Chapter 6
Orthogonal representations II: Minimal dimension

Nachdiplomvorlesung by László Lovász
ETH Zürich, Spring 2014

1 Minimum dimension

Perhaps the most natural way to be "economic" in constructing an orthogonal representation is to minimize the dimension. We can say only a little about the minimum dimension of all orthogonal representations, but we get interesting results if we impose some "non-degeneracy" conditions. We will study three nondegeneracy conditions: general position, faithfulness, and the strong Arnold property.

1.1 Minimum dimension with no restrictions

Let $d_{\min}(G)$ denote the minimum dimension in which G has an orthogonal representation. The following facts are not hard to prove.

Lemma 1.1 *For every graph G,*

$$\vartheta(G) \le d_{\min}(G),$$

and

$$c \log \chi(\overline{G})) \le d_{\min}(G) \le \chi(\overline{G}))$$

for some absolute constant $c > 0$.

Proof. To prove the first inequality, let $\mathbf{u}: V(G) \to \mathbb{R}^d$ be an orthogonal representation of G in dimension $d = d_{\min}(G)$. Then $i \mapsto \mathbf{u}_i \circ \mathbf{u}_i$ is another orthogonal representation; this is in higher dimension, but ha the good "handle"

$$\mathbf{c} = \frac{1}{\sqrt{d}}(\mathbf{e}_1 \circ \mathbf{e}_1 + \cdots + \mathbf{e}_d \circ \mathbf{e}_d),$$

for which we have

$$\mathbf{c}^{\mathsf{T}}(\mathbf{u}_i \circ \mathbf{u}_i) = \frac{1}{\sqrt{d}} \sum_{j=1}^{d} (\mathbf{e}_j^{\mathsf{T}} \mathbf{u}_i)^2 = \frac{1}{\sqrt{d}} |\mathbf{u}_i|^2 = \frac{1}{\sqrt{d}},$$

which proves that $\vartheta(G) \le d$.

The upper bound in the second inequality is trivial; the lower bound follows from the fact that $\mathbb{R}^d$ can be colored by C^d colors for some constant C. $\square$

The following inequality follows by the same tensor product construction as Lemma **??**:

$$d_{\min}(G \boxtimes H) \le d_{\min}(G) d_{\min}(H). \tag{1}$$

It follows by Lemma 1.1 that we can use $d_{\min}(G)$ as an upper bound on $\Theta(G)$; however, it would not be better than $\vartheta(G)$. On the other hand, if we consider orthogonal representations over fields of finite characteristic, the dimension is an bound on the Shannon capacity (this follows from (1), which remains valid), and it may be better than ϑ [3, 1].

1.2 General position orthogonal representations

The first non-degeneracy condition we study is *general position*: we assume that any d of the representing vectors in $\mathbb{R}^d$ are linearly independent. A result of Lovász, Saks and Schrijver [6] finds an exact condition for this type of geometric representability.

Theorem 1.2 *A graph with n nodes has a general position orthogonal representation in $\mathbb{R}^d$ if and only if it is $(n-d)$-connected.*

The condition that the given set of representing vectors is in general position is not easy to check (it is NP-hard). A weaker, but very useful condition will be that the vectors representing the nodes nonadjacent to any node v are linearly independent. We say that such a representation is in *locally general position*.

It is almost trivial to see that every orthogonal representation that is in general position is in locally general position. For this it suffices to notice that every node i has at most $d-1$ non-neighbors; indeed, these are represented in a $(d-1)$-dimensional subspace (orthogonal to the vector representing i), and if they are linearly dependent, then some d of them are linearly dependent, contradicting the condition that the representation is in general position.

Theorem 1.2 is proved in the following slightly more general form:

Theorem 1.3 *If G is a graph with n nodes, then the following are equivalent:*

(i) *G has a general position orthogonal representation in $\mathbb{R}^d$;*

(ii) *G has a locally general position orthogonal representation in $\mathbb{R}^d$.*

(iii) *G is $(n-d)$-connected;*

We describe the proof in a couple of installments. First, to illustrate the connection between connectivity and orthogonal representations, we prove that (ii)$\Rightarrow$(iii). Let $\mathbf{x} : V \to \mathbb{R}^d$ be an orthogonal representation in locally general position. Let V_0 be a cutset of nodes of G, then $V = V_0 \cup V_1 \cup V_2$, where $V_1, V_2 \ne \emptyset$, and no edge connects V_1 and V_2. This implies that the vectors representing V_1 are linearly independent, and similarly, the vectors representing V_2 are linearly independent. Since the vectors representing V_1 and V_2 are mutually orthogonal, all vectors representing $V_1 \cup V_2$ are linearly independent. Hence $d \geq |V_1 \cup V_2| = n - |V_0|$, and so $|V_0| \geq n - d$.

The difficult part of the proof will be the construction of a general position orthogonal (or orthogonal) representation for $(n-d)$-connected graphs, and we describe and analyze the algorithm constructing the representation first. As a matter of fact, the following construction is almost trivial, the difficulty lies in the proof of its validity.

Let $\sigma = (1, ..., n)$ be any ordering of the nodes of $G = (V, E)$. Let us choose vectors $\mathbf{f}_1, \mathbf{f}_2, \ldots$ consecutively as follows. $\mathbf{f}_1$ is any vector of unit length. Suppose that $\mathbf{f}_i$ $(1 \leq i \leq j)$ are already chosen, then we choose $\mathbf{f}_{j+1}$ randomly, subject to the constraints that it has to be orthogonal to all previous vectors $\mathbf{f}_i$ for which $i \notin N(j+1)$. These orthogonality constraints restrict $\mathbf{f}_{j+1}$ to a linear subspace L_{j+1}, and we choose it from the uniform distribution over the unit sphere of L_{j+1}. Note that if G is $(n-d)$-connected, then every node of it has degree at least $n-d$, and hence

$$\dim L \geq d - \#\{i : \ i \leq j, \ \mathbf{v}_i\mathbf{v}_j \notin E\} \geq d - (d-1) = 1,$$

and so $\mathbf{f}_{j+1}$ can always be chosen.

This way we get a random mapping $\mathbf{f} : \ V \to S^{d-1}$, i.e., a probability distribution over $(S^{d-1})^V$, which we denote by μ_σ. We call $\mathbf{f}$ the *random sequential orthogonal representation* of G associated with the ordering $(1, ..., n)$, or shortly, a *sequential representation*. We are going to prove that this construction provides what we need:

Theorem 1.4 *Let G be an $(n-d)$-connected graph. Fix any ordering of its nodes and let f be the sequential representation of G. Then with probability 1, f is in general position.*

The sequential representation may depend on the initial ordering of the nodes. Let us consider a simple example.

Example 1.5 Let G have four nodes a, b, c and d and two edges ac and bd. Consider a sequential representation $\mathbf{f}$ in $\mathbb{R}^3$, associated with the given ordering. Since every node has one neighbor and two non-neighbors, we can always find a vector orthogonal to the earlier non-neighbors. The vectors $\mathbf{f}_b$ and $\mathbf{f}_b$ are orthogonal, and $\mathbf{f}_c$ is constrained to the plane $\mathbf{f}_b^\perp$; almost surely, $\mathbf{f}_c$ will not be parallel to $\mathbf{f}_a$, so together they span the plane $\mathbf{f}^\perp$. This means that $\mathbf{f}_d$, which must be orthogonal to $\mathbf{f}_a$ and $\mathbf{f}_c$, must be parallel to $\mathbf{f}_b$.

Now suppose that we choose $\mathbf{f}_c$ and $\mathbf{f}_d$ in the opposite order: then $\mathbf{f}_d$ will almost surely not be parallel to $\mathbf{f}_b$, but $\mathbf{f}_c$ will be forced to be parallel with $\mathbf{f}_a$. So not only are the two distributions $\mu_{(a,b,c,d)}$ and $\mu_{(a,b,d,c)}$ different, but an event, namely $\mathbf{f}_b \| \mathbf{f}_d$, occurs with probability 0 in one and probability 1 in the other.

Let us modify this example by connecting a and c by an edge. Then the planes $\mathbf{f}_a^\perp$ and $\mathbf{f}_b^\perp$ are not orthogonal any more (almost surely). Choosing $\mathbf{f}_c \in \mathbf{f}_b^\perp$ first still determines the direction of $\mathbf{f}_d$, but now it does not have to be parallel to $\mathbf{f}_b$; in fact, depending on $\mathbf{f}_c$, it can ba any unit vector in $\mathbf{f}_a^\perp$. The distributions $\mu_{(a,b,c,d)}$ and $\mu_{(a,b,d,c)}$ are still different, but any event that occurs with probability 0 in one will also occur with probability 0 in the other (this is not obvious; see Exercise 1.21 below).

This example motivates the following considerations. The distribution of a sequential representation may depend on the initial ordering of the nodes. The key to the proof will be that this dependence is not too strong. We say that two probability measures μ and ν on the same sigma-algebra S are *mutually absolute continuous*, if for any measurable subset A of S, $\mu(A) = 0$ if and only if $\nu(A) = 0$. The crucial step in the prof is the following lemma.

Lemma 1.6 *If G is $(n-d)$-connected, then for any two orderings σ and τ of V, the distributions μ_σ and μ_τ are mutually absolute continuous.*

Before proving this lemma, we have to state and prove a simple technical fact. For a subspace $A \subseteq \mathbb{R}^d$, we denote by $A^\perp$ its orthogonal complement. We need the elementary relations $(A^\perp)^\perp = A$ and $(A+B)^\perp = A^\perp + B^\perp$.

Lemma 1.7 *Let A, B and C be mutually orthogonal linear subspaces of $\mathbb{R}^d$ with $\dim(C) \geq 2$. Select a unit vector $\mathbf{a}_1$ uniformly from $A+C$, and then select a unit vector $\mathbf{b}_1$ uniformly from $(B+C)\cap\mathbf{a}_1^\perp$. Also, select a unit vector $\mathbf{b}_2$ uniformly from $B+C$, and then select a unit vector $\mathbf{a}_2$ uniformly from $(A+C)\cap\mathbf{b}_2^\perp$. Then the distributions of $(\mathbf{a}_1, \mathbf{b}_1)$ and $(\mathbf{a}_2, \mathbf{b}_2)$ are mutually absolute continuous.*

Proof. Let $r = \dim(A)$, $s = \dim(B)$ and $t = \dim(C)$. The special case when $r = 0$ or $s = 0$ is trivial. Suppose that $r, s \geq 1$ (by hypothesis, $t \geq 2$).

Observe that a unit vector $\mathbf{a}$ in $A+C$ can be written uniquely in the form $(\cos\theta)\mathbf{x} + (\sin\theta)\mathbf{y}$, where $\mathbf{x}$ is a unit vector in A, $\mathbf{y}$ is a unit vector in C, and $\theta \in [0, \pi/2]$. Uniform selection of $\mathbf{a}$ means independent uniform selection of $\mathbf{x}$ and $\mathbf{y}$, and an independent selection of θ from a distribution $\zeta_{r,t}$ that depends only on s and t. Using that $s, t \geq 1$, it is easy to see that $\zeta_{s,t}$ is mutually absolute continuous with respect to the uniform distribution on the interval $[0, \pi/2]$.

So the pair $(\mathbf{a}_1, \mathbf{b}_1)$ can be generated through five independent choices: a uniform unit vector $\mathbf{x}_1 \in A$, a uniform unit vector $\mathbf{z}_1 \in B$, a pair of orthogonal unit vectors $(\mathbf{y}_1, \mathbf{y}_2)$ selected from C (uniformly over all such pairs: this is possible since $\dim(C) \geq 2$), and two numbers θ_1 selected according to $\zeta_{s,t}$ and θ_2 is selected according to $\zeta_{r,t-1}$. The distribution of $(\mathbf{a}_2, \mathbf{b}_2)$ is described similarly except that θ_2 is selected according to $\zeta_{r,t}$ and θ_1 is selected according to $\zeta_{s,t-1}$.

Since $t \geq 2$, the distributions $\zeta_{r,t}$ and $\zeta_{r,t-1}$ are mutually absolute continuous and similarly, $\zeta_{s,t}$ and $\zeta_{s,t-1}$ are mutually absolute continuous, from which we deduce that the distributions of $(\mathbf{a}_1, \mathbf{b}_1)$ and $(\mathbf{a}_2, \mathbf{b}_2)$ are mutually absolute continuous. $\square$

Next, we prove our main Lemma.

Proof of Lemma 1.6. It suffices to prove that if τ is the ordering obtained from σ by swapping the nodes in positions j and $j+1$ $(1 \leq j \leq n-1)$, then μ_σ an μ_τ are mutually absolute continuous. Let us label the nodes so that $\sigma = (1, 2, \dots, n)$.

Let $\mathbf{f}$ and $\mathbf{g}$ be sequential representations from the distributions μ_σ and μ_τ. It suffices to prove that the distributions of $\mathbf{f}_1, \dots, \mathbf{f}_{j+1}$ and $\mathbf{g}_1, \dots, \mathbf{g}_{j+1}$ are mutually absolute continuous, since conditioned on any given assignment of vectors to $[j+1]$, the remaining vectors $\mathbf{f}_k$ and $\mathbf{g}_k$ are generated in the same way, and hence the distributions μ_σ and μ_τ are identical. Also, note that the distributions of $\mathbf{f}_1, \dots, \mathbf{f}_{j-1}$ and $\mathbf{g}_1, \dots, \mathbf{g}_{j-1}$ are identical.

We have several cases to consider.

Case 1. j and $j+1$ are adjacent. When conditioned on $\mathbf{f}_1, \dots, \mathbf{f}_{j-1}$, the vectors $\mathbf{f}_j$ and $\mathbf{f}_{j+1}$ are independently chosen, so it does not mater in which order they are selected.

Case 2. j and $j+1$ are not adjacent, but they are joined by a path that lies entirely in $[j+1]$. Let P be a shortest such path and t be its length (number of edges), so $2 \leq t \leq j$. We argue by induction on j and t. Let i be any internal node of P. We swap j and $j+1$ by the following steps (Figure 1):

(1) Interchange i and j, by successive adjacent swaps among the first j elements.

(2) Swap i and $j+1$.

(3) Interchange $j+1$ and j, by successive adjacent swaps among the first j elements.

(4) Swap j and i.

(5) Interchange $j+1$ and i, by successive adjacent swaps among the first j elements.

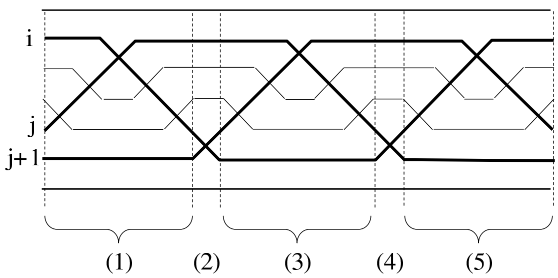

Figure 1: Interchanging j and $j+1$.

In each step, the new and the previous distributions of sequential representations are mutually absolute continuous: in steps (1), (3) and (5) this is so because the swaps take place place among the first j nodes, and in steps (2) and (4), because the nodes swapped are at a smaller distance than t in the graph distance.

Case 3. There is no path connecting j to $j+1$ in $[j+1]$. Let $\mathbf{x}_1, \dots, \mathbf{x}_{j-1}$ any selection of vectors for the first $j-1$ nodes. It suffices to show that the distributions of $(\mathbf{f}_j, \mathbf{f}_{j+1})$ and $(\mathbf{g}_j, \mathbf{g}_{j+1})$, conditioned on $\mathbf{f}_i = \mathbf{g}_i = \mathbf{x}_i$ $(i = 1, \dots, j-1)$, are mutually absolute continuous. Let $J = [j-1]$, $U_0 = N(j) \cap J$ and $U_1 = N(j+1) \cap J$. Then J has a partition $W_0 \cup W_1$ so that $U_0 \subseteq W_0$, $U_1 \subseteq W_1$, and there is no edge between W_0 and W_1. Furthermore, it follows that $V \setminus [j+1]$ is a cutset, whence $n-j-1 \geq n-d$ and so $j \leq d-1$.

For $S \subseteq J$, let $\mathrm{lin}(S)$ denote the linear subspace of $\mathbb{R}^d$ generated by the vectors $\mathbf{x}_i$, $i \in S$. Let

$$L = \mathrm{lin}(J), \quad L_0 = \mathrm{lin}(J \setminus U_0), \quad L_1 = \mathrm{lin}(J \setminus U_1).$$

Then $\mathbf{f}_j$ is selected uniformly from the unit sphere in $L_0^\perp$, and then $\mathbf{f}_{j+1}$ is selected uniformly from the unit sphere in $L_1^\perp \cap \mathbf{f}_j^\perp$. On the other hand, $\mathbf{g}_{j+1}$ is selected uniformly from the unit sphere in $L_1^\perp$, and then $\mathbf{g}_j$ is selected uniformly from the unit sphere in $L_0^\perp \cap \mathbf{g}_{j+1}^\perp$.

Let $A = L \cap L_0^\perp$, $B = L \cap L_1^\perp$ and $C = L^\perp$. We claim that A, B and C are mutually orthogonal. It is clear that $A \subseteq C^\perp = L$, so $A \perp C$, and similarly, $B \perp C$. Furthermore, $L_0 \supseteq \mathrm{lin}(W_1)$, and hence $L_0^\perp \subseteq \mathrm{lin}(W_1)^\perp$. So $A = L \cap L_0^\perp \subseteq L \cap \mathrm{lin}(W_1)^\perp = \mathrm{lin}(W_0)$. Since, similarly, $B \subseteq \mathrm{lin}(W_1)$, it follows that $A \perp B$.

Furthermore, we have $L_0 \subseteq L$, and hence L_0 and $L^\perp$ are orthogonal subspaces. This implies that $(L_0 + L^\perp) \cap L = L_0$, and hence $L_0^\perp = (L_0 + L^\perp)^\perp + L^\perp = (L_0^\perp \cap L) + L^\perp = A + C$. It follows similarly that $L_1^\perp = B + C$.

Finally, notice that $\dim(C) = d - \dim(L) \geq d - |J| \geq 2$. So Lemma 1.7 applies, which completes the proof. $\square$

Proof of Theorem 1.4. Observe that the probability that the first d vectors in a sequential representation are linearly dependent is 0. This event has then probability 0 in any other sequential representation. Since we can start the ordering with any d-tuple of nodes, it follows that the probability that the representation is not in general position is 0. $\square$

Proof of Theorem 1.3. We have seen the implications (i)⇒(ii) and (ii)⇒(iii), and Theorem 1.4 implies that (iii)⇒(i). $\square$

1.3 Faithful orthogonal representations

An orthogonal representation is *faithful* if different nodes are represented by non-parallel vectors and adjacent nodes are represented by non-orthogonal vectors.

We do not know how to determine the minimum dimension of a *faithful* orthogonal representation. It was proved by Maehara and Rödl [7] that if the maximum degree of the complementary graph $\overline{G}$ of a graph G is D, then G has a faithful orthogonal representation in $2D$ dimensions. They conjectured that the bound on the dimension can be improved to $D+1$. We show how to obtain their result from the results in Section 1.2, and that the conjecture is true if we strengthen its assumption by requiring that G is sufficiently connected.

Corollary 1.8 *Every $(n-d)$-connected graph on n nodes has a faithful general position orthogonal representation in $\mathbb{R}^d$.*

Proof. It suffices to show that in a sequential representation, the probability of the event that two nodes are represented by parallel vectors, or two adjacent nodes are represented by orthogonal vectors, is 0. By the Lemma 1.6, it suffices to prove this for the representation obtained from an ordering starting with these two nodes. But then the assertion is obvious. □

Using the elementary fact that a graph with minimum degree $n-D-1$ is at least $(n-2D)$-connected, we get the result of Maehara and Rödl:

Corollary 1.9 *If the maximum degree of the complementary graph $\overline{G}$ of a graph G is D, then G has a faithful orthogonal representation in $2D$ dimensions.*

1.4 Orthogonal representations with the Strong Arnold Property

We survey results about another, nontrivial and deep non-degeneracy condition, with proofs.
Strong Arnold Property. Consider an orthogonal representation $i \mapsto \mathbf{v}_i \in \mathbb{R}^d$ of a graph G. We can view this as a point in the $\mathbb{R}^{d|V|}$, satisfying the quadratic equations

$$\mathbf{v}_i^{\mathsf{T}} \mathbf{v}_j = 0 \qquad (ij \in \overline{E}). \tag{2}$$

Each of these equation defines a hypersurface in $\mathbb{R}^{d|V|}$. We say that the orthogonal representation $i \mapsto \mathbf{v}_i$ has the *Strong Arnold Property* if the surfaces (2) intersect transversally at this point. This means that their normal vectors are linearly independent.

This can be rephrased in more explicit terms as follows. For each nonadjacent pair $i, j \in V$, form the $d \times V$ matrix $V^{ij} = \mathbf{e}_i^{\mathsf{T}} \mathbf{v}_j + \mathbf{e}_j^{\mathsf{T}} \mathbf{v}_i$. Then the Strong Arnold Property says that the matrices V^{ij} are linearly independent.

Another way of saying this is that there is no symmetric $V \times V$ matrix $X \neq 0$ such that $X_{ij} = 0$ if $i = j$ or $ij \in E$, and

$$\sum_j X_{ij} \mathbf{v}_j = 0 \tag{3}$$

for every node i. Since (3) means a linear dependence between the non-neighbors of i, every orthogonal representation of a graph G in locally general position has the Strong Arnold Property. This shows that the Strong Arnold Property can be thought of as some sort of symmetrized version of locally general position. But the two conditions are not equivalent, as the following example shows.

Example 1.10 (Triangular grid) Consider the graph Δ_3 obtained by attaching a triangle on each edge of a triangle (Figure 2). This graph has an orthogonal representation $(\mathbf{v}_i : i = 1, \dots, 6)$ in $\mathbb{R}^3$: $\mathbf{v}_1$, $\mathbf{v}_2$ and $\mathbf{v}_3$ are mutually orthogonal unit vectors, and $\mathbf{v}_4 = \mathbf{v}_1 + \mathbf{v}_2$, $\mathbf{v}_5 = \mathbf{v}_1 + \mathbf{v}_3$, and $\mathbf{v}_6 = \mathbf{v}_2 + \mathbf{v}_3$.

This representation is not in locally general position, since the nodes non-adjacent to (say) node 1 are represented by linearly dependent vectors. But this representation satisfies the Strong Arnold Property. Suppose that a symmetric matrix X satisfies (3). Since node 4 is adjacent to all the other nodes except node 3, $X_{4,j} = 0$ for $j \neq 3$, and therefore case $i = 4$ of (3) implies that $X_{4,3} = 0$. By symmetry, $X_{3,4} = 0$, and hence case $i = 3$ of (3) simplifies to $X_{3,1}\mathbf{v}_1 + X_{3,2}\mathbf{v}_2 = 0$, which implies that $X_{3,j} = 0$ for all J. Going on similarly, we get that all entries of X must be zero.

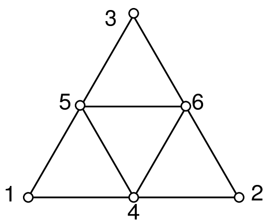

Figure 2: The graph Δ_3 with an orthogonal representation that has the strong Arnold property but is not in locally general position.

Algebraic width. Based on this definition, Colin de Verdière [2] introduced an interesting graph invariant related to connectivity (this is different from the better known "Colin de Verdière number" related to planarity). Let d be the smallest dimension in which G has a faithful orthogonal representation with the Strong Arnold Property, and define $\mathsf{w}_{\mathrm{alg}}(G) = n - d$. We call $\mathsf{w}_{\mathrm{alg}}(G)$ the *algebraic width* of the graph (the name refers to its connection with tree width, see below).

This definition is meaningful, since it is easy to construct a faithful orthogonal representation in $\mathbb{R}^n$ (where $n = |V(G)|$), in which the representing vectors are almost orthogonal and hence linearly independent, which implies that the Strong Arnold Property is also satisfied.

This definition can be rephrased in terms of matrices: consider a matrix $N \in \mathbb{R}^{V \times V}$ with the following properties:

(N1) $N_{ij} \begin{cases} = 0 & \text{if } ij \notin E,\ i \neq j;, \\ \neq 0 & \text{if } ij \in E. \end{cases}$

(N2) N is positive semidefinite;

(N3) [Strong Arnold Property] If X is a symmetric $n \times n$ matrix such that $X_{ij} = 0$ whenever $i = j$ or $ij \in E$, and $NX = 0$, then $X = 0$.

Lemma 1.11 *The algebraic width of a graph G is the maximum corank of a matrix with properties* (N1)–(N3).

Example 1.12 (Complete graphs) The Strong Arnold Property is void for complete graphs, and every representation is orthogonal, so we can use the same vector to represent every node. This shows that $\mathsf{w}_{\text{alg}}(K_n) = n-1$. For every noncomplete graph G, $\mathsf{w}_{\text{alg}}(G) \leq n-2$, since a faithful orthogonal representation requires at least two dimensions.

Example 1.13 (Edgeless graphs) To have a faithful orthogonal representation, all representing vectors must be mutually orthogonal, hence $\mathsf{w}_{\text{alg}}(\overline{K}_n) = 0$. It is not hard to see that every other graph G has $\mathsf{w}_{\text{alg}}(G) \geq 1$.

Example 1.14 (Paths) Every matrix N satisfying (N1) has an $(n-1)\times(n-1)$ nonsingular submatrix, and hence by Lemma 1.11, $\mathsf{w}_{\text{alg}}(P_n) \leq 1$. Since P_n is connected, we know that equality holds here.

Example 1.15 (Triangular grid II) To see a more interesting example, let us have a new look at the graph Δ_3 in Figure 2 (Example 1.10). Nodes 1, 2 and 3 must be represented by mutually orthogonal vectors, hence every faithful orthogonal representation of Δ_3 must have dimension at least 3. On the other hand, we have seen that Δ_3 has a faithful orthogonal representation with the Strong Arnold Property in $\mathbb{R}^3$. It follows that $\mathsf{w}_{\text{alg}}(\Delta_3) = 3$.

We continue with some easy bounds on the algebraic width. The condition that the representation must be faithful implies that the vectors representing a largest stable set of nodes must be mutually orthogonal, and hence the dimension of the representation is at least $\alpha(G)$. This implies that

$$\mathsf{w}_{\text{alg}}(G) \leq n-\alpha(G) = \tau(G). \tag{4}$$

By Theorem 1.3, every k-connected graph G has a faithful general position orthogonal representation in $\mathbb{R}^{n-k}$, and hence

$$\mathsf{w}_{\text{alg}}(G) \geq \kappa(G). \tag{5}$$

We may also use the Strong Arnold Property. There are $\binom{n}{2}-m$ orthogonality conditions, and in an optimal representation they involve $(n-\mathsf{w}_{\text{alg}}(G))n$ variables. If their normal vectors are linearly independent, then $\binom{n}{2}-m \leq (n-\mathsf{w}_{\text{alg}}(G))n$, and hence

$$\mathsf{w}_{\text{alg}}(G) \leq \frac{n+1}{2}+\frac{m}{n}. \tag{6}$$

The most important consequence of the Strong Arnold Property is the following.

Lemma 1.16 *The graph parameter $\mathsf{w}_{\text{alg}}(G)$ is minor-monotone.*

The parameter has other nice properties, of which the following will be relevant:

Lemma 1.17 *Let G be a graph, and let $B \subseteq V(G)$ induce a complete subgraph. Let $G_1, \dots, G_k$ be the connected components of $G \setminus B$, and let H_i be the subgraph induced by $V(G_i) \cup B$. Then*

$$\mathsf{w}_{\text{alg}}(G) = \max_i \mathsf{w}_{\text{alg}}(H_i).$$

Algebraic width, tree-width and connectivity. The *monotone connectivity* $\kappa_{\text{mon}}(G)$ of a graph G is defined as the maximum connectivity of any minor of G.

Tree-width is a parameter related to connectivity, introduced by Robertson and Seymour [8] as an important element in their graph minor theory. Colin de Verdière [2] defines a closely related parameter, which we call the *product-width* $\mathsf{w}_{\text{prod}}(G)$ of a graph G. This is the smallest positive integer r for which G is a minor of a Cartesian sum $K_r \Box T$, where T is a tree.

The difference between the tree-width $\mathsf{w}_{\text{tree}}(G)$ and product-width $\mathsf{w}_{\text{prod}}(G)$ of a graph G is at most 1:

$$\mathsf{w}_{\text{tree}}(G) \le \mathsf{w}_{\text{prod}}(G) \le \mathsf{w}_{\text{tree}}(G) + 1. \tag{7}$$

The lower bound was proved by Colin de Verdière, the upper, by van der Holst [4]. It is easy to see that $\kappa_{\text{mon}}(G) \le \mathsf{w}_{\text{tree}}(G) \le \mathsf{w}_{\text{prod}}(G)$. The parameter $\mathsf{w}_{\text{prod}}(G)$ is clearly minor-monotone.

The algebraic width is sandwiched between two of these parameters:

Theorem 1.18 *For every graph G,*

$$\kappa_{\text{mon}}(G) \le \mathsf{w}_{\text{alg}}(G) \le \mathsf{w}_{\text{prod}}(G).$$

The upper bound was proved by Colin de Verdière [2], while the lower bound follows easily from the results in Section 1.2.

For small values, equality holds in Theorem 1.18; this was proved by Van der Holst [4] and Kotlov [5].

Proposition 1.19 *If $\mathsf{w}_{\text{alg}}(G) \le 2$, then $\kappa_{\text{mon}}(G) = \mathsf{w}_{\text{alg}}(G) = \mathsf{w}_{\text{prod}}(G)$.*

A planar graph has $\kappa_{\text{mon}}(G) \le 5$ (since every simple minor of it has a node with degree at most 5). Since planar graphs can have arbitrarily large treewidth, the lower bound in Theorem 1.18 can be very far from equality. The following example of Kotlov shows that in general, equality does not hold in the upper bound in Theorem 1.18 either. It is not known whether $\mathsf{w}_{\text{prod}}(G)$ can be bounded from above by some function $\mathsf{w}_{\text{alg}}(G)$.

Example 1.20 The k-cube Q^k has $\mathsf{w}_{\text{alg}}(Q^k) = O(2^{k/2})$ but $\mathsf{w}_{\text{prod}}(Q^k) = \Theta(2^k)$. The complete graph $K_{2^{m+1}}$ is a minor of Q^{2m+1} (see Exercise 1.22).

Exercise 1.21 Let Σ and Δ be two planes in $\mathbb{R}^3$ that are neither parallel nor weakly orthogonal. Select a unit vector $\mathbf{a}_1$ uniformly from Σ, and a unit vector $\mathbf{b}_1 \in \Delta \cap \mathbf{a}_1^\perp$. Let the unit vectors $\mathbf{a}_2$ and $\mathbf{b}_2$ be defined similarly, but selecting $\mathbf{b}_2 \in \Delta$ first (uniformly over all unit vectors), and then $\mathbf{a}_2$ from $\Sigma \cap \mathbf{b}_2^\perp$. Prove that the distributions of $(\mathbf{a}_1, \mathbf{b}_1)$ and $(\mathbf{a}_2, \mathbf{b}_2)$ are different, but mutually absolute continuous.

Exercise 1.22 (a) For every bipartite graph G, the graph $G \boxtimes K_2$ is a minor of $G \Box C_4$. (b) K_{2m+1} is a minor of Q^{2m+1}.

References

[1] N. Alon: The Shannon capacity of a union, *Combinatorica* **18** (1998), 301–310.

[2] Y. Colin de Verdière: Multiplicities of eigenvalues and tree-width of graphs, *J. Combin. Theory Ser. B* **74** (1998), 121–146.

[3] W. Haemers: On some problems of Lovász concerning the Shannon capacity of a graph, *IEEE Trans. Inform. Theory* **25** (1979), 231–232.

[4] H. van der Holst: *Topological and Spectral Graph Characterizations*, Ph.D. Thesis, University of Amsterdam, Amsterdam, 1996.

[5] A. Kotlov: A spectral characterization of tree-width-two graphs, *Combinatorica* **20** (2000)

[6] L. Lovász, M. Saks, A. Schrijver: Orthogonal representations and connectivity of graphs, *Linear Alg. Appl.* **114/115** (1989), 439–454; A Correction: Orthogonal representations and connectivity of graphs, *Linear Alg. Appl.* **313** (2000) 101–105.

[7] H. Maehara, V. Rödl: On the dimension to represent a graph by a unit distance graph, *Graphs and Combinatorics* **6** (1990), 365–367.

[8] N. Robertson, P.D. Seymour: Graph minors III: Planar tree-width, *J. Combin. Theory B* **36** (1984), 49–64.