The Gram Dimension of a Graph

Monique Laurent[1,2] and Antonios Varvitsiotis[1]

[1] Centrum Wiskunde & Informatica (CWI), Amsterdam, The Netherlands
{M.Laurent,A.Varvitsiotis}@cwi.nl
[2] Tilburg University, The Netherlands

Abstract. The Gram dimension $\mathrm{gd}(G)$ of a graph is the smallest integer $k \geq 1$ such that, for every assignment of unit vectors to the nodes of the graph, there exists another assignment of unit vectors lying in $\mathbb{R}^k$, having the same inner products on the edges of the graph. The class of graphs satisfying $\mathrm{gd}(G) \leq k$ is minor closed for fixed k, so it can characterized by a finite list of forbidden minors. For $k \leq 3$, the only forbidden minor is K_{k+1}. We show that a graph has Gram dimension at most 4 if and only if it does not have K_5 and $K_{2,2,2}$ as minors. We also show some close connections to the notion of d-realizability of graphs. In particular, our result implies the characterization of 3-realizable graphs of Belk and Connelly [5,6].

1 Introduction

The problem of completing a given partial matrix (where only a subset of entries are specified) to a full positive semidefinite (psd) matrix is one of the most extensively studied matrix completion problems. A particular instance is the completion problem for correlation matrices arising in probability and statistics, and it is also closely related to the completion problem for Euclidean distance matrices with applications, e.g., to sensor network localization and molecular conformation in chemistry. We refer, e.g., to [8,14] and further references therein for additional details.

An important feature of a matrix is its rank which intuitively can be seen as a measure of complexity of the data it represents. As an example, the minimum embedding dimension of a finite metric space can be expressed as the rank of an appropriate matrix [8]. Another problem of interest is to compute low rank solutions to semidefinite programs as they may lead to improved approximations to the underlying discrete optimization problem [2]. Consequently, the problem of computing (approximate) matrix completions is of fundamental importance in many disciplines and it has been extensively studied (see, e.g., [7,20]).

This motivates the following question which we study in this paper: Given a partially specified matrix which admits at least one psd completion, provide guarantees for the existence of small rank psd completions.

Evidently, the (non)existence of small rank completions depends on the values of the prescribed entries of the partial matrix. We approach this problem from a combinatorial point of view and give an answer in terms of the combinatorial

A.R. Mahjoub et al. (Eds.): ISCO 2012, LNCS 7422, pp. 356–367, 2012.

structure of the problem, which is captured by the *Gram dimension* of the graph. Before we give the precise definition, we introduce some notation.

Throughout $\mathcal{S}^n$ denotes the set of symmetric $n \times n$ matrices and $\mathcal{S}^n_+$ (resp., $\mathcal{S}^n_{++}$) is the subset of all positive semidefinite (resp., positive definite) matrices. For a matrix X the notation $X \succeq 0$ means that X is psd. Given a graph $G = (V = [n], E)$, its edges are denoted as (unordered) pairs (i, j) and, for convenience, we will sometimes identify V with the set of all diagonal pairs, i.e., we set $V = \{(i, i) \mid i \in [n]\}$. Moreover, π_{VE} denotes the projection from $\mathcal{S}^n$ onto the subspace $\mathbb{R}^{V \cup E}$ indexed by the diagonal entries and the edges of G.

Definition 1. *The Gram dimension* $\mathrm{gd}(G)$ *of a graph* $G = ([n], E)$ *is the smallest integer* $k \geq 1$ *such that, for any matrix* $X \in \mathcal{S}^n_+$*, there exists another matrix* $X' \in \mathcal{S}^n_+$ *with rank at most* k *and such that* $\pi_{VE}(X) = \pi_{VE}(X')$.

Given a graph $G = ([n], E)$, a partial G-matrix is a partial $n \times n$ matrix whose entries are specified on the diagonal and at positions corresponding to edges of G. Then, if a partial G-matrix admits a psd completion, it also has one of rank at most $\mathrm{gd}(G)$. This motivates the study of bounds for $\mathrm{gd}(G)$.

As we will see in Section 2, the class of graphs with $\mathrm{gd}(G) \leq k$ is closed under taking minors for any fixed k, hence it can be characterized in terms of a finite list of forbidden minors. Our main result is such a characterization for $k \leq 4$.

Main Theorem. For $k \leq 3$, a graph G has $\mathrm{gd}(G) \leq k$ if and only if it has no K_{k+1} minor. For $k = 4$, a graph G has $\mathrm{gd}(G) \leq 4$ if and only if it has no K_5 and $K_{2,2,2}$ minors.

An equivalent way of rephrasing the notion of Gram dimension is in terms of ranks of feasible solutions to semidefinite programs. Indeed, the Gram dimension of a graph $G = (V, E)$ is at most k if and only if the set

$$S(G, a) = \{X \succeq 0 \mid X_{ij} = a_{ij} \text{ for } ij \in V \cup E\}$$

contains a matrix of rank at most k for all $a \in \mathbb{R}^{V \cup E}$ for which $S(G, a)$ is not empty. The set $S(G, a)$ is a typical instance of spectrahedron. Recall that a *spectrahedron* is the convex region defined as the intersection of the positive semidefinite cone with a finite set of linear subspaces, i.e., the feasibility region of a semidefinite program in canonical form:

$$\max \langle A_0, X\rangle \text{ subject to } \langle A_j, X\rangle = b_j,\ (j = 1, \ldots, m), \qquad X \succeq 0. \qquad (1)$$

If the feasibility region of (1) is not empty, it follows from well known geometric results that it contains a matrix X of rank k satisfying $\binom{k+1}{2} \leq m$, that is, $k \leq \lfloor \frac{\sqrt{8m+1}-1}{2} \rfloor$ (see [4]). Applying this to the spectrahedron $S(G, a)$, we obtain the bound $\mathrm{gd}(G) = O(\sqrt{|V| + |E|})$, which is however weak in general.

As an application, the Gram dimension can be used to bound the rank of optimal solutions to semidefinite programs. Indeed consider a semidefinite program in canonical form (1). Its *aggregated sparsity pattern* is the graph G with node set $[n]$ and whose edges are the pairs corresponding to the positions where at least one of the matrices A_j ($j \geq 0$) has a nonzero entry. Then, whenever (1) attains its maximum, it admits an optimal solution of rank at most $\text{gd}(G)$. Results ensuring existence of low rank solutions are important, in particular, for approximation algorithms. Indeed semidefinite programs are widely used as convex tractable relaxations to hard combinatorial problems. Then the rank one solutions typically correspond to the desired optimal solutions of the discrete problem and low rank solutions can lead to improved performance guarantees (see e.g. the result of [2] for max-cut).

As an illustration, consider the max-cut problem for graph G and its standard semidefinite programming relaxation:

$$\max \frac{1}{4}\langle L, X\rangle \text{ subject to } X_{ii} = 1 \ (i = 1, \ldots, n), \quad X \succeq 0, \tag{2}$$

where L denotes the Laplacian matrix of G. Clearly, the aggregated sparsity pattern of program (2) is equal to G. In particular, our main Theorem implies that if G does not have K_5 and $K_{2,2,2}$ minors, then program (2) has an optimal solution of rank at most four. Of course, this is not of great interest since for K_5 minor free graphs, the max-cut problem can be solved in polynomial time ([3]).

In a similar flavor, for a graph $G = ([n], E)$ and $w \in \mathbb{R}_+^{V \cup E}$, the problem of computing bounded rank solutions to semidefinite programs of the form

$$\max \sum_{i=1}^{n} w_i X_{ii} \ \text{ s.t. } \sum_{i,j=1}^{n} w_i w_j X_{ij} = 0, \ X_{ii} + X_{ij} - 2X_{ij} \leq w_{ij} \ (ij \in E), \ X \succeq 0,$$

has been studied in [12]. In particular, it is shown in [12] that there always exists an optimal solution of rank at most the tree-width of G plus 1. There are numerous other results related to geometric representations of graphs; we refer, e.g., to [13,17,18] for further results and references.

Yet another, more geometrical, way of interpreting the Gram dimension is in terms of graph embeddings in the spherical metric space. For this, consider the unit sphere $\mathbf{S}^{k-1} = \{x \in \mathbb{R}^k \mid \|x\| = 1\}$, equipped with the distance

$$d_{\mathbf{S}}(x, y) = \arccos(x^T y) \ \text{ for } x, y \in \mathbf{S}^{k-1}.$$

Here, $\|x\|$ denotes the usual Euclidean norm. Then $(\mathbf{S}^{k-1}, d_{\mathbf{S}})$ is a metric space, known as the *spherical metric space*. A graph $G = ([n], E)$ has Gram dimension at most k if and only if, for any assignment of vectors $p_1, \ldots, p_n \in \mathbf{S}^d$ (for some $d \geq 1$), there exists another assignment $q_1, \ldots, q_n \in \mathbf{S}^{k-1}$ such that

$$d_{\mathbf{S}}(p_i, p_j) = d_{\mathbf{S}}(q_i, q_j), \text{ for } ij \in E.$$

In other words, this is the question of deciding whether a partial matrix can be embedded in the $(k-1)$-dimensional spherical space. The analogous question for the Euclidean metric space $(\mathbb{R}^k, \|\cdot\|)$ has been extensively studied. In particular, Belk and Connelly [5,6] show the following result for the graph parameter $\text{ed}(G)$, the analogue of $\text{gd}(G)$ for Euclidean embeddings, introduced in Definition 5.

Theorem 1. *For $k \leq 2$, $\mathrm{ed}(G) \leq k$ if and only if G has no K_{k+2} minor. For $k = 3$, $\mathrm{ed}(G) \leq 3$ if and only if G does not have K_5 and $K_{2,2,2}$ minors.*

There is a striking similarity between our main Theorem and Theorem 1 above. This is no coincidence, since these two parameters are very closely related as we will see in Section 5.

The paper is orgranized as follows. In Section 2 we give definitions and establish some basic properties of the graph parameter $\mathrm{gd}(G)$. In Section 3 we sketch the proof of our main Theorem. In Section 4 we show how we can use semidefinite programming in order to prove that $\mathrm{gd}(V_8)$ and $\mathrm{gd}(C_5 \times C_2)$ are both at most four. In Section 5 we will elaborate between the similarities and differences between the two graph parameters $\mathrm{gd}(G)$ and $\mathrm{ed}(G)$. Section 6 discusses the complexity of the natural decision problem associated with the graph parameter $\mathrm{gd}(G)$.

Note. The extended version of this paper is available at [15]. Complexity issues associated with the parameter $\mathrm{gd}(G)$ are further discussed in [10].

2 Basic Definitions and Properties

For a graph $G = (V = [n], E)$ let $\mathcal{S}_+(G) = \pi_{VE}(\mathcal{S}_+^n) \subseteq \mathbb{R}^{V \cup E}$ denote the projection of the positive semidefinite cone onto $\mathbb{R}^{V \cup E}$, whose elements can be seen as the partial G-matrices that can be completed to a psd matrix. Let $\mathcal{E}_n$ denote the set of matrices in $\mathcal{S}_+^n$ with an all-ones diagonal (aka the correlation matrices), and let $\mathcal{E}(G) = \pi_E(\mathcal{E}_n) \subseteq \mathbb{R}^E$ denote its projection onto the edge subspace $\mathbb{R}^E$, known as the *elliptope* of G; we only project on the edge set since all diagonal entries are implicitly known and equal to one for matrices in $\mathcal{E}_n$.

Definition 2. *Given a graph $G = (V, E)$ and a vector $a \in \mathbb{R}^{V \cup E}$, a Gram representation of a in $\mathbb{R}^k$ consists of a set of vectors $p_1, \ldots, p_n \in \mathbb{R}^k$ such that*

$$p_i^T p_j = a_{ij} \ \forall ij \in V \cup E.$$

The Gram dimension of $a \in \mathcal{S}_+(G)$, denoted as $\mathrm{gd}(G, a)$, is the smallest integer k for which a has a Gram representation in $\mathbb{R}^k$.

Definition 3. *The Gram dimension of a graph $G = (V, E)$ is defined as*

$$\mathrm{gd}(G) = \max_{a \in \mathcal{S}_+(G)} \mathrm{gd}(G, a). \tag{3}$$

We denote by $\mathcal{G}_k$ the class of graphs for which $\mathrm{gd}(G) \leq k$. Clearly, the maximization in (3) can be restricted to be taken over all $a \in \mathcal{E}(G)$ (where all diagonal entries are implicitly taken to be equal to 1).

We now investigate the behavior of the graph parameter $\mathrm{gd}(G)$ under some simple graph operations.

Lemma 1. *The graph parameter $\mathrm{gd}(G)$ is monotone nondecreasing with respect to edge deletion and contraction. That is, if H is a minor of G (denoted as $H \preceq G$), then $\mathrm{gd}(H) \leq \mathrm{gd}(G)$.*

Proof. Let $G = ([n], E)$ and $e \in E$. It is clear that $\text{gd}(G\backslash e) \le \text{gd}(G)$. We show that $\text{gd}(G/e) \le \text{gd}(G)$. Say e is the edge $(1, n)$ and $G/e = ([n-1], E')$. Consider $X \in \mathcal{S}_+^{n-1}$; we show that there exists $X' \in \mathcal{S}_+^{n-1}$ with rank at most $k = \text{gd}(G)$ and such that $\pi_{E'}(X) = \pi_{E'}(X')$. For this, extend X to the matrix $Y \in \mathcal{S}_+^n$ defined by $Y_{nn} = X_{11}$ and $Y_{in} = X_{1i}$ for $i \in [n-1]$. By assumption, there exists $Y' \in \mathcal{S}_+^n$ with rank at most k such that $\pi_E(Y) = \pi_E(Y')$. Hence $Y'_{1i} = Y'_{ni}$ for all $i \in [n]$, so that the principal submatrix X' of Y' indexed by $[n-1]$ has rank at most k and satisfies $\pi_{E'}(X') = \pi_{E'}(X)$. □

Let $G_1 = (V_1, E_1)$, $G_2 = (V_2, E_2)$ be two graphs, where $V_1 \cap V_2$ is a clique in both G_1 and G_2. Their *clique sum* is the graph $G = (V_1 \cup V_2, E_1 \cup E_2)$, also called their *clique k-sum* when $k = |V_1 \cap V_2|$. The following result follows from well known arguments (used already, e.g., in [11]).

Lemma 2. *If G is the clique sum of two graphs G_1 and G_2, then*

$$\text{gd}(G) = \max\{\text{gd}(G_1), \text{gd}(G_2)\}.$$

As a direct application, one can bound the Gram dimension of partial k-trees. Recall that a graph G is a *k-tree* if it is a clique k-sum of copies of K_{k+1} and a *partial k-tree* if it is a subgraph of a k-tree (equivalently, G has tree-width k). Partial 1-trees are exactly the forests and partial 2-trees (aka series-parallel graphs) are the graphs with no K_4 minor (see [9]).

Lemma 3. *If G is a partial k-tree then* $\text{gd}(G) \le k + 1$.

For example, for the complete graph K_n, $\text{gd}(K_n) = n$, and $\text{gd}(K_n \setminus e) = n - 1$ for any edge e of K_n. Moreover, for the complete bipartite graph $K_{n,m}$ $(n \le m)$, $\text{gd}(K_{n,m}) = n + 1$ (since $K_{n,m}$ is a partial n-tree and contains a K_{n+1} minor).

In view of Lemma 1, the class $\mathcal{G}_k$ of graphs with Gram dimension at most k is closed under taking minors. Hence, by the celebrated graph minor theorem [21], it can be characterized by finitely many minimal forbidden minors. The simple properties we just established suffice to characterize $\mathcal{G}_k$, for $k \le 3$.

Theorem 2. *For $k \le 3$,* $\text{gd}(G) \le k$ *if and only if G has no minor K_{k+1}.*

The next natural question is to characterize the graphs with Gram dimension at most 4, which we address in the next section.

3 Characterizing Graphs with Gram Dimension at Most 4

In this section we characterize the class of graphs with Gram dimension at most 4. Clearly, K_5 is a minimal forbidden minor for $\mathcal{G}_4$. We now show that this is also the case for the complete tripartite graph $K_{2,2,2}$.

Lemma 4. *The graph $K_{2,2,2}$ is a minimal forbidden minor for $\mathcal{G}_4$.*

Proof. First we construct $x \in \mathcal{E}(K_{2,2,2})$ with $\text{gd}(K_{2,2,2}, x) = 5$. For this, let $K_{2,2,2}$ be obtained from K_6 by deleting the edges $(1, 4)$, $(2, 5)$ and $(3, 6)$. Let $e_1, \ldots, e_5$ denote the standard unit vectors in $\mathbb{R}^5$, let X be the Gram matrix of the vectors e_1, e_2, e_3, e_4, e_5 and $(e_1 + e_2)/\sqrt{2}$ labeling the nodes $1, \ldots, 6$, respectively, and let $x \in \mathcal{E}(K_{2,2,2})$ be the projection of X. We now verify that X is the unique psd completion of x which shows that $\text{gd}(K_{2,2,2}) \geq 5$. Indeed the chosen Gram labeling of the matrix X implies the following linear dependency: $C_6 = (C_4 + C_5)/\sqrt{2}$ among its columns C_4, C_5, C_6 indexed respectively by $4, 5, 6$; this implies that the unspecified entries X_{14}, X_{25}, X_{36} are uniquely determined in terms of the specified entries of X.

On the other hand, one can easily verify that $K_{2,2,2}$ is a partial 4-tree, thus $\text{gd}(K_{2,2,2}) \leq 5$. Moreover, deleting or contracting an edge in $K_{2,2,2}$ yields a partial 3-tree, thus with Gram dimension at most 4. □

By Lemma 3 we know that all partial 3-trees belong to $\mathcal{G}_4$. Moreover, it is known that partial 3-trees can be characterized in terms of four forbidden minors as stated below.

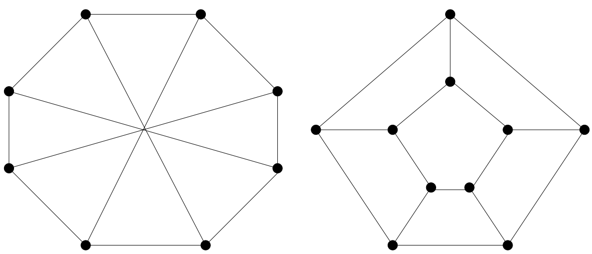

Fig. 1. The graphs V_8 and $C_5 \times C_2$

Theorem 3. *[1] A graph G is a partial 3-tree if and only if G does not have $K_5, K_{2,2,2}, V_8$ and $C_5 \times C_2$ as a minor.*

The graphs V_8 and $C_5 \times C_2$ are shown in Figure 1. The forbidden minors for partial 3-trees are natural candidates for being obstructions to the class $\mathcal{G}_4$. We have already seen that for K_5 and $K_{2,2,2}$ this is indeed the case. However, this is not the true for V_8 and $C_5 \times C_2$. Indeed, in the extended version of the paper, it is proven that $\text{gd}(V_8) = \text{gd}(C_5 \times C_2) = 4$ [15]. Using this, we can now complete our characterization of the class $\mathcal{G}_4$.

Theorem 4. *For a graph G, $\text{gd}(G) \leq 4$, if and only if G does not have K_5 or $K_{2,2,2}$ as a minor.*

Proof. The 'only if' part follows from Lemmas 1 and 4. The 'if part' follows from the fact that $\text{gd}(V_8) = \text{gd}(C_5 \times C_2) = 4$ and Lemmas 1, 2, combined with the following graph theoretical result, shown in [6]: If G is a graph with no K_5, $K_{2,2,2}$ minors, then G is a subgraph of a clique sum of copies of K_4, V_8 and $C_5 \times C_2$. □

4 Using Semidefinite Programming

In this section we sketch the approach which we will follow in order to bound the Gram dimension of the two graphs V_8 and $C_5 \times C_2$.

Definition 4. *Given a graph $G = (V = [n], E)$, a configuration of G is an assignment of vectors $p_1, \ldots, p_n \in \mathbb{R}^k$ (for some $k \geq 1$) to the nodes of G; the pair $(G, \mathbf{p})$ is called a framework, where we use the notation $\mathbf{p} = \{p_1, \ldots, p_n\}$. Two configurations $\mathbf{p}, \mathbf{q}$ of G (not necessarily lying in the same space) are said to be equivalent if $p_i^T p_j = q_i^T q_j$ for all $ij \in V \cup E$.*

Our objective is to show that the two graphs $G = V_8$, $C_5 \times C_2$ belong to $\mathcal{G}_4$. That is, we must show that, given any $a \in \mathcal{S}_+(G)$, one can construct a Gram representation $\mathbf{q}$ of (G, a) lying in the space $\mathbb{R}^4$.

Along the lines of [5] (which deals with Euclidean distance realizations), our strategy to achieve this is as follows: First, we select an initial Gram representation $\mathbf{p}$ of (G, a) obtained by 'stretching' as much as possible along a given pair (i_0, j_0) which is not an edge of G; more precisely, $\mathbf{p}$ is a representation of (G, a) which maximizes the inner product $p_{i_0}^T p_{j_0}$. As suggested in [24] (in the context of Euclidean distance realizations), this configuration $\mathbf{p}$ can be obtained by solving a semidefinite program; then $\mathbf{p}$ corresponds to the Gram representation of an optimal solution X to this program.

In general we cannot yet claim that $\mathbf{p}$ lies in $\mathbb{R}^4$. However, we can derive useful information about $\mathbf{p}$ by using an optimal solution Ω (which will correspond to a 'stress matrix') to the dual semidefinite program. Indeed, the optimality condition $X\Omega = 0$ will imply some linear dependencies among the p_i's that can be used to show the existence of an equivalent representation $\mathbf{q}$ of (G, a) in low dimension. Roughly speaking, most often, these dependencies will force the majority of the p_i's to lie in $\mathbb{R}^4$, and one will be able to rotate each remaining vector p_j about the space spanned by the vectors labeling the neighbors of j into $\mathbb{R}^4$. Showing that the initial representation $\mathbf{p}$ can indeed be 'folded' into $\mathbb{R}^4$ as just described makes up the main body of the proof.

We now sketch how to model the 'stretching' procedure using semidefinite programming and how to obtain a 'stress matrix' via semidefinite programming duality.

Let $G = (V = [n], E)$ be a graph and let $e_0 = (i_0, j_0)$ be a non-edge of G (i.e., $i_0 \neq j_0$ and $e_0 \notin E$). Let $a \in \mathcal{S}_{++}(G)$ be a partial positive semidefinite matrix for which we want to show the existence of a Gram representation in a small dimensional space. For this consider the semidefinite program:

$$\max \langle E_{i_0 j_0}, X \rangle \ \text{ such that } \langle E_{ij}, X \rangle = a_{ij} \ (ij \in V \cup E), \quad X \succeq 0, \tag{4}$$

where $E_{ij} = (e_i e_j^T + e_j e_i^T)/2$ and $e_1, \ldots, e_n$ are the standard unit vectors in $\mathbb{R}^n$. The dual semidefinite program of (4) reads:

$$\min \sum_{ij \in V \cup E} w_{ij} a_{ij} \text{ such that } \Omega = \sum_{ij \in V \cup E} w_{ij} E_{ij} - E_{i_0 j_0} \succeq 0. \tag{5}$$

As the program (5) is strictly feasible, there is no duality gap and the optimal values are attained in both programs. Consider now a pair (X, Ω) of primal-dual optimal solutions. Then (X, Ω) satisfies the optimality condition, i.e., $X\Omega = 0$. This condition can be reformulated as

$$w_{ii}p_i + \sum_{j|ij\in E\cup\{e_0\}} w_{ij}p_j = 0 \text{ for all } i \in [n], \tag{6}$$

where $\Omega = (w_{ij})$ and $X = \text{Gram}(p_1, \ldots, p_n)$. Using the local information provided by the 'equilibrium' conditions (6) about the configuration $\mathbf{p}$ and examining all possible cases for the support of the stress matrix, one can construct an equivalent configurations in $\mathbb{R}^4$ for the graphs V_8 and $C_5 \times C_2$.

For the full proof the reader is referred to the extended version of the paper [15].

5 Links to Euclidean Graph Realizations

In this section we investigate the links between the notion of Gram dimension and graph realizations in Euclidean spaces which will, in particular, enable us to relate our result from Theorem 4 to the result of Belk and Connelly (Theorem 1).

Recall that a matrix $D = (d_{ij}) \in \mathcal{S}^n$ is a *Euclidean distance matrix* (EDM) if there exist vectors $p_1, \ldots, p_n \in \mathbb{R}^k$ (for some $k \geq 1$) such that $d_{ij} = \|p_i - p_j\|^2$ for all $i, j \in [n]$. Then EDM_n denotes the cone of all $n \times n$ Euclidean distance matrices and, for a graph $G = ([n], E)$, $\text{EDM}(G) = \pi_E(\text{EDM}_n)$ is the set of partial G-matrices that can be completed to a Euclidean distance matrix.

Definition 5. *Given a graph $G = ([n], E)$ and $d \in \mathbb{R}^E_+$, a Euclidean (distance) representation of d in $\mathbb{R}^k$ consists of a set of vectors $p_1, \ldots, p_n \in \mathbb{R}^k$ such that*

$$\|p_i - p_j\|^2 = d_{ij} \;\; \forall ij \in E.$$

Then, $\text{ed}(G, d)$ *is the smallest integer k for which d has a Euclidean representation in $\mathbb{R}^k$ and the graph parameter* $\text{ed}(G)$ *is defined as*

$$\text{ed}(G) = \max_{d\in\text{EDM}(G)} \text{ed}(G, d). \tag{7}$$

There is a well known correspondence between psd and EDM completions (for details and references see, e.g., [8]). Namely, for a graph G, let ∇G denote its *suspension graph*, obtained by adding a new node (the *apex* node, denoted by 0), adjacent to all nodes of G. Consider the one-to-one map $\phi : \mathbb{R}^{V\cup E(G)} \mapsto \mathbb{R}^{E(\nabla G)}_+$, which maps $x \in \mathbb{R}^{V\cup E(G)}$ to $d = \phi(x) \in \mathbb{R}^{E(\nabla G)}_+$ defined by

$$d_{0i} = x_{ii} \;\; (i \in [n]), \quad d_{ij} = x_{ii} + x_{jj} - 2x_{ij} \;\; (ij \in E(G)).$$

Then the vectors $u_1, \ldots, u_n \in \mathbb{R}^k$ form a Gram representation of x if and only if the vectors $u_0 = 0, u_1, \ldots, u_n$ form a Euclidean representation of $d = \phi(x)$ in $\mathbb{R}^k$. This shows:

Lemma 5. *Let* $G = (V, E)$ *be a graph. Then,* $\mathrm{gd}(G, x) = \mathrm{ed}(\nabla G, \phi(x))$ *for any* $x \in \mathbb{R}^{V \cup E}$ *and thus* $\mathrm{gd}(G) = \mathrm{ed}(\nabla G)$.

For the Gram dimension of a graph one can show the following property:

Lemma 6. *Consider a graph* $G = ([n], E)$ *and let* $\nabla G = ([n] \cup \{0\}, E \cup F)$, *where* $F = \{(0, i) \mid i \in [n]\}$. *Given* $x \in \mathbb{R}^E$, *its* 0-extension *is the vector* $y = (x, 0) \in \mathbb{R}^{E \cup F}$. *If* $x \in \mathcal{S}_+(G)$, *then* $y \in \mathcal{S}_+(\nabla G)$ *and* $\mathrm{gd}(G, x) = \mathrm{gd}(\nabla G, y)$. *Moreover,* $\mathrm{gd}(\nabla G) = \mathrm{gd}(G) + 1$.

Proof. The first part is clear and implies $\mathrm{gd}(\nabla G) \geq \mathrm{gd}(G) + 1$. Set $k = \mathrm{gd}(G)$; we show the reverse inequality $\mathrm{gd}(\nabla G) \leq k + 1$. For this, let $X \in \mathcal{S}_+^{n+1}$, written in block-form as $X = \begin{pmatrix} \alpha & a^T \\ a & A \end{pmatrix}$, where $A \in \mathcal{S}_+^n$ and the first row/column is indexed by the apex node 0 of ∇G. If $\alpha = 0$ then $a = 0$, $\pi_{VE}(A)$ has a Gram representation in $\mathbb{R}^r$ and thus $\pi_{V(\nabla G)E(\nabla G)}(X)$ too. Assume now $\alpha > 0$ and without loss of generality $\alpha = 1$. Consider the Schur complement Y of X with respect to the entry $\alpha = 1$, given by $Y = A - aa^T$. As $\mathrm{gd}(G) = k$, there exists $Z \in \mathcal{S}_+^n$ such that $\mathrm{rank}(Z) \leq k$ and $\pi_{VE}(Z) = \pi_{VE}(Y)$. Define the matrix

$$X' := \begin{pmatrix} 1 & a^T \\ a & aa^T \end{pmatrix} + \begin{pmatrix} 0 & 0 \\ 0 & Z \end{pmatrix}.$$

Then, $\mathrm{rank}(X') = \mathrm{rank}(Z) + 1 \leq k + 1$. Moreover, X' and X coincide at all diagonal entries as well as at all entries corresponding to edges of ∇G. This concludes the proof that $\mathrm{gd}(\nabla G) \leq k + 1$. □

We do not know whether the analogous property is true for the graph parameter $\mathrm{ed}(G)$. On the other hand, one can prove the following partial result, whose proof was communicated to us by A. Schrijver.

Theorem 5. *For a graph* G, $\mathrm{ed}(\nabla G) \geq \mathrm{ed}(G) + 1$.

Proof. Set $\mathrm{ed}(\nabla G) = k$; we show $\mathrm{ed}(G) \leq k - 1$. We may assume that G is connected (else deal with each connected component separately). Let $d \in \mathrm{EDM}(G)$ and let $p_1 = 0, p_2, \ldots, p_n$ be a Euclidean representation of d in $\mathbb{R}^m$ $(m \geq 1)$. Extend the p_i's to vectors $\widehat{p_i} = (p_i, 0) \in \mathbb{R}^{m+1}$ by appending an extra coordinate equal to zero, and set $\widehat{p_0}(t) = (0, t) \in \mathbb{R}^{m+1}$ where t is any positive real scalar. Now consider the distance $\widehat{d}(t) \in \mathrm{EDM}(\nabla G)$ with Euclidean representation $\widehat{p_0}(t), \widehat{p_1}, \ldots, \widehat{p_n}$.

As $\mathrm{ed}(\nabla G) = k$, there exists another Euclidean representation of $\widehat{d}(t)$ by vectors $q_0(t), q_1(t), \ldots, q_n(t)$ lying in $\mathbb{R}^k$. Without loss of generality, we can assume that $q_0(t) = \widehat{p_0}(t) = (0, t)$ and $q_1(t)$ is the zero vector; for $i \in [n]$, write $q_i(t) = (u_i(t), a_i(t))$, where $u_i(t) \in \mathbb{R}^{k-1}$ and $a_i(t) \in \mathbb{R}$. Then $\|q_i(t)\| = \|\widehat{p_i}\| = \|p_i\|$ whenever node i is adjacent to node 1 in G. As the graph G is connected, this implies that, for any $i \in [n]$, the scalars $\|q_i(t)\|$ $(t \in \mathbb{R}_+)$ are bounded. Therefore there exists a sequence $t_m \in \mathbb{R}_+$ $(m \in \mathbb{N})$ converging to $+\infty$ and for which the sequence $(q_i(t_m))_m$ has a limit. Say $q_i(t_m) = (a_i(t_m), u_i(t_m))$ converges to $(u_i, a_i) \in \mathbb{R}^k$ as $m \to +\infty$, where $u_i \in \mathbb{R}^{k-1}$ and $a_i \in \mathbb{R}$. The condition $\|q_0(t) - q_i(t)\|^2 = \widehat{d}(t)_{0i}$ implies that $\|p_i\|^2 + t^2 = \|u_i(t)\|^2 + (a_i(t) - t)^2$ and thus

$$a_i(t_m) = \frac{a_i^2(t_m) + \|u_i(t_m)\|^2 - \|p_i\|^2}{2t_m} \quad \forall m \in \mathbb{N}.$$

Taking the limit as $m \to \infty$ we obtain that $\lim_{m\to\infty} a_i(t_m) = 0$ and thus $a_i = 0$. Then, for $i, j \in [n]$, $d_{ij} = \widehat{d}(t_m)_{ij} = \|(a_i(t_m), u_i(t_m)) - (a_j(t_m), u_j(t_m))\|^2$ and taking the limit as $m \to +\infty$ we obtain that $d_{ij} = \|u_i - u_j\|^2$. This shows that the vectors $u_1, \dots, u_n$ form a Euclidean representation of d in $\mathbb{R}^{k-1}$. □

This raises the following question: Is it true that $\text{ed}(\nabla G) \le \text{ed}(G)+1$? A positive answer would imply that our characterization for the graphs with Gram dimension 4 (Theorem 4) is equivalent to the characterization of Belk and Connelly for the graphs having a Euclidean representation in $\mathbb{R}^3$ (Theorem 1). In any case, we have that:

$$\text{gd}(G) = \text{ed}(\nabla G) \ge \text{ed}(G) + 1. \tag{8}$$

In the full version of the paper it is proven that $\text{gd}(V_8) = \text{gd}(C_5 \times C_2) = 4$ [15]. This fact combined with (8) implies that $\text{ed}(V_8) = \text{ed}(C_2 \times C_5) = 3$, which was the main part in the proof of Belk [5] to characterize graphs with $\text{ed}(G) \le 3$.

6 Some Complexity Results

Consider the natural decision problem associated with the graph parameter $\text{gd}(G)$: Given a graph G and a rational vector $x \in \mathcal{E}(G)$, determine whether $\text{gd}(G, x) \le k$, where $k \ge 1$ is some fixed integer. In this section we show that this is a hard problem for any $k \ge 3$, already when x is the zero vector. Further results concerning complexity issues associated with the graph parameter $\text{gd}(G)$ are discussed in [10].

Recall that an orthogonal representation of dimension k of $G = ([n], E)$ is a set of nonzero vectors $v_1, \dots, v_n \in \mathbb{R}^k$ such that $v_i^T v_j = 0$ for all pairs $ij \notin E$. Clearly, the minimum dimension of an orthogonal representation of the complementary graph $\overline{G}$ coincides with $\text{gd}(\overline{G}, 0)$; this graph parameter is called the *orthogonality dimension* of G, also denoted by $\xi(G)$. Note that it satisfies the inequalities $\omega(G) \le \xi(G) \le \chi(G)$, where $\omega(G)$ and $\chi(G)$ are the clique and chromatic numbers of G (see [16]).

One can easily verify that, for $k = 1, 2$, $\xi(G) \le k$ if and only if $\chi(G) \le k$, which can thus be tested in polynomial time. On the other hand, for $k = 3$, Peeters [19] gives a polynomial time reduction of the problem of testing $\xi(G) \le 3$ to the problem of testing $\chi(G) \le 3$; moreover this reduction preserves graph planarity. As a consequence, it is NP-hard to check whether $\text{gd}(G, 0) \le 3$, already for the class of planar graphs.

This hardness result for the zero vector extends to any $k \ge 3$, using the operation of adding an apex node to a graph. For a graph G, $\nabla^k G$ is the new graph obtained by adding iteratively k apex nodes to G.

Theorem 6. *For any fixed $k \geq 3$, it is NP-hard to decide whether* $\text{gd}(G, 0) \leq k$, *already for graphs G of the form $G = \nabla^{k-3}H$ where H is planar.*

Proof. Use the result of Peeters [19] for $k = 3$, combined with the first part of Lemma 6 for $k \geq 4$. □

Combining with Lemma 5 this implies that, for any fixed $k \geq 3$, it is NP-hard to decide whether $\text{ed}(G, d) \leq k$, already when $G = \nabla^{k-2}H$ where H is planar and $d \in \{1, 2\}^E$. In comparison, Saxe [22] showed NP-hardness for any $k \geq 1$ and for $d \in \{1, 2\}^E$.

Acknowledgements. We thank M. E.-Nagy for useful discussions and A. Schrijver for his suggestions for the proof of Theorem 5.

References

1. Arnborg, S., Proskurowski, A., Corneil, D.G.: Forbidden minors characterization of partial 3-trees. Disc. Math. 8(1), 1–19 (1990)
2. Avidor, A., Zwick, U.: Rounding Two and Three Dimensional Solutions of the SDP Relaxation of MAX CUT. In: Chekuri, C., Jansen, K., Rolim, J.D.P., Trevisan, L. (eds.) APPROX 2005 and RANDOM 2005. LNCS, vol. 3624, pp. 14–25. Springer, Heidelberg (2005)
3. Barahona, F.: The max-cut problem on graphs not contractible to K_5. Operations Research Letters 2(3), 107–111 (1983)
4. Barvinok, A.: A remark on the rank of positive semidefinite matrices subject to affine constraints. Disc. Comp. Geom. 25(1), 23–31 (2001)
5. Belk, M.: Realizability of graphs in three dimensions. Disc. Comput. Geom. 37, 139–162 (2007)
6. Belk, M., Connelly, R.: Realizability of graphs. Disc. Comput. Geom. 37, 125–137 (2007)
7. Candes, E.J., Recht, B.: Exact matrix completion via convex optimization. Foundations of Computational Mathematics 9(6), 717–772 (2009)
8. Deza, M., Laurent, M.: Geometry of Cuts and Metrics. Springer (1997)
9. Duffin, R.J.: Topology of series-parallel networks. Journal of Mathematical Analysis and Applications 10(2), 303–313 (1965)
10. E.-Nagy, M., Laurent, M., Varvitsiotis, A.: Complexity of the positive semidefinite matrix completion problem with a rank constraint (preprint, 2012)
11. Grone, R., Johnson, C.R., Sá, E.M., Wolkowicz, H.: Positive definite completions of partial Hermitian matrices. Linear Algebra and its Applications 58, 109–124 (1984)
12. Göring, F., Helmberg, C., Wappler, M.: The rotational dimension of a graph. J. Graph Theory 66(4), 283–302 (2011)
13. Hogben, L.: Orthogonal representations, minimum rank, and graph complements. Linear Algebra and its Applications 428, 2560–2568 (2008)
14. Laurent, M.: Matrix completion problems. In: Floudas, C.A., Pardalos, P.M. (eds.) The Encyclopedia of Optimization, vol. III, pp. 221–229. Kluwer (2001)
15. Laurent, M., Varvitsiotis, A.: A new graph parameter related to bounded rank positive semidefinite matrix completions (preprint, 2012)
16. Lovász, L.: On the Shannon capacity of a graph. IEEE Trans. Inform. Th. IT-25, 1–7 (1979)

17. Lovász, L.: Semidefinite programs and combinatorial optimization. Lecture Notes (1995), http://www.cs.elte.hu/~lovasz/semidef.ps
18. Lovász, L.: Geometric representations of graphs. Lecture Notes (2001), http://www.cs.elte.hu/~lovasz/geomrep.pdf
19. Peeters, R.: Orthogonal representations over finite fields and the chromatic number of graphs. Combinatorica 16(3), 417–431 (1996)
20. Recht, B., Fazel, M., Parrilo, P.A.: Guaranteed minimum-rank solutions of linear matrix equations via nuclear norm minimization. SIAM Review 52(3), 471–501 (2010)
21. Robertson, N., Seymour, P.D.: Graph minors. XX. Wagners conjecture. J. Combin. Theory Ser. B 92(2), 325–357 (2004)
22. Saxe, J.B.: Embeddability of weighted graphs in k-space is strongly NP-hard. In: Proc. 17th Allerton Conf. Comm. Control Comp., pp. 480–489 (1979)
23. Man-Cho So, A.: A semidefinite programming approach to the graph realization problem. PhD thesis, Stanford (2007)
24. Man-Cho So, A., Ye, Y.: A semidefinite programming approach to tensegrity theory and realizability of graphs. In: Proceedings of the 17th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), pp. 766–775 (2006)