Nullspace Embeddings for Outerplanar Graphs

László Lovász and Alexander Schrijver

Dedicated to the memory of Jiří Matoušek

Abstract We study relations between geometric embeddings of graphs and the spectrum of associated matrices, focusing on outerplanar embeddings of graphs. For a simple connected graph $G = (V, E)$, we define a "good" G-matrix as a $V \times V$ matrix with negative entries corresponding to adjacent nodes, zero entries corresponding to distinct nonadjacent nodes, and exactly one negative eigenvalue. We give an algorithmic proof of the fact that if G is a 2-connected graph, then either the nullspace representation defined by any "good" G-matrix with corank 2 is an outerplanar embedding of G, or else there exists a "good" G-matrix with corank 3.

1 Introduction

We study relations between geometric embeddings of graphs, the spectrum of associated matrices and their signature, and topological properties of associated cell complexes. We focus in particular on 1-dimensional and 2-dimensional embeddings of graphs, in the hope that the techniques can be extended to higher dimensions.

The research leading to these results has received funding from the European Research Council under the European Union's Seventh Framework Programme (FP7/2007–2013)/ERC grant agreement no. 227701.

The research leading to these results has received funding from the European Research Council under the European Union's Seventh Framework Programme (FP7/2007–2013)/ERC grant agreement no. 339109.

L. Lovász
Eötvös Loránd University, Budapest, Hungary

A. Schrijver (✉)
University of Amsterdam and CWI, Science Park 107, 1098 XG Amsterdam, The Netherlands
e-mail: lex@cwi.nl

M. Loebl et al. (eds.), *A Journey Through Discrete Mathematics*,
DOI 10.1007/978-3-319-44479-6_23

Spectral parameters of graphs The basic connection between graphs, matrices, and geometric embeddings considered in this paper can be described as follows. We define a *G-matrix* for an undirected graph $G = (V, E)$ as a symmetric real-valued $V \times V$ matrix M with $M_{ij} = 0$ if i and j are distinct nonadjacent nodes. The matrix is *well-signed* if $M_{ij} < 0$ for adjacent nodes i and j. (There is no condition on the diagonal entries.) If, in addition, M has exactly one negative eigenvalue, then let us call it *good* (for the purposes of this introduction). Let $\kappa(G)$ denote the largest d for which there exists a good G-matrix with corank d. (The corank is the dimension of the nullspace.)

The parameter κ is closely tied to certain topological properties of the graph. Combining results of [1, 5, 7, 9] and [8], one gets the following facts:

If G is connected, then $\kappa(G) \leq 1 \Leftrightarrow G$ is a path,
If G is 2-connected, then $\kappa(G) \leq 2 \Leftrightarrow G$ is outerplanar,
If G is 3-connected, then $\kappa(G) \leq 3 \Leftrightarrow G$ is planar,
If G is 4-connected, then $\kappa(G) \leq 4 \Leftrightarrow G$ is linklessly embeddable in $\mathbb{R}^3$.

We study algorithmic aspects of the first two facts. Let us discuss here the second, which says that if G is a 2-connected graph, then either it has an embedding in the plane as an outerplanar map, or else there exists a good G-matrix with corank 3 (and so the graph is not outerplanar). To construct an outerplanar embedding, we use the nullspace of any good G-matrix with corank 2.

Nullspace representations To describe this construction, suppose that a G-matrix M has corank d. Let $U \in \mathbb{R}^{d\times n}$ be a matrix whose rows form a basis of the nullspace of M. This matrix satisfies the equation $UM = 0$. Let u_i be the column of U corresponding to node $i \in V$. The mapping $u : V \to \mathbb{R}^d$ is called the *nullspace representation of V defined by M*. It is unique up to linear transformations of $\mathbb{R}^d$. (For the purist: the map $V \to \ker(M)^*$ is canonically defined; choosing the basis in $\ker(M)$ just identifies $\ker(M)^*$ with $\mathbb{R}^d$.)

If $G = (V, E)$ is a graph and $u : V \to \mathbb{R}^d$ is any map, we can extend it to the edges by mapping the edge ij to the straight line segment between u_i and u_j. If u is the nullspace representation of V defined by M, then this extension gives the *nullspace representation of G defined by M*.

In this paper we give algorithmic proofs of two facts:

1. If G is a connected graph with $\kappa(G) = 1$, then the nullspace representation defined by any good G-matrix with corank 1 yields an embedding of G in the line.
2. If G is 2-connected and $\kappa(G) = 2$, then the nullspace representation defined by any good G-matrix with corank 2 yields an outerplanar embedding of G.

(The word "yields" above hides some issues concerning normalization, to be discussed later.) The proofs are algorithmic in the sense that (say, in the case of (2)) for every 2-connected graph we either construct an outerplanar embedding or a good G-matrix with corank 3 in polynomial time. The alternative proof that can be derived from the results of [6] uses the minor-monotonicity of the Colin de Verdière parameter (see below), and this way it involves repeated reference to the Implicit Function Theorem, and does not seem to be implementable in polynomial time. Our algorithms use exact real arithmetic and a subroutine for finding roots of one-variable polynomials, which are steps that can be easily turned into polynomial-time algorithms (say, in binary arithmetic).

Suppose that the input to our algorithm is a 3-connected planar graph. Then the algorithm outputs a good G-matrix with corank at least 3. Paper [6] also contains the analogous result for planar graphs, which was extended in [4]:

3. If G is 3-connected and $\kappa(G) = 3$, then the nullspace representation defined by any good G-matrix with corank 3 yields a representation of G as the skeleton of a convex 3-polytope.

Thus computing the nullspace representation defined by the matrix M, and performing node-scaling as described in [4], we get a representation of G as the skeleton of a 3-polytope.

Unfortunately, the proof of (3) uses the minor-monotonicity of the Colin de Verdière parameter and the Implicit Function Theorem, and hence it does not yield an efficient algorithm: if the input is not a planar graph, then it does *not* provide a polynomial-time algorithm to compute a good G-matrix with corank at least 4. It would be interesting to see whether our approach can be extended to the case $\kappa \geq 3$. (While we focus on the case $\kappa, \leq 2$, some of our results do bear upon higher dimensions, in particular the results in Sect. 2.2 below.)

A further extension to dimension 4 would be particularly interesting, since for 4-connected graphs G, linkless embeddability is characterized by the property that $\kappa(G) \leq 4$, but it is not known whether the nullspace representation obtained from a good G-matrix of corank 4 yields a linkless embedding of the graph.

The Strong Arnold Hypothesis and the Colin de Verdière number We conclude this introduction with a discussion of the connection between the parameter $\kappa(G)$ and the graph parameter $\mu(G)$ introduced by Colin de Verdière (cf. [11]). This latter is defined similarly to κ as the maximum corank of a good G-matrix M, where it is required, in addition, that M has a nondegeneracy property called the *Strong Arnold Property*. There are several equivalent forms of this property; let us formulate one that is related to our considerations in the sense that it uses any nullspace representation u defined by M: if a symmetric $d \times d$ matrix N satisfies $u_i^\mathsf{T} N u_i = 0$ for all $i \in V$ and $u_i^\mathsf{T} N u_j = 0$ for each edge ij of G, then $N = 0$. In more geometric terms this means that the nullspace representation of the graph defined by M is not contained in any nontrivial homogeneous quadric.

The relationship between μ and κ is not completely clarified. Trivially $\mu(G) \leq \kappa(G)$. Equality does not hold in general: consider the graph $G_{l,m}$ made from an $(l + m)$-clique by removing the edges of an m-clique. If $l \geq 1$ and $m \geq 3$, then $\mu(G_{l,m}) = l+1$ whereas $\kappa(G_{l,m}) = l+m-2$. (Note that $G_{l,m}$ is not $l+1$-connected.)

Colin de Verdière's parameter has several advantages over κ. First, it is minor-monotone, while $\kappa(G)$ is not minor-monotone, not even subgraph-monotone: any path P satisfies $\kappa(P) \leq 1$, but a disjoint union of paths can have arbitrarily large

$\kappa(G)$. Furthermore, the connection with topological properties of graphs holds for μ without connectivity conditions:

$\mu(G) \leq 1 \Leftrightarrow G$ is a disjoint union of paths,
$\mu(G) \leq 2 \Leftrightarrow G$ is outerplanar,
$\mu(G) \leq 3 \Leftrightarrow G$ is planar,
$\mu(G) \leq 4 \Leftrightarrow G$ is linklessly embeddable in $\mathbb{R}^3$.

Our use of κ is motivated by its easier definition and by the (slightly) stronger, algorithmic results.

We see from the facts above that by requiring that G is $\mu(G)$-connected, we have $\mu(G) = \kappa(G)$ for $\mu(G) \leq 4$. In fact, it was shown by Van der Holst [10] that if G is 2-connected outerplanar or 3-connected planar, then *every* good G-matrix has the Strong Arnold Property. This also holds true for 4-connected linklessly embeddable graphs [8]. One may wonder whether this remains true for $\mu(G)$-connected graphs with larger $\mu(G)$. This would imply that $\mu(G) = \kappa(G)$ for every $\mu(G)$-connected graph.

Remark 1 Our setup is related to rigidity theory of bar-and-joint structures. To formulate just one connection, let G be a graph, M a well-signed G-matrix, and $u : V(G) \to \mathbb{R}^d$ a nullspace representation, considered as specifying a position for each node. Replace the edges by rubber bands of strength M_{ij} (i.e., stretching an edge to length t results in a force of $-M_{ij}t$ pulling the endpoints together). Add "braces" (rigid bars) from the origin to each node; these braces can carry an arbitrary force, as long as it is parallel to the brace. Then the equation $UM = 0$ just says that the structure is in equilibrium (where, as before, U is the matrix with columns u_i). The matrix M is called a (braced) *stress matrix* on the structure (G, u).

Other conditions like the rank of the matrix M, its signature and its Strong Arnold Property also play a significant role in rigidity theory; see [2, 3].

2 *G*-Matrices

2.1 Nullspace Representations

Let us fix a connected graph $G = (V, E)$ on node set $V = [n]$, and an integer $d \geq 1$. We denote by $\mathcal{W}$ the set of well-signed G-matrices with corank at least d, and by $\mathcal{W}^=$, the set of well-signed G-matrices with corank exactly d. We denote by $\mathcal{W}^1$ the set of G-matrices in $\mathcal{W}$ with exactly one negative eigenvalue (counted with multiplicity).

Suppose that we are also given a vector labeling $u : V \to \mathbb{R}^d$, which we can encode as a $d \times V$ matrix U, whose column corresponding to $i \in V$ is the vector u_i. For $p \in \mathbb{R}^d$, let us write $u - p$ for the representation $(u_1 - p, \ldots, u_n - p)$. We denote by $\mathcal{M}_u$ the linear space of G-matrices M with $UM = 0$, by $\mathcal{W}_u$, the set of well-signed G-matrices in $\mathcal{M}_u$, by $\mathcal{W}_u^1$, the set of matrices in $\mathcal{W}_u$ with exactly one negative eigenvalue, and by $\mathcal{W}_u^2$, the set of matrices in $\mathcal{W}_u$ with at least two negative eigenvalues.

We can always perform a linear transformation of $\mathbb{R}^d$, i.e., replace U by AU, where A is any nonsingular $d \times d$ matrix. In the case when $\text{corank}(M) = d$ (which will be the important case for us), the matrix U is determined by M up to such a linear transformation of $\mathbb{R}^d$.

Another simple transformation we use is "node scaling": replacing U by $U' = UD$ and M by $M' = D^{-1}MD^{-1}$, where D is a nonsingular diagonal matrix with positive diagonal. Then M' is a G-matrix and $U'M' = 0$. Moreover, it maintains well-signedness of M. Through this transformation, we may assume that every nonzero vector u_i has unit length. We call such a representation *normalized*.

One of our main tools will be to describe more explicit solutions of the basic equation $UM = 0$ in dimensions 1 and 2. More precisely, given a graph $G = (V, E)$ and a vector labeling $u : V \to \mathbb{R}^2$, our goal is to describe all G-matrices M with $UM = 0$. Note that if the vector labels are nonzero, then it suffices to find the off-diagonal entries: if M_{ij} is given for $ij \in E$ in such a way that $\sum_{j \in N(i)} M_{ij} u_j$ is a scalar multiple of u_i for every node i, then there is a unique choice of diagonal entries M_{ii} that gives a matrix with $UM = 0$:

$$M_{ii} = -\sum_j M_{ij} \frac{u_j^\top u_i}{u_i^\top u_i}. \tag{1}$$

2.2 G-Matrices and Eigenvalues

In this section we consider eigenvalues of well-signed G-matrices; we consider the connected graph G and the dimension parameter d fixed. We start with a couple of simple observations.

Lemma 2 *Let M be a well-signed G-matrix and let $U \in \mathbb{R}^{d \times n}$ such that $UM = 0$ and* $\text{rank}(U) = d$.

(a) *If M is positive semidefinite, then* $\text{corank}(M) = d = 1$*, and all entries of U are nonzero and have the same sign.*
(b) *If M has a negative eigenvalue, then the origin is an interior point of the convex hull of the columns of U.*

Proof Let λ be the smallest eigenvalue of M. As G is connected, λ has multiplicity one by the Perron–Frobenius theorem, and M has a positive eigenvector v belonging to λ. If $\lambda = 0$, then this multiplicity is $d = 1$, and U consists of a single row parallel to v. If $\lambda < 0$, then every row of U, being in the nullspace of M, is orthogonal to v. Thus the entries of v provide a representation of 0 as a convex combination of the columns of U with positive coefficients. □

Lemma 3 *If $d \geq 2$, then the set $\mathcal{W}^1$ is relatively closed in $\mathcal{W}$, and $\mathcal{W}^1 \cap \mathcal{W}^=$ is relatively open in $\mathcal{W}$.*

Proof Let $\lambda_i(M)$ denote the i-th smallest eigenvalue of the matrix M. We claim that for any $M \in \mathcal{W}$,

$$M \in \mathcal{W}^1 \;\Leftrightarrow\; \lambda_2(M) \geq 0. \qquad (2)$$

Indeed, if $M \in \mathcal{W}^1$, then trivially $\lambda_2(M) \geq 0$. Conversely, if $\lambda_2(M) \geq 0$, then M has at most one negative eigenvalue. By Lemma 2(a), it has exactly one, that is, $M \in \mathcal{W}^1$. This proves (2). Since $\lambda_2(M)$ is a continuous function of M, the first assertion of the lemma follows.

We claim that if $d \geq 2$, for any $M \in \mathcal{W}$,

$$M \in \mathcal{W}^1 \cap \mathcal{W}^= \;\Leftrightarrow\; \lambda_{d+2}(M) > 0. \qquad (3)$$

Indeed, if $M \in \mathcal{W}^1 \cap \mathcal{W}^=$, then M has one negative eigenvalue and exactly d zero eigenvalues, and so $\lambda_{d+2}(M) > 0$. Conversely, assume that $\lambda_{d+2}(M) > 0$. Since M has at least d zero eigenvalues and at least one negative eigenvalue (by Lemma 2(a)), we must have equality in both bounds, which means that $M \in \mathcal{W}^1 \cap \mathcal{W}^=$. This proves (3). Continuity of $\lambda_{d+2}(M)$ implies the second assertion. □

This last lemma implies that each nonempty connected subset of $\mathcal{W}^=$ is either contained in $\mathcal{W}^1$ or is disjoint from $\mathcal{W}^1$. We formulate several consequences of this fact.

Lemma 4 *Suppose that G is 2-connected, and let M be a well-signed G-matrix with one negative eigenvalue and with corank $d = \kappa(G)$. Let u be the nullspace representation defined by M, let $v \in \mathbb{R}^d$, and let $J := \{i : u_i = v\}$. If $|J| \geq 2$, then the origin 0 belongs to the convex hull of $u(V \setminus J)$.*

Proof For $i \in V$, let e_i be the i-th unit basis vector, and for $i, j \in V$, let D^{ij} be the matrix $(e_i - e_j)(e_i - e_j)^\mathsf{T}$. Define

$$M^\alpha := M + \alpha \sum_{\substack{ij \in E \\ i,j \in J}} M_{ij} D^{ij} \qquad (\alpha \in [0, 1]).$$

The definition of J implies that $\ker(M) \subseteq \ker(D^{ij})$ for all $i, j \in J$, and hence $\ker(M) \subseteq \ker(M^\alpha)$ for each $\alpha \in [0, 1]$. So $\operatorname{corank}(M^\alpha) \geq \operatorname{corank}(M) = \kappa(G)$ for each $\alpha \in [0, 1]$. Moreover, M^α is a well-signed G-matrix for each $\alpha \in [0, 1)$. So $M^\alpha \in \mathcal{W}$ for each $\alpha \in [0, 1)$. As $\kappa(G) = d$, we know $\mathcal{W}^1 \subseteq \mathcal{W}^=$, hence $\mathcal{W}^1 \cap \mathcal{W}^= = \mathcal{W}^1$. So by Lemma 3, $\mathcal{W}^1$ is relatively open and closed in $\mathcal{W}$. Since $M = M^0 \in \mathcal{W}^1$, this implies that $M^\alpha \in \mathcal{W}^1$ for each $\alpha \in [0, 1)$. By the continuity of eigenvalues, M^1 has at most one negative eigenvalue. Note that $M^1_{ij} = 0$ for any two distinct $i, j \in J$.

Assume that 0 does not belong to the convex hull of $\{u_i : i \notin J\}$. Then there exists $c \in \mathbb{R}^{\kappa(G)}$ such that $u_i^\mathsf{T} c < 0$ for each $i \notin J$. As 0 belongs to interior of the convex hull of $u(V)$ by Lemma 2(b), this implies that $u_i^\mathsf{T} c = v^\mathsf{T} c > 0$ for each $i \in J$.

As $|J| \geq 2$, the 2-connectivity of G implies that J contains two distinct nodes, say nodes 1 and 2, that have neighbors outside J. Since $\ker(M) \subseteq \ker(M^1)$, we have $\sum_j M^1_{1j} u_j = 0$, and hence

$$M^1_{11} u_1^\mathsf{T} c = -\sum_{j \neq 1} M^1_{1j} u_j^\mathsf{T} c = -\sum_{j \notin J} M^1_{1j} u_j^\mathsf{T} c.$$

As $u_1^\mathsf{T} c > 0$ and $u_j^\mathsf{T} c < 0$ for $j \notin J$, and as $M^1_{1j} \leq 0$ for all $j \notin J$, and $M^1_{1j} < 0$ for at least one $j \notin J$, this implies $M^1_{11} < 0$. Similarly, $M^1_{22} < 0$. As $M^1_{12} = 0$, the first two rows and columns of M^1 induce a negative definite 2×2 submatrix of M^1. This contradicts the fact that M^1 has at most one negative eigenvalue. □

For the next step we need a simple lemma from linear algebra.

Lemma 5 *Let A and M be symmetric $n \times n$ matrices. Assume that A is 0 outside a $k \times k$ principal submatrix, and let M_0 be the complementary $(n-k) \times (n-k)$ principal submatrix of M. Let a and b denote the number of negative eigenvalues of A and M_0, respectively. Then for some $s > 0$, the matrix $sM + A$ has at least $a + b$ negative eigenvalues.*

Proof We may assume $A = \begin{pmatrix} A_0 & 0 \\ 0 & 0 \end{pmatrix}$ and $M = \begin{pmatrix} M_1 & M_2^\mathsf{T} \\ M_2 & M_0 \end{pmatrix}$, with A_0 and M_1 having order $k \times k$. By scaling the last $n-k$ rows and columns of $sM + A$ by $1/\sqrt{s}$, we get the matrix $\begin{pmatrix} sM_1 + A_0 & \sqrt{s}M_2^\mathsf{T} \\ \sqrt{s}M_2 & M_0 \end{pmatrix}$. Letting $s \to 0$, this tends to $B = \begin{pmatrix} A_0 & 0 \\ 0 & M_0 \end{pmatrix}$. Clearly, B has $a + b$ negative eigenvalues, and by the continuity of eigenvalues, the lemma follows. □

Lemma 6 *Let M be a well-signed G-matrix with one negative eigenvalue and with corank $d = \kappa(G)$, let u be the nullspace representation defined by M, and let C be a clique in G of size at most $\kappa(G)$ such that the origin belongs to the convex hull of $u(C)$. Then $G - C$ is disconnected.*

Proof Since the origin belongs to the convex hull of $u(C)$, we can write $0 = \sum_i a_i u_i$ with $a_i \geq 0$, $\sum_i a_i = 1$, and $a_i = 0$ if $i \notin C$. Let A be the matrix $-aa^\mathsf{T}$. Since a is nonzero, A has a negative eigenvalue.

Since $\sum_i a_i u_i = 0$, we have $\ker(M) \subseteq \ker(M + sA)$ for each s. This implies that $\mathrm{corank}(M + sA) \geq \mathrm{corank}(M)$ for each s. Moreover, $M + sA$ is a well-signed G-matrix for $s \geq 0$. So $M + sA \in \mathcal{W}$ for each $s \geq 0$. Hence, as $M \in \mathcal{W}^1$ and as $\mathcal{W}^1 \subseteq \mathcal{W}^=$ (since $d = \kappa(G)$), we know by Lemma 3 that $M + sA \in \mathcal{W}^1$ for every $s \geq 0$. In other words, $M + sA$ has one negative eigenvalue for every $s \geq 0$.

Let M_0 be the matrix obtained from M by deleting the rows and columns with index in C. Note that M_0 has no negative eigenvalue: otherwise by Lemma 5, $M + sA$ has at least two negative eigenvalues for some $s > 0$, a contradiction.

Now suppose that $G - C$ is connected. As $u(C)$ is linearly dependent and $|C| \leq \text{corank}(M)$, $\ker(M)$ contains a nonzero vector x with $x_i = 0$ for all $i \in C$. Then by the Perron–Frobenius theorem, $\text{corank}(M_0) = 1$ and $\ker(M_0)$ is spanned by a positive vector y. As G is connected, x is orthogonal to the positive eigenvector belonging to the negative eigenvalue of M. So x has both positive and negative entries. On the other hand, $x|_{V \setminus C} \in \ker(M_0)$, and so $x|_{V \setminus C}$ must be a multiple of y, a contradiction. □

Taking C a singleton, we derive:

Corollary 7 *Let G be a 2-connected graph, let $M \in \mathcal{W}^1$ have corank $\kappa(G)$, and let u be the nullspace representation defined by M. Then $u_i \neq 0$ for every node i.*

Equivalently, the nullspace representation defined by M can be normalized by node scaling.

2.3 Auxiliary Algorithms

Now we turn to the algorithmic part, starting with some auxiliary algorithms. The following general argument will be needed repeatedly.

Algorithm 1 (Interpolation)

Input: a continuous family of full-row-rank matrices $U(t) \in \mathbb{R}^{d \times n}$, and a continuous family of symmetric matrices $M(t) \in \mathbb{R}^{n \times n}$ $(0 \leq t \leq 1)$ such that $U(t)M(t) = 0$, $M(0)$ has exactly one negative eigenvalue and $M(1)$ has at least two negative eigenvalues.

Output: a value $t \in [0, 1]$ for which $M(t)$ has at most one negative eigenvalue and at least $d + 1$ zero eigenvalues.

Let $X := \{t \mid \lambda_2(M(t)) \geq 0\}$ and $Y := \{t \mid \lambda_{d+2}(M(t)) \leq 0\}$. Since $U(t)M(t) = 0$ and $U(t)$ has full row rank, every matrix $M(t)$ has at least d zero eigenvalues. Hence $X \cup Y = [0, 1]$. Therefore, as X and Y are closed, and as X is a nonempty proper subset of $[0, 1]$ (since $0 \in X$, $1 \notin X$), we have $X \cap Y \neq \emptyset$, that is, $t \in X \cap Y$ for some t.

How to compute such a value of t? By binary search, we can compute it with arbitrary precision. In our applications, we can do better, since the entries of the families $U(t)$ and $M(t)$ will be (very simple) rational functions of t. We can find those values of t for which $M(t)$ has corank at least $d + 1$ by considering any nonsingular $(n - d) \times (n - d)$ submatrix of $M(0)$, and finding the roots of $\det(B(t)) = 0$, where $B(t)$ is the corresponding submatrix of $M(t)$. Then every value of t with $\text{corank}(M(t)) > d$ is one of these roots. The smallest such value of t will give a matrix $M(t)$ with corank at least $d+1$. Since the matrices $M(s)$ with $s < t$ have at most one negative eigenvalue (as otherwise $[0, t) \cap Y \neq \emptyset$ (since $X \cup Y = [0, 1]$), hence $[0, t) \cap X \cap Y \neq \emptyset$ (as $0 \in X$ and X and Y are closed), so $\text{corank}(M(s)) > d$ for some $s < t$), the matrix $M(t)$ has at most one.

We describe two simple applications of this general method.

Algorithm 2 (Double zero node)

Input: a connected graph $G = (V, E)$, a full-dimensional vector labeling u in $\mathbb{R}^d$, two nodes i and j with $u_i = u_j = 0$, and a matrix $M \in \mathcal{W}_u^1$.

Output: a matrix $M' \in \mathcal{W}_u^1$ with $\text{corank}(M') \geq d + 1$.

Subtract $t > 0$ from both diagonal entries M_{ii} and M_{jj}, to get a matrix $M(t)$. Trivially $M(t) \in \mathcal{W}_u$. Furthermore, if $t > 2\max\{|M_{ii}|, |M_{jj}|, |M_{ij}|\}$, then the submatrix of $M(t)$ formed by rows and columns i and j has negative trace and positive determinant, and so it has two negative eigenvalues. This implies by Interlacing Eigenvalues that $M(t)$ has at least two negative eigenvalues. Calling Algorithm 1, we get a $0 \leq s \leq t$ such that $M(s)$ has at most one negative eigenvalue and at least $d+1$ zeroes. Lemma 2 implies that $M(s)$ cannot be positive semidefinite, so $M(s) \in \mathcal{W}_u^1$.

Algorithm 3 (Zero node)

Input: a 2-connected graph $G = (V, E)$, a full-dimensional vector labeling u in $\mathbb{R}^d$, a node i with $u_i = 0$, and a matrix $M \in \mathcal{W}_u^1$.

Output: a matrix $M'' \in \mathcal{W}_u^1$ with $\text{corank}(M'') \geq d + 1$.

We may assume $i = 1$. Let N be the matrix obtained from M by deleting row and column 1. Any coordinate of the vectors u_j $(j \neq 1)$ is in the nullspace of N. Since $G \setminus 1$ is connected, the Perron–Frobenius Theorem implies that N is not positive semidefinite (otherwise $d = 1$ by Lemma 2(a), and then $(u_i \mid i \neq 1)$ would be the eigenvector of N belonging to the smallest eigenvalue 0, while this vector is not constant in sign). So N has a negative eigenvalue λ, with eigenvector y $(|y| = 1)$. Replacing M_{11} by a sufficiently small negative number s, we get a matrix $M' \in \mathcal{W}_u$ with two negative eigenvalues. Simple linear algebra shows that $s < \left(e_1^\mathsf{T} M \binom{0}{y}\right)^2 / \lambda$ suffices. We conclude by calling Algorithm 1 as before.

3 1-Dimensional Nullspace Representations

As a warmup, let us settle the case $d = 1$. For every connected graph $G = (V, E)$, it is easy to construct a singular G-matrix with exactly one negative eigenvalue: start with any G-matrix, and subtract an appropriate constant from the main diagonal. Our goal is to show that unless the graph is a path and the nullspace representation is a monotone embedding in the line, we can modify the matrix to get a G-matrix with one negative eigenvalue and with corank at least 2.

3.1 Nullspace and Neighborhoods

We start with noticing that given a vector $u \in \mathbb{R}^V$, it is easy to describe the matrices in $\mathcal{W}_u$. Indeed, consider any matrix $M \in \mathcal{M}_u$. Then for every node i with $u_i = 0$, we have

$$\sum_{j\in N(i)} M_{ij}u_j = \sum_j M_{ij}u_j = 0. \tag{4}$$

Furthermore, for every node i with $u_i \neq 0$, we have

$$M_{ii} = -\frac{1}{u_i}\sum_{j\in N(i)} M_{ij}u_j. \tag{5}$$

Conversely, if we specify the off-diagonal entries of a G-matrix M so that (4) is satisfied for each i with $u_i = 0$, then we can define M_{ii} for nodes $i \in \mathrm{supp}(u)$ according to (5), and for nodes i with $u_i = 0$ arbitrarily, we get a matrix in $\mathcal{M}_u$.

As an application of this construction, we prove the following lemma.

Lemma 8 *Let $u \in \mathbb{R}^V$. Then $\mathcal{W}_u \neq \emptyset$ if and only if for every node i with $u_i = 0$, either all its neighbors satisfy $u_j = 0$, or it has neighbors both with $u_j < 0$ and $u_j > 0$.*

Proof By the remark above, it suffices to specify negative numbers M_{ij} for the edges ij so that (4) is satisfied for each i with $u_i = 0$. The edges between two nodes with $u_i = 0$ play no role, and so the conditions (4) can be considered separately. For a fixed i, the single linear equation for the M_{ij} can be satisfied by negative numbers if and only if the condition in the lemma holds. □

We need the following fact about the neighbors of the other nodes.

Lemma 9 *Let $u \in \mathbb{R}^V$, $M \in \mathcal{W}_u$, and suppose that M has a negative eigenvalue $\lambda < 0$, with eigenvector $\pi > 0$. Then every node i with $u_i > 0$ has a neighbor j for which $u_j/\pi_j < u_i/\pi_i$.*

Proof Suppose not. Then $u_j \geq \pi_j u_i/\pi_i$ for every $j \in N(i)$, and so

$$0 = \sum_j M_{ij}u_j \leq M_{ii}u_i + \sum_{j\in N(i)} M_{ij}\frac{\pi_j}{\pi_i}u_i = \frac{u_i}{\pi_i}\Big(\sum_j M_{ij}\pi_j\Big) = \lambda u_i < 0,$$

a contradiction. □

Algorithm 4 (Double cover)

Input: a vector $u \in \mathbb{R}^V$, two edges ab and cd with $u_a u_b \leq 0$, $u_c u_d \leq 0$, $b \neq d$, $u_a \neq 0$, $u_c \neq 0$, and a matrix $M \in \mathcal{W}_u^1$.

Output: a matrix $M' \in \mathcal{W}_u^1$ of corank at least 2.

Define the symmetric matrix $N^{ab} \in \mathbb{R}^{V\times V}$ by

$$(N^{ab})_{ij} = \begin{cases} u_a u_b, & \text{if } \{i,j\} = \{a,b\}, \\ -u_b^2, & \text{if } i = j = a, \\ -u_a^2, & \text{if } i = j = b, \\ 0, & \text{otherwise,} \end{cases}$$

and define N^{cd} analogously. Then $N^{ab}u = N^{cd}u = 0$, and so $M' = M+tN^{ab}+tN^{cd} \in \mathcal{W}_u$ for every $t > 0$. Moreover, $N^{ab} + N^{cd}$ has two negative eigenvalues, as one may (case-)check. So $M + tN^{ab} + tN^{cd} \in \mathcal{W}_u^2$ for some t, by Lemma 5. So with the Interpolation Algorithm 1 we find M' as required.

3.2 Embedding in the Line

Now we come to the main algorithm for dimension 1.

Algorithm 5

Input: A connected graph $G = (V, E)$.

Output: Either an embedding $u : V \to \mathbb{R}$ of G (then G is a path), or a well-signed G-matrix with one negative eigenvalue and corank at least 2.

Preparation We find a matrix $M \in \mathcal{W}^1$. This is easy by creating any well-signed G-matrix and subtracting its second smallest eigenvalue from the diagonal. We may assume that $\text{corank}(M) = 1$, else we are done.

Let $u \neq 0$ be a vector in the nullspace of M, and let π be an eigenvector belonging to its negative eigenvalue. Then the matrix $M' = \text{diag}(\pi)M\text{diag}(\pi)$ is in $\mathcal{W}^1(G)$ and the vector $w = (u_i/\pi_i : i \in V)$ is in its nullspace. By Lemma 9, this means that if we replace M by M' and u by w, then we get a vector $u \in \mathbb{R}^n$ and a matrix $M \in \mathcal{W}_u^1$ such that every node i with $u_i > 0$ has a neighbor j with $u_j < u_i$, and every node i with $u_i < 0$ has a neighbor j with $u_j > u_i$.

If $u_i = u_j = 0$ for some distinct i, j, we can apply Algorithm 2. So we can assume that $u_i = 0$ for at most one i.

Let us define a *cell* as an open interval between two consecutive points u_i. If every cell is covered by only one edge, then G is a path and u defines an embedding of G in the line, and we are done. Indeed, suppose first $u_i = u_j$ with $j \neq i$. By assumption $u_i \neq 0$. If $u_i > 0$, then both i and j have a neighbour i' and j' respectively, with $u_{i'} < u_i$ and $u_{j'} < u_j$, hence some cell is covered twice by edges. Similarly if $u_i < 0$. So the u_i are all distinct. Assuming that each cell is covered at most once by an edge, u must be an embedding of G into $\mathbb{R}$, and so G is a path.

So we can assume that there exists a cell (a, b) covered by at least two edges. We choose (a, b) nearest to the origin. Replacing u by $-u$ if necessary, we may assume that $b > 0$.

Main step Below, we are going to maintain the following conditions. We have a vector $u \in \mathbb{R}^V$ and a matrix $M \in \mathcal{W}_u^1$; every node i with $u_i > 0$ has a neighbor j with $u_j < u_i$; there is a cell (a, b) with $b > 0$ that is doubly covered, and that is nearest the origin among such cells.

We have to distinguish some cases.

Case 1. If $a < 0$, then we use the Double Cover Algorithm 4 to obtain a matrix with the desired properties.

Case 2. If $a \geq 0$, then let u_p be the smallest nonnegative entry of u.

Case 2.1. Assume that $u_p = 0$. Let $(0, c)$ be the cell incident with 0 and with $c > 0$, and let M' be obtained from M by replacing the (p, p) diagonal entry by 0. Then $M' \in \mathcal{W}_u$. It follows by Lemma 2(a) that M' is not positive semidefinite. If M' has more than one negative eigenvalue, then we can run the Interpolation Algorithm 1. So we may assume that $M' \in \mathcal{W}_u^1$.

For $0 < t < c$, consider the G-matrices $A(t)$ defined for edges ij by

$$A(t)_{ij} = A(t)_{ji} = \begin{cases} M_{ij}, & \text{if } i, j \neq p, \\ \dfrac{u_j}{u_j - t} M_{pj}, & \text{if } i = p, \end{cases}$$

and on the diagonal by

$$A(t)_{ii} = -\frac{1}{u_i - t} \sum_{j \in N(i)} A(t)_{ij}(u_j - t).$$

Clearly $A(t) \in \mathcal{W}_{u-t}$. Lemma 2(a) implies that $A(t)$ has at least one negative eigenvalue. Furthermore, if $t \to 0$, then $A(t)_{ij} \to M_{ij}$; this is trivial except for $i = j = p$, when, using that $\sum_{j \in N(p)} M_{pj} u_j = -M_{pp} u_p = 0$, we have

$$A(t)_{pp} = \frac{1}{t} \sum_{j \in N(p)} M_{pj} u_j = 0.$$

Thus defining $A(0) = M'$ the family $A(t)$ remains continuous.

If the matrix $A(c/2)$ has one negative eigenvalue, then replace M by $A(c/2)$ and u by $u - c/2$, and return to the Main Step. Note that the number of nodes with $u_i \geq 0$ has decreased, while those with $u_i > 0$ did not change.

If $A(c/2)$ has at least two negative eigenvalues, then the Interpolation Algorithm 1 can be applied to the families $(A(t))$ and $(u - t)$ to get a number $0 \leq s \leq c/2$ with $A(s) \in \mathcal{W}_{u-t}^1$ and $\text{corank}(A(s)) > 1$.

Case 2.2. Assume that $u_p > 0$. Let σ and τ denote the cells to the left and to the right of u_p (so $0 \in \sigma$). There is no other node q with $u_q = u_p$ (since from both nodes, an edge would start to the left, whereas 0 is covered only once). From u_p, there is an edge starting to the left, and also one to the right (since by connectivity, there is an edge covering τ, and this must start at p, since σ is covered only once). Therefore, $\mathcal{W}_{u-u_p} \neq \emptyset$ by Lemma 8. Following the proof of this Lemma, we can construct a matrix $B \in \mathcal{W}_{u-u_p}$ with $B_{pp} = 0$. Since $u - u_p$ has a zero entry, Lemma 2(a) implies that B has at least one negative eigenvalue.

For $t \in [0, u_p)$, consider the G-matrices $B(t)$ defined for edges ij by

$$B(t)_{ij} = B(t)_{ji} = \begin{cases} B_{ij}, & \text{if } i, j \neq p, \\ \dfrac{u_j - u_p}{u_j - t} B_{pj}, & \text{if } i = p, \end{cases}$$

and on the diagonal by

$$B(t)_{ii} = -\frac{1}{u_i - t} \sum_{j \in N(i)} B(t)_{ij}(u_j - t).$$

Clearly, $B(t) \in \mathcal{W}_{u-t}$. Furthermore, $\lim_{t \to u_p} B(t) = B$.

If B has one negative eigenvalue, then replace M by B and u by $u - u_p$, and go to the Main Step. Note that the number of nodes with $u_i > 0$ has decreased, while those with $u_i \geq 0$ did not change.

If B has more than one negative eigenvalue and $B(0)$ has only one, then the Interpolation Algorithm 1 gives a value $0 \leq s < u_p$ such that $B(s) \in \mathcal{W}^1_{u-s}$ and $\text{corank}(B(s)) > 1$.

Finally, if $B(0)$ has more than one negative eigenvalue, then we call the Interpolation Algorithm 1 for the family of matrices $(1-t)M + tB(0)$, keeping u fixed.

4 2-Dimensional Nullspace Representations

4.1 G-Matrices and Circulations

Our goal in this section is to provide a characterization of G-matrices and their nullspace representations in dimension 2.

A *circulation* on an undirected simple graph G is a real $V \times V$ matrix f such that it is supported on adjacent pairs, is skew-symmetric and satisfies the flow conditions:

$$f_{ij} = 0 \ (ij \notin E), \quad f_{ij} = -f_{ji} \ (i,j \in V), \quad \sum_j f_{ij} = 0 \ (i \in V).$$

If we fix an orientation of the graph, then it suffices to specify the values of f on the oriented edges; the values on the reversed edges follow by skew symmetry. A *positive circulation* on an oriented graph (V, A) is a circulation on the underlying undirected graph that takes positive values on the arcs in A.

For any representation $u : V \to \mathbb{R}^2$, we define its *area-matrix* as the (skew-symmetric) matrix $T = T(u)$ by $T_{ij} := \det(u_i, u_j)$. This number is the signed area of the parallelogram spanned by u_i and u_j. If R denotes counterclockwise rotation by 90°, then $T_{ij} = u_i^{\mathsf{T}} R u_j$.

Given a graph G and a representation $u : V \to \mathbb{R}^2$ by nonzero vectors, we define a directed graph (V, A_u) and an undirected graph (V, E_u) by

$$A_u := \{(i,j) \in V \times V \mid ij \in E, T(u)_{ij} > 0\}$$

$$E_u := \{ij \in E \mid T(u)_{ij} = 0\}.$$

So E is partitioned into A_u and E_u, where (V, A_u) is an oriented graph in which each edge is oriented counterclockwise as seen from the origin. The graph (V, E_u) consists of all edges that are contained in a line through the origin.

Given a representation $u : V \to \mathbb{R}^2$, a circulation f on (V, A_u), and a function $g : E_u \to \mathbb{R}$, we define a G-matrix $M(u, f, g)$ by

$$M(u,f,g)_{ij} = \begin{cases} -f_{ij}/T(u)_{ij}, & \text{if } ij \in A_u, \\ g(ij), & \text{if } ij \in E_u. \end{cases}$$

We define the diagonal entries by (1), and let the other entries be 0. The first main ingredient of our proof and algorithm is the following representation of G-matrices with a given nullspace.

Lemma 10 *Let $G = (V, E)$ be a graph and let $u : V \to \mathbb{R}^2$ be a labeling of V by nonzero vectors. Then*

$$\mathcal{M}_u = \{M(u,f,g) : f \textit{ is a circulation on } (V, A_u) \text{ and } g : E_u \to \mathbb{R}\}.$$

Proof First, we prove that $M(u,f,g) \in \mathcal{M}_u$ for every circulation on (V, A_u) and every $g : E_u \to \mathbb{R}$. Using that $M(u,f,g) = M(u,f,0) + M(u,0,g)$, it suffices to prove that $M(u,f,g) \in \mathcal{M}_u$ if either $f = 0$ or $g = 0$. If $M = M(u,f,0)$, then using that f is a circulation, we have

$$\Big(\sum_j M_{ij}u_j\Big)^{\mathsf{T}} Ru_i = \sum_j f_{ij} = 0.$$

This means that $\sum_j M_{ij}u_j^{\mathsf{T}}$ is orthogonal to Ru_i, and so parallel to u_i. As remarked above, this means that $M(u,f,0) \in \mathcal{M}_u$. If $M = M(u,0,g)$, then for every $i \in V$,

$$\sum_{j \in N(i)} M_{ij}u_j = \sum_{j:\, ij \in E_u} g(ij)u_j$$

This vector is clearly parallel to u_i, proving that $M(u,0,g) \in \mathcal{M}_u$.

Second, given a matrix $M \in \mathcal{M}_u$, define $f_{ij} = -T_{ij}M_{ij}$ for $ij \in A_u$ and $g_{ij} = M_{ij}$ for $ij \in E_u$. Then f is a circulation. Indeed, for $i \in V$,

$$\sum_{ij \in A_u} f_{ij} = -\sum_{ij \in A_u} M_{ij}u_j^{\mathsf{T}} Ru_i = -\sum_{j \in V} M_{ij}u_j^{\mathsf{T}} Ru_i = \Big(-\sum_{j \in V} M_{ij}u_j\Big)^{\mathsf{T}} Ru_i = 0.$$

Furthermore, $M(u,f,g) = M$ by simple computation. □

Note that the G-matrix $M(u,f,g)$ is well-signed if and only if f is a positive circulation on (V, A_u) and $g < 0$. Thus,

Corollary 11 *Let $G = (V,E)$ be a graph, let $u: V \to \mathbb{R}^2$ be a representation of V by nonzero vectors. Then*

$$\mathcal{W}_u = \big\{M(u,f,g) : f \text{ is a positive circulation on } (V,A_u), \\ g: E_u \to \mathbb{R},\ g < 0\big\}.$$

In particular, it follows that $\mathcal{W}_u \neq \emptyset$ if and only if A_u carries a positive circulation. This happens if and only if each arc in A_u is contained in a directed cycle in A_u; that is, if and only if each component of the directed graph (V,A_u) is strongly connected.

The signature of eigenvalues of $M(u,f,g)$ is a more difficult question, but we can say something about $M(u,0,g)$ if $g < 0$. Let H be a connected component of the graph (V,E_u), and let M_H be the submatrix of $M(u,0,g)$ formed by the rows and columns whose index belongs to $V(H)$. Then M_H is a well-signed H-matrix. The vectors u_i representing nodes $i \in V(H)$ are contained in a single line through the origin. Lemma 2 implies that M_H has at least one negative eigenvalue unless $u(V(H))$ is contained in a semiline starting at the origin. Let us call such a component *degenerate*. Then we can state:

Lemma 12 *Let $u: V \to \mathbb{R}^2$ be a representation of V with nonzero vectors, and let $g: E_u \to \mathbb{R}$ be a function with negative values. Then the number of negative eigenvalues of $M(u,0,g)$ is at least the number of nondegenerate components of (V,E_u).*

4.2 Shifting the Origin

Consider the cell complex made by the (two-way infinite) lines through distinct points u_i and u_j with $ij \in E$. The 1- and 2-dimensional cells are called 1-*cells* and 2-*cells*, respectively. Two cells c and d are *incident* if $d \subseteq \overline{c} \setminus c$ or $c \subseteq \overline{d} \setminus d$.

Two points p and q belong to the same cell if and only if $A_{u-p} = A_{u-q}$ and $E_{u-p} = E_{u-q}$. Hence, for any cell c, we can write A_c and E_c for A_{u-p} and E_{u-p}, where p is an arbitrary element of c. For any cell c, set $\mathcal{W}_c := \bigcup_{p \in c} \mathcal{W}_{u-p}$. It follows by Lemma 10 that if $\mathcal{W}_c \neq \emptyset$, then $\mathcal{W}_{u-p} \neq \emptyset$ for every $p \in c$. It also follows that $\mathcal{W}_c$ is connected for each cell c, as it is the range of the continuous function $M(u-p,f,g)$ on the connected topological space of triples (p,f,g) where $p \in c, f$ is a positive circulation on A_c, and g is a negative function on E_c.

The following lemma is an essential tool in the proof.

Lemma 13 *Let c be a cell with $\mathcal{W}_c \neq \emptyset$ and let $q \in \overline{c}$. Then $M(u-q,0,g) \in \overline{\mathcal{W}_c}$ for some negative function g on E_{u-q}.*

Proof Choose any $p \in c$. Note that $q \in \overline{c}$ implies that $E_{u-p} \subseteq E_{u-q}$. Let $M \in \mathcal{W}_{u-p}$. Define $g(ij) = M_{ij}$ for $ij \in E_{u-q}$ and let $N = M(u-q, 0, g)$. We prove that N belongs to $\overline{\mathcal{W}}_c$.

By Lemma 10 we can write $M = M(u-p, f, g')$ with some positive circulation f on A_{u-p} and negative function g' on E_{u-p}. Define $g(ij) = M_{ij}$ for $ij \in E_{u-q}$ and let $N = M(u-q, 0, g)$. For $\alpha \in (0, 1]$, define $p_\alpha = (1-\alpha)q + \alpha p$, and consider the G-matrices $M_\alpha = M(u-p_\alpha, \alpha f, g')$. Clearly $M_\alpha \in \mathcal{W}_c$. We show $\lim_{\alpha \to 0} M_\alpha = N$.

Let $f_0 = \max_{ij \in E} |f_{ij}|$, $\ell = \min_{u_i \neq u_j} |u_i - u_j|$, $\beta = \max_{i,j} |u_i|/|u_j|$, and let δ denote the distance of q from the closest edge in A_{u-q}. Let $0 < \alpha \leq \delta/|q-p|$. It suffices to prove that

$$\|M_\alpha - N\|_\infty \leq \frac{4\alpha\beta f_0}{\delta\ell}, \tag{6}$$

which implies that $M_\alpha \to N$ as $\alpha \to 0$.

- If $ij \in E_{u-p}$, then $(M_\alpha)_{ij} = N_{ij} = g'(ij)$, independently of α.
- If $ij \in E_{u-q} \setminus E_{u-p}$, then for each $\alpha \in (0, 1]$ we have $ij \notin E_{u-\alpha p}$. The points u_i, u_j, and q are collinear, hence $T(u-p_\alpha)_{ij} = \alpha T(u-p)_{ij}$ for each $\alpha \in (0, 1]$. Thus

$$(M_\alpha)_{ij} = \frac{-\alpha f_{ij}}{T(u-p_\alpha)_{ij}} = \frac{-f_{ij}}{T(u-p)_{ij}} = M_{ij} = g_{ij} = N_{ij}. \tag{7}$$

- If $ij \in E \setminus E_{u-q}$, then $N_{ij} = 0$ and

$$|T(u-p_\alpha)_{ij}| \geq |T(u-q)_{ij}| - \frac{1}{2}|q-p_\alpha|\,|u_i-u_j| \geq \frac{1}{2}(\delta - \alpha|q-p|)|u_i-u_j| \geq \frac{1}{4}\delta\ell.$$

  So

$$|N_{ij} - (M_\alpha)_{ij}| = |(M_\alpha)_{ij}| \leq \alpha\frac{4f_0}{\delta\ell}.$$

- If $i, j \in V$ with $ij \notin E$ and $i \neq j$, then $(M_\alpha)_{ij} = 0 = N_{ij}$.
- For the diagonal, (1) gives that

$$|N_{ii} - (M_\alpha)_{ii}| \leq \sum_{j \in N(i)} |N_{ij} - (M_\alpha)ij| \frac{|u_j^\mathsf{T} u_i|}{u_i^\mathsf{T} u_i} \leq \alpha\beta\frac{4f_0}{\delta\ell}.$$

This proves (6). □

Corollary 14 *Let c be a cell with $\mathcal{W}_c \neq \emptyset$ and $q \in \overline{c}$. Then for every matrix $M \in \mathcal{W}_{u-q}$ there is a matrix $M' \in \mathcal{W}_{u-q} \cap \overline{\mathcal{W}}_c$ that differs from M only on entries corresponding to edges in E_{u-q} and on the diagonal entries.*

Proof By Lemma 10 we can write $M = M(u-q, f, g)$ with some positive circulation f on A_{u-q} and negative function g on E_{u-q}. By Lemma 13, there is a negative function g' on E_{u-q} such that $M(u-q, 0, g') \in \overline{\mathcal{W}_c}$. There are points $p_k \in c$ and matrices $M_k \in \mathcal{W}_{u-p_k}$ such that $M_k \to M(u-q, 0, g')$ as $k \to \infty$. Then $M_k + M(u-p_k, f, 0)$ belongs to $\mathcal{W}_{u-p_k}$ and $M_k + M(u-p_k, f, 0) \to M(u-q, 0, g') + M(u-q, f, 0) = M(u-q, f, g')$ as $k \to \infty$, showing that $M' = M(u-q, f, g')$ belongs to $\overline{\mathcal{W}_c}$. Furthermore, $M - M' = M(u-q, 0, g-g')$ is nonzero on entries in E_{u-q} and on the diagonal entries only. □

Corollary 15 *If c and d are incident cells, then $\mathcal{W}_c \cup \mathcal{W}_d$ is connected.*

Proof We may assume that $d \subseteq \overline{c} \setminus c$, and that both $\mathcal{W}_c$ and $\mathcal{W}_d$ are nonempty (otherwise the assertion follows from the connectivity of $\mathcal{W}_c$ and $\mathcal{W}_d$).

Choose $q \in d$. Since $\mathcal{W}_d \neq \emptyset$, Corollary 14 implies that $\mathcal{W}_d$ and $\overline{\mathcal{W}_c}$ intersect, and by the connectivity of $\mathcal{W}_c$ and $\mathcal{W}_d$, this implies that $\mathcal{W}_c \cup \mathcal{W}_d$ is connected. □

Call a segment σ in the plane *separating*, if σ connects points u_a and u_b for some $a, b \in V$, with the property that $V \setminus \{a, b\}$ can be partitioned into two nonempty sets X and Y such that no edge of G connects X and Y and such that the sets $\{u_i \mid i \in X\}$ and $\{u_i \mid i \in Y\}$ are on distinct sides of the line through σ. Note that this implies that σ is a 1-cell.

Lemma 16 *Let G be a connected graph, and let σ be a separating segment connecting u_i and u_j, with incident 2-cells R and Q. If $\mathcal{W}_\sigma \cup \mathcal{W}_R \neq \emptyset$, then A_Q contains a directed circuit traversing ij.*

Proof We may assume that σ connects u_1 and u_2, and that edge 12 of G is oriented from 1 to 2 in A_Q. Let ℓ be the line through σ, and let H and H' be the open halfplanes with boundary ℓ containing Q and R, respectively.

Choose $p \in \sigma \cup R$ with $\mathcal{W}_{u-p} \neq \emptyset$. Note that A_Q and A_{u-p} differ only for edge 12. Any edge $ij \neq 12$ has the same orientation in A_Q as in A_{u-p}.

Since H contains points u_i, since G is connected, and since ℓ crosses no u_iu_j with $ij \in E$, G has an edge $1k$ or $2k$ with $u_k \in H$. By symmetry, we can assume that $2k$ is an edge. Then in A_{u-p}, edge $2k$ is oriented from 2 to k. As $\mathcal{W}_{u-p} \neq \emptyset$, A_{u-p} has a positive circulation. So A_{u-p} contains a directed circuit D containing $2k$. The edge preceding $2k$, say $j2$, must have $u_j \in H'$, as p belongs to $\sigma \cup R$. Therefore, since $\{1, 2\}$ separates nodes k and j, D traverses node 1. So the directed path in D from 2 to 1 together with the edge 12 forms the required directed circuit C in A_{u-q}. □

Corollary 17 *Let G be a connected graph, let σ be a separating segment, and let R be a 2-cell incident with σ. Then $\mathcal{W}_\sigma \neq \emptyset$ if and only if $\mathcal{W}_R \neq \emptyset$.*

Proof Let σ connect u_1 and u_2. If $\mathcal{W}_\sigma \neq \emptyset$, then A_σ has a positive circulation f'. By Lemma 16, A_R contains a directed circuit C traversing 12. Let f be the incidence vector of C. Then $f' + f$ is a positive circulation on A_R. So $\mathcal{W}_R \neq \emptyset$.

Conversely, if $\mathcal{W}_R \neq \emptyset$, then A_R has a positive circulation f. By Lemma 16, A_R contains a directed cycle through the arc 21, which gives a directed path P from 1 to 2 not using 12. It follows that by rerouting f_{12} over P, we obtain a positive circulation on A_σ, showing that $\mathcal{W}_\sigma \neq \emptyset$. □

4.3 Outerplanar Nullspace Embeddings

Let $G = (V, E)$ be a graph. A mapping $u : V \to \mathbb{R}^2$ is called *outerplanar* if its extension to the edges gives an embedding of G in the plane, and each u_i is incident with the unbounded face of this embedding.

Theorem 18 *Let G be a 2-connected graph with $\kappa(G) = 2$. Then the normalized nullspace representation defined by any well-signed G-matrix with one negative eigenvalue and with corank 2 is an outerplanar embedding of G.*

Proof Let u be such a normalized nullspace representation (this exists by Corollary 7). Let K be the convex hull of $u(V)$. Since all u_i have unit length, each u_i is a vertex of K. We define a *diagonal edge* as the line segment connecting points $u_i \neq u_j$, where $ij \in E$. We don't know at this point that the points u_i are different and that diagonal edges do not cross; so the same diagonal edge may represent several edges of G, and may consist of several 1-cells.

Let P denote the set of points $p \in \mathbb{R}^2 \setminus u(V)$ with $\mathcal{W}^1_{u-p} \neq \emptyset$. Clearly, the origin belongs to P. Lemma 2(b) implies that

Claim 1 *P is contained in the interior of K.*
(It will follow below that P is equal to the interior of K.)

Consider again the cell complex into which the diagonal edges cut K. By the connectivity of the sets $\mathcal{W}_c$ and by Lemma 3, P is a union of cells.

Claim 2 *$\overline{P}$ cannot contain a point $u_i = u_j$ for two distinct nodes i and j.*

Indeed, since $u_i = u_j$ is a vertex of the convex hull of $u(V)$, we can choose $p \in P$ close enough to u_i so that it is not in the convex hull of $u(V) \setminus \{u_i\}$. This, however, contradicts Lemma 4.

Claim 3 *No point $p \in \overline{P} \setminus u(V)$ is contained in two different diagonal edges.*

Indeed, consider any cell $c \subseteq P$ with $p \in \overline{c}$. Since $\mathcal{W}_c \neq \emptyset$, Lemma 13 implies that there is a negative function g on E_{u-p} such that $M(u - p, 0, g) \in \overline{\mathcal{W}_c}$. As all matrices in $\mathcal{W}_c$ have exactly one negative eigenvalue, $M(u - p, 0, g)$ has at most one negative eigenvalue. Lemma 12 implies that (V, E_{u-p}) has at most one nondegenerate component. But every diagonal containing p is contained in a nondegenerate component of (V, E_{u-p}), and these components are different for different diagonals, so p can be contained in at most one diagonal. This proves Claim 3.

It is easy to complete the proof now. Clearly, P is bounded by one or more polygons. Let p be a vertex of $\overline{P}$, and assume that $p \notin u(V)$. Then p belongs to two diagonals (defining the edges of P incident with p), contradicting Claim 3. Thus all vertices of P are contained in $u(V)$. This implies that $\overline{P}$ is a convex polygon spanned by an appropriate subset of $u(V)$.

To show that $\overline{P} = K$, assume that the boundary of P has an edge σ contained in the interior of K and let $R \subseteq P$ be a 2-cell incident with σ, and let Q be the 2-cell incident with σ on the other side. Clearly, $\mathcal{W}_R \neq \emptyset$, and by Corollary 17, $\mathcal{W}_\sigma \neq \emptyset$ and by the same Corollary, $\mathcal{W}_Q \neq \emptyset$. The sets $\mathcal{W}_\sigma \cup \mathcal{W}_R$ and $\mathcal{W}_\sigma \cup \mathcal{W}_Q$

are connected by Corollary 15, and hence so is $\mathcal{W}_\sigma \cup \mathcal{W}_R \cup \mathcal{W}_Q$. We also know that $\mathcal{W}^1 \cap \mathcal{W}_R \neq \emptyset$. Since $\mathcal{W}^1$ is open and closed in $\mathcal{W}$ (Lemma 3, note that in this case $\mathcal{W}^1 = \mathcal{W}^1 \cap \mathcal{W}^=$ as $\kappa(G) = 2$), we conclude that $\mathcal{W}^1 \cap \mathcal{W}_Q \neq \emptyset$, i.e., $Q \subseteq P$. But this contradicts the definition of σ.

Thus P is equal to the interior of K. Claim 2 implies that the points u_i are all different, and Claim 3 implies that the diagonals do not cross. ■

4.4 Algorithm

The considerations in this section give rise to a polynomial algorithm achieving the following.

Algorithm 6

Input: A 2-connected graph $G = (V, E)$.

Output: Either an outerplanar embedding $u : V \to \mathbb{R}^2$ of G, or a well-signed G-matrix with one negative eigenvalue and corank at least 3.

The algorithm progresses along the same lines as the algorithm in Sect. 3.2. We describe the main steps, omitting some details. It will be useful to remember that by Lemma 2(a), no well-signed G-matrix with two zero eigenvalues is positive semidefinite.

Step 1. We call Algorithm 5, which returns a well-signed G-matrix M with one negative and at least two zero eigenvalues (since the graph is not a path). If it has three zero eigenvalues, we are done, so suppose that this is not the case. We compute its nullspace representation u. We compute a positive circulation f on (G, u) and a negative function g on E_u such that $M = M(u, f, g)$, following the simple formulas in the proof of Lemma 10.

If $M(u, f, 0)$ has two negative eigenvalues, then the Interpolation Algorithms, applied with the matrix family $M(s) = (1 - t)M + tM(u, f, 0)$, returns a number $0 \leq s < 1$ for which $M(s)$ a well-signed G-matrix with one negative eigenvalue and corank at least 3. So suppose that $M(u, f, 0)$ has one negative eigenvalue.

Step 2. If there is an i with $u_i = 0$, then Algorithm 3 gives a matrix $M'' \in \mathcal{W}_u$ with one negative and at least three zero eigenvalues. So we may assume that $u_i \neq 0$ for every i. We scale M so that $|u_i| = 1$. (All we are going to use of this condition is that every u_i is a vertex of the convex hull K of the vectors u_i.) Lemma 2 implies that $0 \in \text{int}(K)$; let c be the cell containing 0 (this may be a point, and edge, or a polygon).

If u is an outerplanar embedding, we are done. Otherwise, we have either two nodes $i, j \in V$ with $u_i = u_j$, or two (diagonal) edges that intersect. Let $z \in K$ be a point that is either the intersection point of two diagonal edges, or $z = u_i = u_j$ for two nodes i and j. Choose z so that the number of diagonal edges separating z from 0 is minimal.

Step 3. If $z = 0$ (equivalently, c is 0-dimensional), then the origin is the intersection point of two diagonal edges, and hence $M(u, 0, -1)$ has at least two negative eigenvalues. So we can apply Algorithm 1 with the matrix family $tM + (1-t)M(u, 0, -1)$ (keeping u fixed).

Step 4. Suppose that we find two matrices $M \in \mathcal{W}^1_{u-p}$ and $M' \in \mathcal{W}^2_{u-q}$ where $p, q \in c$. Since p and q belong to the same cell, the matrix $M(u-q, f, g)$ is well defined and $M(u-q, f, g) \in \mathcal{W}_{u-q}$. If $M(u-q, f, g)$ has one negative eigenvalue, then we invoke Algorithm 1 with the family $(1-t)M' + tM(u-q, f, g)$ (keeping $u-q$ fixed). If $M(u-q, f, g)$ has at least two negative eigenvalues, then similarly invoke Algorithm 1 with the family $M(u - tp - (1-t)q, f, g)$ and $(u - tp - (1-t)q)$ for $0 \le t \le 1$.

Step 5. Suppose that no diagonal edge separates z from the origin, and z is the intersection point of at least two diagonal edges. Choose a number α such that

$$0 < \alpha < \min\left\{1, \frac{\delta}{|z|}, \frac{\delta \ell n}{4\beta f_0}\right\},$$

where the numbers β, f_0, δ, ℓ are defined as in the proof of Lemma 13 and are easily computed. As in the proof of Lemma 13, we construct a negative function g on E_{u-z} and a matrix $M_\alpha \in \mathcal{W}_{u-(1-\alpha)z}$ such that the matrix $N = M(u-(1-\alpha)z, 0, g)$ satisfies $\|M_\alpha - N\|_\infty < 4\alpha\beta f_0/(\delta\ell)$. The matrix N has at least two negative eigenvalues. Then elementary linear algebra gives that the matrix M_α has at least two negative eigenvalues. We conclude by Step 4.

Step 6. Suppose that every vertex of c is in $u(V)$, and c has a vertex $z = u_i = u_j$. Let q be a point in the interior of c but not in $\mathrm{conv}(u(V) \setminus \{z\}$. Then by Lemma 4, the matrix $M(u-q, f, g)$ has either corank at least 3 or two negative eigenvalues. In the first case, we are done; in the second, we invoke Step 4. So we may assume that z is not a vertex of c.

Step 7. If $c = [u_i, u_j]$ is a diagonal (intersecting no other diagonal), then let $\varepsilon > 0$ be small enough so that εz belongs to a region R bounded by c. By the construction in the proof of Lemma 17, we find a directed cycle C in G that passes every edge in the positive direction when viewed from R. Let h denote the unit flow around C, and let $M' = M(u - \varepsilon z, f + \varepsilon^2 h, 0) \in \mathcal{W}_{\varepsilon z}$.

If M' has one negative eigenvalue, then we can replace M by M' and u by $u - \varepsilon z$, to get an instance where the segment $[\varepsilon z, z]$ intersects fewer diagonal edges than $[0, z]$. If $M' \in \mathcal{W}^2_{u-\varepsilon z}$, then we apply the interpolation argument to the family $M(t) = M(u - tz, f + t^2 h, 0)$, using that $M(0)$ has one negative eigenvalue (as it is the limit of $M(u, f, \beta g)$ as $\beta \to 0$) and $M(1) = M'$. (The coefficient of h is t^2 to make sure that $M(t)$ depends continuously on t at $t = 0$.)

Step 8. So we may assume that c is a 2-dimensional polygon, 0 is an internal point of it, every vertex of c is the position of exactly one node, and so every edge of c is a full diagonal edge. Let q be the intersection point of $[0, z]$ with the boundary of c. Let $ij \in A_u$ be the edge for which $q \in [u_i, u_j]$, and let Q be the region on the other side of e, let C be a cycle through e in A_Q whose edges are counterclockwise when viewed from Q (constructed as in Lemma 17). Let h denote the unit flow around C, and let $M' = M(u - q, f + f_{ij}h, -1)$. Then $M' \in \mathcal{W}_{u-q}$. If it has one negative eigenvalue, then we can replace M by M' and u by $u - q$. If $M' \in \mathcal{W}^2_{u-q}$, then we apply the Interpolation Algorithm 1 to the matrix family $M(t) = M(tq, f + tf_{ij}h, 0)$.

Acknowledgements We thank Bart Sevenster for helpful discussion on $\kappa(G)$, and two referees for useful suggestions improving the presentation.

References

1. Y. Colin de Verdière, Sur un nouvel invariant des graphes et un critère de planarité. J. Combin. Theory Ser. B **50**, 11–21 (1990) [English transl.: On a new graph invariant and a criterion for planarity, in: *Graph Structure Theory*, ed. by N. Robertson, P. Seymour (American Mathematical Society, Providence, 1993), pp. 137–147]
2. R. Connelly, Rigidity and energy. Invent. Math. **66**, 11–33 (1982)
3. M. Laurent, A. Varvitsiotis, Positive semidefinite matrix completion, universal rigidity and the Strong Arnold Property. Linear Algebra Appl. **452**, 292–317 (2014)
4. L. Lovász, Steinitz representations and the Colin de Verdière number. J. Combin. Theory B **82**, 223–236 (2001)
5. L. Lovász, A. Schrijver, A Borsuk theorem for antipodal links and a spectral characterization of linklessly embeddable graphs. Proc. Am. Math. Soc. **126**, 1275–1285 (1998)
6. L. Lovász, A. Schrijver, On the null space of a Colin de Verdière matrix. Annales de l'Institut Fourier, Université de Grenoble **49**, 1017–1026 (1999)
7. N. Robertson, P. Seymour, R. Thomas, Sachs' linkless embedding conjecture. J. Combin. Theory Ser. B **64**, 185–227 (1995)
8. A. Schrijver, B. Sevenster, The Strong Arnold Property for 4-connected flat graphs. Linear Algebra Appl. **522**, 153–160 (2017)
9. H. van der Holst, A short proof of the planarity characterization of Colin de Verdière. J. Combin. Theory Ser. B **65**, 269–272 (1995)
10. H. van der Holst, Topological and spectral graph characterizations. Ph.D. thesis, University of Amsterdam, Amsterdam (1996)
11. H. van der Holst, L. Lovász, A. Schrijver, The Colin de Verdière graph parameter, in *Graph Theory and Combinatorial Biology*. Bolyai Society Mathematical Studies, vol. 7 (János Bolyai Mathematical Society, Budapest, 1999), pp. 29–85