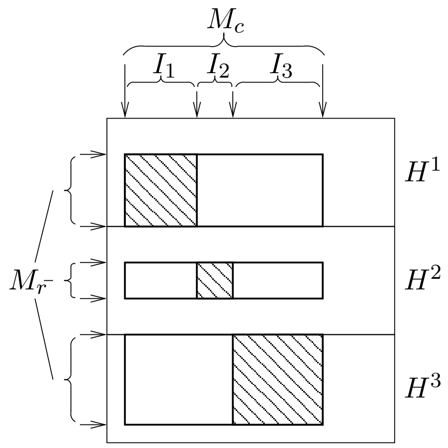

FIGURE 4.1. A schematic drawing (with the rows and columns reordered for ease of drawing) of the shape of permutations that contribute to a monomial of $\det H$.

The rank of H is the maximum size of a nonsingular minor of H. Let M be some nonsingular minor, denoting the row set by $M_r = \operatorname{row} M$ and the column set by $M_c = \operatorname{col} M$. It splits up among the H^i as shown in Figure 4.1 on page 9.

Since the entries of $G_d(A)$ are algebraically independent, we might as well replace them by indeterminates; call these indeterminates g_{ij} with $i \in [d], j \in [|E|]$. Since M is nonsingular, $\det M$ is a nonzero polynomial in these g_{ij}. Therefore, some monomial m has a nonzero coefficient. The determinant is a sum over permutations

$$\det M = \sum_{\sigma \in S_{|M|}} \operatorname{sgn} \sigma \prod_{j \in M_c} (H)_{\sigma(j),j},$$

where σ maps M_c to M_r in the natural way. So, every monomial in the sum has exactly one g_{ij} for every $j \in M_c$; and, for a fixed i, the number of g_{ij} in the monomial is $|\operatorname{col} H^i \cap M_c|$. For the chosen monomial m and some $i \in [d]$, let $I^i = \{j | g_{ij} \text{ in monomial } m\}$. Then, m splits up as follows:

$$m = \prod_{i=1}^{d} \prod_{j \in I^i} g_{ij},$$

and the I^i are pairwise disjoint (see Figure 4.1 on page 9). Any permutation confined to the sub-minors $M^i = M|_{H^i}$ (shaded squares in the figure) will contribute to the coefficient of m; thus, the coefficient is:

$$\prod_{i=1}^{d} \det M^i \neq 0.$$

So, $\det M^i \neq 0$. Therefore, we went from a nonsingular minor of H to a set of nonsingular minors M^i of H^i with disjoint column sets. Now, we show that the converse procedure works.

We will run the above argument backwards, keeping all the set-up, including the submatrices H^i, the indeterminates g_{ij}, the monomial m, and start with the above minors M^i. We construct the minor M of H by combining the row and column

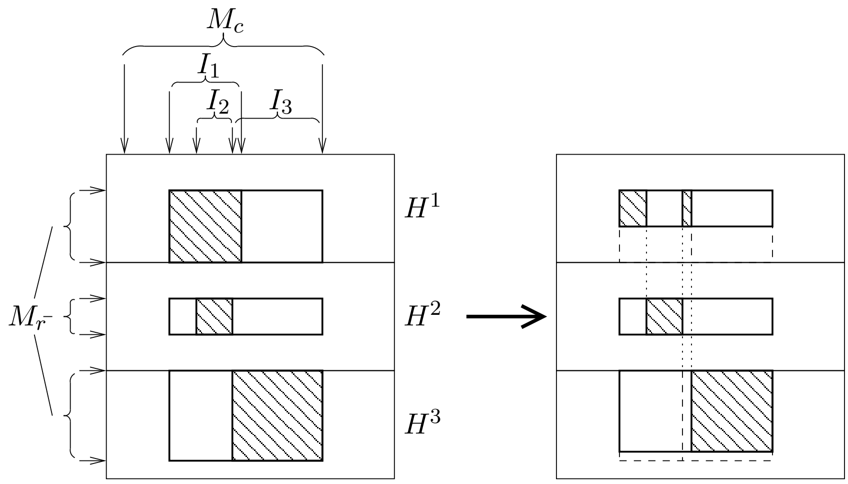

FIGURE 4.2. A variant of Figure 4.1 on page 9 with minors overlapping in columns, and a possible way to resolve this. The resulting minor at the top of the right-hand-side matrix, though drawn split, is a square.

sets of M^i. The determinant of M is again a polynomial. As discussed above, the coefficient of monomial m in $\det M$ will be exactly

$$\prod_{i=1}^{d} \det M^i \neq 0, \tag{4.1}$$

and so M will be nonsingular.

We needed to find a nonsingular M of the largest possible size. The previous paragraph tells us that we can instead find a set of d nonsingular M^i, minors of H^i with disjoint column sets, such that the union of their column sets is maximal. But, in fact, we can omit the disjointness condition.

Claim 4.4. Let the minors M^i be as above, except that some column sets overlap. Then, the combined minor has full rank.

Proof. Suppose some column sets overlap; the row sets still do not, so the combined minor M is not a square, but a rectangle with fewer columns than rows. A possible scenario is pictured in Figure 4.2 on page 10. The solution is to drop some rows and columns from some of the minors. The following procedure produces nonsingular minors with nonoverlapping column sets, and an unchanged column set union:

(1) Suppose $\operatorname{col} M^i \cap \operatorname{col} M^j \neq \emptyset$. We'll (arbitrarily) decide to remove columns from M^i.
(2) Remove $\operatorname{col} M^j$ from M^i.
(3) The remaining non-square matrix has full rank, so it has a nonsingular minor of the maximum possible size. Remove the rows not in this minor.

We repeatedly apply this procedure to remove all column overlaps. Reasoning just like in (4.1) proves this claim.

□

Therefore, we seek a set of nonsingular minors M^i, one from every H^i, such that

$$\left|\bigcup_{i=1}^{d} \operatorname{col} M^i\right|$$

is maximal. Recall that a minor of H^i is nonsingular iff the corresponding minor of A is nonsingular. It's thus enough to find a set of d nonsingular minors A^i of A covering a maximal number of columns. We might as well look for *maximal* nonsingular minors – this won't reduce the number of columns they cover. Therefore, the rank of H is the maximum size of the union of d bases of A.

□

4.2. **Rank increments of $A \otimes G_d(A)$ form a partition.**

Proposition 4.5. *If A is a represented matroid,*

$$\operatorname{rank} A \otimes_c G_d(A) - \operatorname{rank} A \otimes_c G_{d-1}(A), d = 1, 2, 3, \ldots$$

forms a partition.

Proof. Since we are, in effect, multiplying columns of copies of A by indeterminates, $H_d = A \otimes_c G_d(A)$ is $H_{d-1} = A \otimes_c G_{d-1}(A)$ after adjoining on the bottom A_d, a modified copy of A. We will denote this $H_d = \frac{H_{d-1}}{A_d}$. Then, in turn, $H_{d-1} = \frac{H_{d-2}}{A_{d-1}}$. The intuition is that the rows of A_d and A_{d-1} are completely symmetric with respect to linear independence. Herse's what that means formally – all the indeterminates used by the two matrices are algebraically independent. So, consider a row r of A_d and the corresponding row r' of A_{d-1}. If r is involved in some linear dependence with rows of H_{d-2} and A_d, then r' participates in the symmetric linear dependence with rows of H_{d-2} and A_{d-1} (and vice versa). In particular, adjoining A_d to H_{d-2} yields exactly the same rank increase, call it r, as adjoining A_{d-1} to H_{d-1}. But then, adjoining A_d to $\frac{H_{d-2}}{A_{d-1}}$ cannot increase the rank by more than r – the additional presence of A_{d-1} can only increase the number of linear dependences involving rows of A_d. This proves the claim. □

5. Proofs of Results

5.1. **Rank of a generic H_G in combinatorial terms.** In the terminology of Section 4, our algebraically independent edge directions $v_{e,i}$ form the d-generic matroid $G_d(\tilde{M}^G)$ for $\tilde{M}^G$. Then, as noted before, the parallel rigidity matrix H_G equals the column-wise tensor $\tilde{M}^G \otimes_c G(\tilde{M}^G)$. So, 4.3 applies, and the rank of H is the maximum size of the union of d bases of $\tilde{M}^G$. From 3.5, we know that these bases are complements of spanning trees of G, and so

$$\operatorname{rank} H_G = \max_{T_1,\ldots,T_d} \left|\bigcup_{i=1}^{d} \overline{T_i}\right| = \max_{T_1,\ldots,T_d} \left|\overline{\bigcap_{i=1}^{d} T_i}\right|.$$

The maximal complement of the intersection occurs exactly when the intersection itself is minimal. There are $|E| - \left|\bigcap_{i=1}^{d} T_i\right|$ edges in this complement, so we get that

$$\operatorname{rank} H_G = |E| - \min_{T_1,\ldots,T_d} \left|\overline{\bigcap_{i=1}^{d} T_i}\right| \Rightarrow \dim\ker H_G = \min_{T_1,\ldots,T_d} \left|\overline{\bigcap_{i=1}^{d} T_i}\right|.$$

This proves Theorem 2.3: the degree of parallel flexibility in d dimensions is the minimal size of the intersection of d spanning trees.

5.2. **Behavior of $D_d(G)$ as d grows.** The degree of parallel flexibility decreases with larger dimensions: cycles with generic edge directions become progressively more rigid. In 2 dimensions, a generic 4-cycle has one degree of freedom; in 3, it's rigid; in 4, one can only be drawn with all edges having zero length.

Alternatively, in terms of Theorem 2.3, adding more trees to the intersection can only reduce its size. So, $D_d(G), d = 2, 3, \ldots$ is a partition. For d large enough, any edge contained in a cycle is contracted to a point. The only remaining edges are bridges, each of which confers a degree of freedom. Thus, $D_d(G), d \to \infty$ stabilizes at the number of bridges in G.

Also, 4.5 proves 2.4, which states that $D_d(G) - D_{d+1}(G)$ is a partition.

5.3. **Drawings with non-coinciding vertices.** As noted before, generic cycles in a large enough dimension can only be drawn with zero-length edges. So, a natural question is: when can a graph be drawn generically in d dimensions so that all vertices are distinct? Alternatively, when do the edge directions of a graph have no algebraic dependences among them?

We can use Theorem 2.3 to answer this. If an edge e of G is forced to have zero length by an algebraic dependence, then $D_d(G/e) = D_d(G)$. This is because contracting an edge has exactly the same effect on the drawing as forcing it to have zero length.

The converse is also true. It follows from the construction of 3.6 that the space of possible edge lengths $L_d(G)$ is a linear subspace. Contracting e is the same as taking the intersection of $L_d(G)$ with the hyperplane $\{length(e) = 0\}$. That preserves the dimension only if $L_d(G) \subseteq \{length(e) = 0\}$, which is to say that e is already forced to have zero length.

So, we want to have $D_d(G/e) < D_d(G)$ for every edge e. That means that minimal intersections of d spanning trees of G/e have fewer edges than in G. We can convert a spanning tree for G/e into one for G by adding back e. So, if some minimal d-intersection of spanning trees for G contains e, the same spanning trees without e will give a minimal d-intersection for G/e with $D_d(G) - 1$ edges. If the minimal intersection for G/e were smaller than $D_d(G) - 1$, we could add e to all the trees, and get a smaller intersection for G.

Thus, we proved 2.5 – a graph has a $d-$dimensional drawing with non-coinciding vertices if and only if every edge is contained in a minimal intersection of spanning trees. The corollary has several related results in the literature, which we discuss in the next section.

5.4. $D_d(G)$ **from the point of view of the parallel rigidity matroid.** The classical approach to graph rigidity (parallel, as well as edge length-preserving) involves studying a so-called rigidity matroid. The matroid is considered for a fixed number of vertices n, and describes which edge sets on these vertices are rigid. The edges of K_n are its ground set. A basis of the matroid is a minimally rigid set of edges that covers all vertices. Removing edges from a basis results in a non-rigid subgraph; adding edges produces an overdetermined graph with algebraic dependences between the edges. Both Whiteley and Develin-Martin-Reiner work with the generic parallel rigidity matroid $P_{d,n}$ [1, 2].

We can restate what $D_d(G)$ means in terms of the rigidity matroid. In our model, every rigidity circuit collapses into a point, and does not contribute to $D_d(G)$. So, we prepare the graph by first cotracting circuits one by one until none remain. What remains is an independent set I, and its degree of flexibility is a function of its rank. If the independent set has full rank, it's rigid and $D_d(G) = 1$; if one needs to add k edges to I to make it full rank, $D_d(G) = k + 1$. In other words,

$$D_d(G) = \operatorname{rank} P_d - \operatorname{rank} I + 1.$$

So, one can, somewhat inelegantly, compute $D_d(G)$ from any other description of the parallel rigidity matroid. Of course, 2.5 is itself one such description; by combining it with the requirement that $D_d(G) = 1$, we get:

Corollary 5.1. *A graph G is a basis of $P_{d,|G|}$ if and only if for every edge e, there are some d spanning trees T_i such that $\bigcap_{i=1}^{d} T_i = \{e\}$, and no d such trees have empty intersection.*

We list below three results from the literature that appear closely related to ours. Specializing Theorem 3.6 from [1] to graphs, we obtain a result similar to 5.1:

Corollary 5.2 (Develin, Martin, Reiner). *Take a graph G, and convert it to a multigraph G' with $d-1$ parallel edges replacing each edge of G. Then, choose $e \in G'$, and double it to make G'' (so there are now d parallel edges between e's vertices). G'' can be partitioned into d disjoint trees if and only if G is an independent set of $P_{d,|G|}$.*

Theorem 8.2.2 from [2] is similar to both 5.1 and the Develin-Martin-Reiner result:

Theorem 5.3 (Whiteley, d-Parallel Graphs Theorem). *If D is a subset of edges of K_n, the following are equivalent:*

(1) *D is a basis of $P_{d,n}$.*
(2) *$|D| = dn - (d+1)$ and for all nonempty subsets $D' \subset D$, $|D'| \leq d|V(D')| - (d+1)$, where $V(D')$ denotes the set of vertices incident with D'.*
(3) *D can be partitioned into $d+1$ edge-disjoint trees, exactly d incident with each vertex, but not d non-empty subtreest span the same subset of vertices.*

We also noticed a strange coincidence. Whiteley defined a notion of 3-dimensional *body-and-hinge rigidity.* One of his results contains a condition of exactly the same form as Theorem 2.3. Here is an excerpt from Theorem 12.2.1 in [2]:

Theorem 5.4 (Extract of the 3-Hinge Theorem). *For a graph G, the following are equivalent:*

(1) *This graph is generically 3-hinge rigid.*
(2) *If each edge of the graph is replaced by five copies, the resulting multigraph $5G$ contains six edge-disjoint spanning trees.*

The second condition is equivalent to: "there exist 6 spanning trees in G with empty intersection." In other words, a graph is generically 3-hinge rigid if and only if has $D_6(G) = 0$! Would something of this form hold true for a generalization of 3-hinge rigidity to higher dimensions?

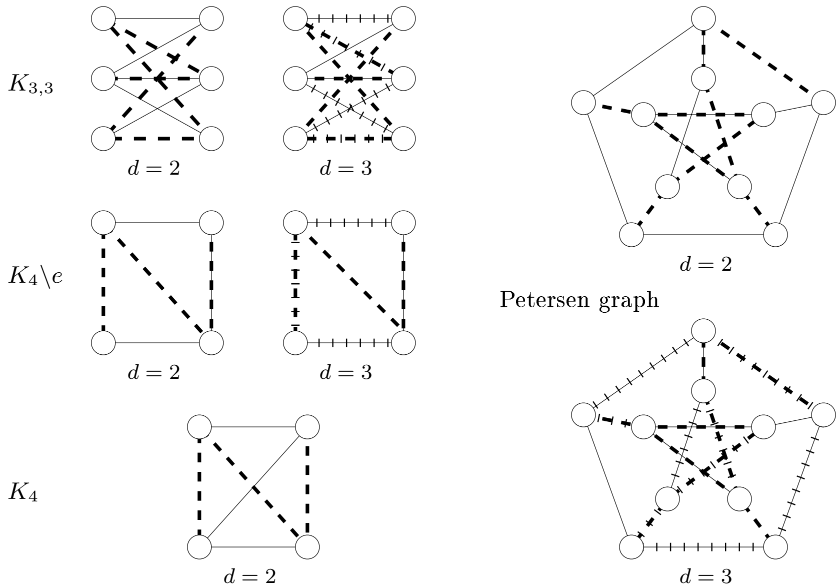

Figure 5.1. Examples of d minimally overlapping spanning trees in various graphs.

References

[1] Mike Develin, Jeremy Martin and Victor Reiner, *Rigidity theory for matroids*, Commentarii Mathematici Helvetici **82** (2007), 197–233.

[2] Walter Whiteley, *Matroids from Discrete Geometry*, in Matroid Theory, J. Bonin, J. Oxley and B. Servatius (eds.), AMS Contemporary Mathematics, (1997), 171-313.