Algorithms for Graph Rigidity and Scene Analysis

Alex R. Berg[1] and Tibor Jordán[2]*

[1] BRICS (Basic Research in Computer Science, Centre of the Danish National Research Foundation), University of Aarhus, Ny Munkegade, building 540, Aarhus C, Denmark. aberg@brics.dk
[2] Department of Operations Research, Eötvös University, Pázmány sétány 1/C, 1117 Budapest, Hungary. jordan@cs.elte.hu

Abstract. We investigate algorithmic questions and structural problems concerning graph families defined by 'edge-counts'. Motivated by recent developments in the unique realization problem of graphs, we give an efficient algorithm to compute the rigid, redundantly rigid, M-connected, and globally rigid components of a graph. Our algorithm is based on (and also extends and simplifies) the idea of Hendrickson and Jacobs, as it uses *orientations* as the main algorithmic tool.
We also consider families of bipartite graphs which occur in parallel drawings and scene analysis. We verify a conjecture of Whiteley by showing that $2d$-connected bipartite graphs are d-tight. We give a new algorithm for finding a maximal d-sharp subgraph. We also answer a question of Imai and show that finding a maximum size d-sharp subgraph is NP-hard.

1 Introduction

A d-dimensional *framework* is a straight line embedding of a graph $G = (V, E)$ in the d-dimensional Euclidean space. We may think of the edges as rigid bars and vertices as rotatable joints, and say that the framework is 'rigid' if it has no non-trivial deformations (for precise definitions, in terms of the 'rigidity matrix' of the framework, see [18]). If the coordinates are 'generic' then rigidity depends only on the graph. Characterising rigidity of frameworks (and graphs) is an old problem in statics (and graph theory). In two dimensions there is a combinatorial characterisation, see Theorem 2.

Rigidity of graphs plays an interesting role in unique graph realizations. Two frameworks of the same graph G are *equivalent* if corresponding edges of the two frameworks have the same length, and they are *congruent* if the distance between all corresponding pairs of vertices is the same. We say that a framework is a *unique realization* of G in R^d if every equivalent framework of G is also congruent.

* Supported by the Hungarian Scientific Research Fund no. T037547, F034930 and FKFP grant no. 0143/2001. Part of this research was done while the second author visited BRICS, University of Aarhus, Denmark.

G. Di Battista and U. Zwick (Eds.): ESA 2003, LNCS 2832, pp. 78–89, 2003.

Hendrickson [6] proved that if G has a unique generic realization then G is $(d+1)$-connected and redundantly rigid (that is, $G-e$ is rigid for all $e \in E(G)$). See also [2]. It was proved in [10] that for $d = 2$ these conditions are sufficient to guarantee that every generic framework of G is a unique realization, in which case the graph is called *globally rigid.* Thus unique realizability is a generic property in two dimensions. This recent result motivated us to give an algorithm for finding the finer structure of rigid, redundantly rigid, and M-connected components of a graph in two dimensions (see Section 3 for the definitions), and for identifying the maximal globally rigid subgraphs of G.

The basic algorithmic questions concerning rigidity were answered many years ago. There exist polynomial algorithms for testing rigidity as well as for computing the 'degree of freedom' in $O(n^2)$ time, where n is the number of vertices of G. The algorithms of Sugihara [14] and Hendrickson [6] used techniques from matching theory, while Imai [9] used network flows. Gabow and Westermann [4] used matroid sums and showed that the M-connected components can also be found in $O(n^2)$ time.

Our algorithm also runs in $O(n^2)$ time and in some sense it is similar to the previous algorithms. However, it requires no auxiliary graphs or digraphs, or matroids, to work on. Thus it is perhaps simpler and easier to visualize. It works with *orientations* of G and only performs reachability searches (and reorients directed paths). The idea of such an algorithm is due to Hendrickson and Jacobs [7], who gave a basic version of this algorithm as a so-called "pebble game" (see also [11]) for finding the rigid components in $O(n^2)$ time. We simplify the terminology of [7] and make it more suitable for finding other substructures as well as for extensions to other families of graphs defined by edge counts.

A different area of discrete applied geometry, where such families occur, is parallel drawings. A *polyhedral incidence structure* is a bipartite graph $S = (V, F; I)$ where the colour classes V and F represent the *vertices* and *faces*, respectively, and the edge set I represents the vertex-face incidences. A *d-scene* of S assigns points of R^d to the vertices and hyperplanes to the faces such that all the incidence constraints are satisfied. For a set of generic normals for the faces, one may ask whether there is a d-scene of S with the given normals. The existence of a such a (non-trivial) d-scene depends only on S. Whiteley [16,17] characterised the bipartite graphs whose edge set is a base in the corresponding matroid (called *minimally d-tight* graphs) as well as the graphs for which such a d-scene exists with distinct points assigned to distinct vertices (called *d-sharp*). A minimally d-tight graph satisfies the count $|I| = d|V| + |F| - d$ and $|I'| \leq d|V(I')| + |F(I')| - d$, for all $\emptyset \neq I' \subseteq I$, where $V(I')$ and $F(I')$ denotes the number of vertices induced by I' in V and F, respectively. A d-sharp graph has $|I'| \leq d|V(I')| + |F(I')| - d$, for all $I' \subseteq I$ with $|V(I')| \geq 2$. Whiteley conjectured that every $2d$-connected incidence graph has a minimally d-tight spanning subgraph [18, p.211]. We prove this conjecture.

We also investigate d-sharp graphs, but in a different context. The basic problem of scene analysis of polyhedral pictures is how to reconstruct a 3-dimensional polyhedron from a given planar projection, a so-called line drawing, by setting

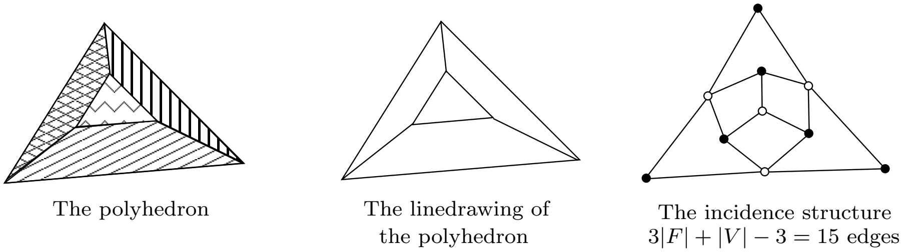

Fig. 1. A line drawing with a minimally 3-tight incidence graph.

the third coordinates of the vertices so that vertices that belong to the same face in the line drawing are coplanar (and such that each pair of faces sharing a vertex have distinct planes). Assuming generic points for the line drawing Sugihara [15] and Whiteley [17] characterised the reconstructible line drawings in terms of their bipartite incidence graph $P = (V, F; I)$, which encodes the vertex-face incindences. This characterisation turns out to be 'polar' to parallel drawings: there is a proper 'lifting' for P if and only if the bipartite graph obtained from P by interchanging the role of V and F is 3-sharp. Due to numerical errors, line drawings produced by computers are 'generic', and may produce non-reconstructible pictures. To cope with this problem Sugihara [15] developed an algorithm which attempts to reconstruct the polyhedron from a maximal 3-sharp subgraph of the incidence graph of the line drawing, see also Imai [9]. To illustrate the applicability of the orientation based approach to other edge counts, we give a different algorithm for finding a maximal d-sharp subgraph. Furthermore, answering an open question posed by Imai [9], we show that finding a maximum size d-sharp subgraph is NP-hard for $d \geq 2$.

2 Orientations with Upper Bounds on the In-Degrees

Let $G = (V, E)$ be a graph. An *orientation* $D = (V, A)$ of G is a directed graph obtained from G by replacing each edge uv by a directed edge (directed from u to v or from v to u). For a subset $X \subseteq V$ let $G[X]$ denote the subgraph of G induced by X, and let $i_G(X)$ (or simply $i(X)$) denote the number of edges in $G[X]$. If $D = (V, A)$ is a directed graph and $X \subseteq V$ then $\rho_D(X)$ denotes the number of directed edges entering X. This is the *in-degree* of X. The in-degree of a vertex v is denoted by $\rho_D(v)$. Let $g : V \to Z_+$ assign non-negative integers to the vertices of G. For $X \subseteq V$ we use the notation $g(X) := \sum_{v \in X} g(v)$. We say that an orientation D of G is a *g-orientation* if $\rho_D(v) \leq g(v)$ holds for all $v \in V$.

The next result, due to Frank and Gyárfás, characterises when G has an orientation satisfying given upper bounds on the in-degrees. We present its simple proof, since it illustrates the basic steps of our algorithm: searching for reachable vertices and reorienting directed paths.

Theorem 1. *[3] Let $G = (V, E)$ be a graph and $g : V \to Z_+$. Then G has a g-orientation if and only if*

$$i(X) \leq g(X) \quad \textit{for all} \quad X \subseteq V. \tag{1}$$

Proof. To see necessity suppose that D is a g-orientation of G and let $X \subseteq V$. Then $i(X) = \sum_{v \in X} \rho_D(v) - \rho_D(X) \leq g(X)$.

To prove sufficiency suppose that (1) holds and choose an orientation D' of G for which $h(D') := \sum_{v \in V}(\rho(v) - g(v))^+$ is as small as possible, where $x^+ := \max\{x, 0\}$ for some integer x. If $h(D') = 0$ then D' is a g-orientation. Otherwise there is a vertex s with $\rho_{D'}(s) > g(s)$. Let S denote the set of vertices from which there is a directed path to s in D'. Clearly, $\rho_{D'}(S) = 0$. If there is a vertex $t \in S$ with $\rho_{D'}(t) < g(t)$ then by reorienting the edges of a directed path from t to s we obtain an orientation D'' with $h(D'') = h(D') - 1$, contradicting the choice of D'. Thus we have $\rho_{D'}(v) \geq g(v)$ for each vertex $v \in S$, and hence, since $\rho_{D'}(s) > g(s)$, we obtain $i(S) = \sum_{v \in S} \rho_{D'}(v) - \rho_{D'}(S) > \sum_{v \in S} g(v) = g(S)$, contradicting (1). This proves the theorem. □

This proof leads to an algorithm for finding a g-orientation, if exists. It shows that if (1) holds then any orientation D' of G can be turned into a g-orientation by finding and reorienting directed paths $h(D')$ times. Such an elementary step (which decreases h by one) can be done in linear time.

3 Rigid Graphs and the Rigidity Matroid

The following combinatorial characterization of two-dimensional rigidity is due to Laman. A graph G is said to be *minimally rigid* if G is rigid, and $G - e$ is not rigid for all $e \in E$. A graph is *rigid* if it has a minimally rigid spanning subgraph.

Theorem 2. *[12] $G = (V, E)$ is minimally rigid if and only if $|E| = 2|V| - 3$ and*

$$i(X) \leq 2|X| - 3 \textit{ for all } X \subseteq V \textit{ with } |X| \geq 2. \tag{2}$$

In fact, Theorem 2 characterises the bases of the *rigidity matroid* of the complete graph on vertex set V. In this matroid a set of edges S is *independent* if the subgraph induced by S satisfies (2). The *rigidity matroid of* G, denoted by $\mathcal{M}(G) = (E, \mathcal{I})$, is the restriction of the rigidity matroid of the complete graph to E. Thus G is rigid if and only if E has rank $2|V| - 3$ in $\mathcal{M}(G)$. If G is rigid and $H = (V, E')$ is a spanning subgraph of G, then H is minimally rigid if and only if E' is a base in $\mathcal{M}(G)$.

3.1 A Base, the Rigid Components, and the Rank

To test whether G is rigid (or more generally, to compute the rank of $\mathcal{M}(G)$) we need to find a base of $\mathcal{M}(G)$. This can be done greedily, by building up a maximal independent set by adding (or rejecting) edges one by one. The key of this procedure is the independence test: given an independent set I and an edge e, check whether $I+e$ is independent. With Theorem 1 we can do this in linear time as follows (see also [7]).

Let $g_2 : V \to Z_+$ be defined by $g_2(v) = 2$ for all $v \in V$. For two vertices $u, v \in V$ let $g_2^{uv} : V \to Z_+$ be defined by $g_2^{uv}(u) = g_2^{uv}(v) = 0$, and $g_2^{uv}(w) = 2$ for all $w \in V - \{u, v\}$.

Lemma 1. *Let $I \subset E$ be independent and let $e = uv$ be an edge, $e \in E - I$. Then $I + e$ is independent if and only if (V, I) has a g_2^{uv}-orientation.*

Proof. Let $H = (V, I)$ and $H' = (V, I + e)$. First suppose that $I + e$ is not independent. Then there is a set $X \subseteq V$ with $i_{H'}(X) \geq 2|X| - 2$. Since I is independent, we must have $u, v \in X$ and $i_H(X) = 2|X| - 3$. Hence $i_H(X) = 2|X| - 3 > g_2^{uv}(X) = 2|X| - 4$, showing that H has no g_2^{uv}-orientation.

Conversely, suppose that $I+e$ is independent, but H has no g_2^{uv}-orientation. By Theorem 1 this implies that there is a set $X \subseteq V$ with $i_H(X) > g_2^{uv}(X)$. Since $i_H(X) \leq 2|X|-3$ and $g_2^{uv}(X) = 2|X|-2|X \cap \{u, v\}|$, this implies $u, v \in X$ and $i_H(X) = 2|X|-3$. Then $i_{H'}(X) = 2|X|-2$, contradicting the fact that $I+e$ is independent. □

A *weak g_2^{uv}-orientation* D of G satisfies $\rho_D(w) \leq 2$ for all $w \in V - \{u, v\}$ and has $\rho_D(u) + \rho_D(v) \leq 1$. It follows from the proof that a weak g_2^{uv}-orientation of (V, I) always exists.

If we start with a g_2-orientation of $H = (V, I)$ then the existence of a g_2^{uv}-orientation of H can be checked by at most four elementary steps (reachability search and reorientation) in linear time. Note also that H has $O(n)$ edges, since I is independent.

This gives rise to a simple algorithm for computing the rank of E in $\mathcal{M}(G)$. By maintaining a g_2-orientation of the subgraph of the current independent set I, testing an edge needs only $O(n)$ time, and hence the total running time is $O(nm)$, where $m = |E|$. We shall improve this to $O(n^2)$ by identifying large rigid subgraphs.

We say that a maximal rigid subgraph of G is a *rigid component* of G. Clearly, every edge belongs to some rigid component, and rigid components are induced subgraphs. Since the union of two rigid subgraphs sharing an edge is also rigid, the edge sets of the rigid components partition E.

We can maintain the rigid components of the set of edges considered so far as follows. Let I be an independent set, let $e = uv$ be an edge with $e \in E - I$, and suppose that $I + e$ is independent. Let D be a g_2^{uv}-orientation of (V, I). Let $X \subseteq V$ be the maximal set with $u, v \in X, \rho_D(X) = 0$, and such that $\rho_D(x) = 2$ for all $x \in X - \{u, v\}$. Clearly, such a set exists, and it is unique. It can be found by identifying the set $V_1 = \{x \in V - \{u, v\} : \rho_D(x) \leq 1\}$, finding the set $\hat{V}_1$ of vertices reachable from V_1 in D, and then taking $X = V - \hat{V}_1$. The next lemma shows how to update the set of rigid components when a new edge e is added to I.

Lemma 2. *Let $H' = (V, I + e)$. Then $H'[X]$ is a rigid component of H'.*

Thus, when we add e to I, the set of rigid components is updated by adding $H'[X]$ and deleting each component whose edge set is contained by the edge set of $H'[X]$. Maintaing this list can be done in linear time. Furthermore, we can reduce the total running time to $O(n^2)$ by performing the independence test for $I + e$ only if e is not spanned by any of the rigid components on the current list (and otherwise rejecting e, since $I + e$ is clearly dependent).

3.2 The M-Circuits and the Redundantly Rigid Components

Given a graph $G = (V, E)$, a subgraph $H = (W, C)$ is said to be an *M-circuit in G* if C is a circuit (i.e. a minimal dependent set) in $\mathcal{M}(G)$. G is an *M-circuit* if E is a circuit in $\mathcal{M}(G)$. By using (2) one can deduce the following properties.

Lemma 3. *Let $G = (V, E)$ be a graph without isolated vertices. Then G is an M-circuit if and only if $|E| = 2|V| - 2$ and $G - e$ is minimally rigid for all $e \in E$.*

A subgraph $H = (W, F)$ is *redundantly rigid* if H is rigid and $H - e$ is rigid for all $e \in F$. M-circuits are redundantly rigid by Lemma 3(b). A *redundantly rigid component* is either a maximal redundantly rigid subgraph of G (in which case the component is *non-trivial*) or a subgraph consisting of a single edge e, when e is contained in no redundantly rigid subgraph of G (in which case it is *trivial*). The redundantly rigid components are induced subgraphs and their edge sets partition the edge set of G. See Figure 2 for an example. An edge $e \in E$ is a *bridge* if e belongs to all bases of $\mathcal{M}(G)$. It is easy to see that each bridge e is a trivial redundantly rigid component. Let $B \subseteq E$ denote the set of bridges in G. The key to finding the redundantly rigid components efficiently is the following lemma.

Lemma 4. *The set of non-trivial redundantly rigid components of G is equal to the set of rigid components of $G' = (V, E - B)$.*

Thus we can identify the redundantly rigid components of G by finding the bridges of G and then finding the rigid components of the graph $G - B$.

3.3 The M-Connected Components and Maximal Globally Rigid Subgraphs

Given a matroid $\mathcal{M} = (E, \mathcal{I})$, one can define a relation on E by saying that $e, f \in E$ are related if $e = f$ or there is a circuit C in $\mathcal{M}$ with $e, f \in C$. It is well-known that this is an equivalence relation. The equivalence classes are called the *components* of $\mathcal{M}$. If $\mathcal{M}$ has at least two elements and only one component then $\mathcal{M}$ is said to be *connected*. Note that the trivial components (containing only one element) of $\mathcal{M}$ are exactly the bridges of G.

We say that a graph $G = (V, E)$ is *M-connected* if $\mathcal{M}(G)$ is connected. The *M-connected components* of G are the subgraphs of G induced by the components of $\mathcal{M}(G)$. The M-connected components are also edge-disjoint induced subgraphs. They are redundantly rigid.

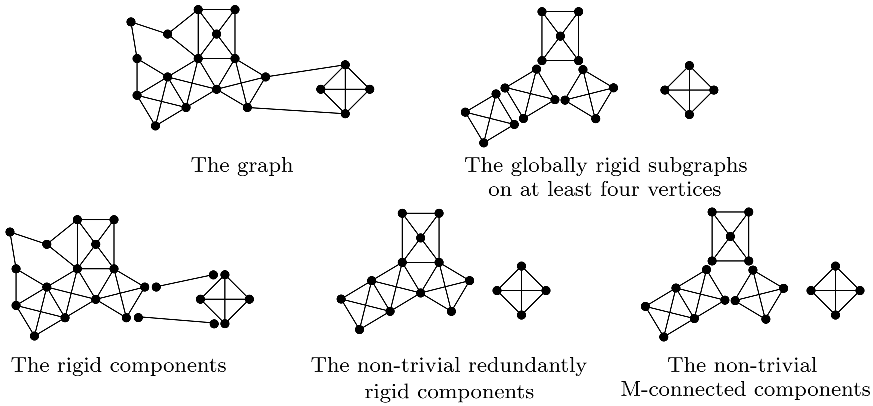

Fig. 2. Decompositions of a graph.

To find the bridges and M-connected components we need the following observations. Suppose that I is independent but $I + e$ is dependent. The *fundamental circuit* of e with respect to I is the (unique) circuit contained in $I + e$. Our algorithm will also identify a set of fundamental circuits with respect to the base I that it outputs. To find the fundamental circuit of $e = uv$ with respect to I we proceed as follows. Let D be a weak g_2^{uv}-orientation of (V, I) (with $\rho_D(v) = 1$, say). As we noted earlier, such an orientation exists. Let $Y \subseteq V$ be the (unique) minimal set with $u, v \in Y, \rho_D(Y) = 0$, and such that $\rho_D(x) = 2$ for all $x \in Y - \{u, v\}$. This set exists, since $I + e$ is dependent. Y is easy to find: it is the set of vertices that can reach v in D.

Lemma 5. *The edge set induced by Y in $(V, I + e)$ is the fundamental circuit of e with respect to I.*

Thus if $I + e$ is dependent, we can find the fundamental circuit of e in linear time. Our algorithm will maintain a list of M-connected components and compute the fundamental circuit of $e = uv$ only if u and v are not in the same M-connected component. Otherwise e is classified as a non-bridge edge. When a new fundamental circuit is found, its subgraph will be merged into one new M-connected component with all the current M-connected components whose edge set intersects it. It can be seen that the final list of M-connected components will be equal to the set of M-connected components of G, and the edges not induced by any of these components will form the set of bridges of G. It can also be shown that the algorithm computes $O(n)$ fundamental circuits, so the total running time is still $O(n^2)$. The algorithm can also determine an *ear-decomposition* of $\mathcal{M}(G)$ (see [10]), for an M-connected graph G, within the same time bound.

Thus to identify the maximal globally rigid subgraphs on at least four vertices we need to search for the maximal 3-connected subgraphs of the M-connected components of G. This can be done in linear time by using the algorithm of Hopcroft and Tarjan [8] which decomposes the graph into its 3-connected blocks.

4 Tight and Sharp Bipartite Graphs

Let $G = (A, B; E)$ be a bipartite graph. For subsets $W \subseteq A \cup B$ and $F \subseteq E$ let $W(F)$ denote the set of those vertices of W which are incident to edges of F. We say that G is *minimally d-tight* if $|E| = d|A| + |B| - d$ and for all $\emptyset \neq E' \subseteq E$ we have

$$|E'| \leq d|A(E')| + |B(E')| - d. \tag{3}$$

G is called *d-tight* if it has a minimally d-tight spanning subgraph. It is not difficult to show that the subsets $F \subseteq E$ for which every $\emptyset \neq E' \subseteq F$ satisfies (3) form the independent sets of a matroid on groundset E. By calculating the rank function of this matroid we obtain the following characterization.

Theorem 3. *[17] $G = (A, B; E)$ is d-tight if and only if*

$$\sum_{i=1}^{t}(d \cdot |A(E_i)| + |B(E_i)| - d) \geq d|A| + |B| - d \tag{4}$$

for all partitions $\mathcal{E} = \{E_1, E_2, ..., E_t\}$ of E.

4.1 Highly Connected Graphs Are *d*-Tight

Lovász and Yemini [13] proved that 6-connected graphs are rigid. A similar result, stating that $2d$-connected bipartite graphs are d-tight, was conjectured by Whiteley [16,18]. We prove this conjecture by using an approach similar to that of [13]. We say that a graph $G = (V, E)$ is *k-connected in W*, where $W \subseteq V$, if there exist k openly disjoint paths in G between each pair of vertices of W.

Theorem 4. *Let $G = (A, B; E)$ be $2d$-connected in A, for some $d \geq 2$, and suppose that there is no isolated vertex in B. Then G is d-tight.*

Proof. For a contradiction suppose that G is not d-tight. By Theorem 3 this implies that there is a partition $\mathcal{E} = \{E_1, E_2, ..., E_t\}$ of E with $\sum_{i=1}^{t}(d \cdot |A(E_i)| + |B(E_i)| - d) < d|A| + |B| - d$. Since G is $2d$-connected in A and there is no isolated vertex in B, we have $d|A(E)| + |B(E)| - d = d|A| + |B| - d$. Thus $t \geq 2$ must hold.

Claim. Suppose that $A(E_i) \cap A(E_j) \neq \emptyset$ for some $1 \leq i < j \leq t$. Then $d|A(E_i)| + |B(E_i)| - d + d|A(E_j)| + |B(E_j)| - d \geq d|A(E_i \cup E_j)| + |B(E_i \cup E_j)| - d$.

The claim follows from the inequality: $d|A(E_i)| + |B(E_i)| - d + d|A(E_j)| + |B(E_j)| - d = d|A(E_i) \cup A(E_j)| + d|A(E_i) \cap A(E_j)| + |B(E_i) \cup B(E_j)| + |B(E_i) \cap B(E_j)| - 2d \geq d|A(E_i \cup E_j)| + |B(E_i \cup E_j)| - d$, where we used $d|A(E_i) \cap A(E_j)| \geq d$ (since $A(E_i) \cap A(E_j) \neq \emptyset$), and $|B(E_i) \cap B(E_j)| \geq 0$.

By the Claim we can assume that $A(E_i) \cap A(E_j) = \emptyset$ for all $1 \leq i < j \leq t$. Let $B' \subseteq B$ be the set of those vertices of B which are incident to edges from at least two classes of partition $\mathcal{E}$. Since $A(E_i) \cap A(E_j) = \emptyset$ for all $1 \leq i < j \leq t$, and $t \geq 2$, the vertex set $B'(E_i)$ separates $A(E_i)$ from $\cup_{j \neq i} A(E_j)$ for all $E_i \in \mathcal{E}$. Hence, since G is $2d$-connected in A, we must have

$$|B'(E_i)| \geq 2d \text{ for all } 1 \leq i \leq t. \tag{5}$$

To finish the proof we count as follows. Since $A(E_i) \cap A(E_j) = \emptyset$ for all $1 \leq i < j \leq t$, we have $\sum_1^t |A(E_i)| = |A|$. Hence $\sum_1^t (|B(E_i)| - d) < |B| - d$ follows, which gives

$$\sum_1^t (|B'(E_i)| - d) < |B'| - d. \tag{6}$$

Furthermore, it follows from (5) and the definition of B' that for every vertex $b \in B'$ we have

$$\sum_{E_i : b \in B(E_i)} (1 - \frac{d}{|B'(E_i)|}) \geq 2(1 - \frac{d}{2d}) = 1.$$

Thus

$$|B'| \leq \sum_{b \in B'} \sum_{E_i : b \in B(E_i)} (1 - \frac{d}{|B'(E_i)|}) = \sum_1^t |B'(E_i)|(1 - \frac{d}{|B'(E_i)|}) = \sum_1^t (|B'(E_i)| - d),$$

which contradicts (6). This proves the theorem. □

4.2 Testing Sharpness and Finding Large Sharp Subgraphs

By modifying the count in (3) slightly, we obtain a family of bipartite graphs which plays a central role in scene analysis (for parameter $d = 3$). We say that a bipartite graph $G = (A, B; E)$ is *d-sharp*, for some integer $d \geq 1$, if

$$|E'| \leq d|A(E')| + |B(E')| - (d + 1) \tag{7}$$

holds for all $E' \subseteq E$ with $|A(E')| \geq 2$. A set $F \subseteq E$ is *d-sharp* if it induces a d-sharp subgraph.

As it was pointed out by Imai [9], the count in (7) does not always define a matroid on the edge set of G. Hence to test d-sharpness one cannot directly apply the general framework which works well for rigidity and d-tightness. Sugihara [15] developed an algorithm for testing 3-sharpness and, more generally, for finding a *maximal* 3-sharp subset of E. Imai [9] improved the running time to $O(n^2)$. Their algorithms are based on network flow methods.

An alternative approach is as follows. Let us call a maximal d-tight subgraph of G a *d-tight component*. As in the case of rigid components, one can show that the d-tight components are pairwise edge-disjoint and their edge sets partition E. Moreover, by using the appropriate version of our orientation based algorithm, they can be identified in $O(n^2)$ time. The following lemma shows how to use these components to test d-sharpness (and to find a maximal d-sharp edge set) in $O(n^2)$ time.

Lemma 6. *Let $G = (A, B; E)$ be a bipartite graph and $d \geq 1$ be an integer. Then G is d-sharp if and only if each d-tight component H satisfies $|V(H) \cap A| = 1$.*

Proof. Necessity is clear from the definition of d-tight and d-sharp graphs. To see sufficiency suppose that each d-tight component H satisfies $|V(H) \cap A| = 1$, but G is not d-sharp. Then there exists a set $I \subseteq E$ with (i) $|A(I)| \geq 2$ and (ii) $|I| \geq d|A(I)| + |B(I)| - d$. Let I be a minimal set satisfying (i) and (ii). Suppose that I satisfies (ii) with strict inequality and let $e \in I$ be an edge. By the minimality of I the set $I - e$ must violate (i). Thus $|A(I - e)| = 1$. By (i), and since $d \geq 2$, this implies $|I| = |I - e| + 1 = |B(I - e)| + 1 \leq |B(I)| + 1 \leq d|A(I)| + |B(I)| - d$, a contradiction. Thus $|I| = d|A(I)| + |B(I)| - d$. The minimality of I (and the fact that each set I' with $|A(I')| = 1$ trivially satisfies $|I'| = d|A(I')| + |B(I')| - d$) implies that (3) holds for each non-empty subset of I. Thus I induces a d-tight subgraph H', which is included by a d-tight component H with $|V(H) \cap A| \geq |A(I)| \geq 2$, contradicting our assumption. □

Imai [9] asked whether a *maximum size* 3-sharp edge set of G can be found in polynomial time. We answer this question by showing that the problem is NP-hard.

Theorem 5. *Let $G = (A, B; E)$ be a bipartite graph, and let $d \geq 2$, $N \geq 1$ be integers. Deciding whether G has a d-sharp edge set $F \subseteq E$ with $|F| \geq N$ is NP-complete.*

Proof. We shall prove that the NP-complete VERTEX COVER problem can be reduced to our decision problem. Consider an instance of VERTEX COVER, which consists of a graph $D = (V, J)$ and an integer M (and the question is whether D has a vertex cover of size at most M). Our reduction is as follows. First we construct a bipartite graph $H_0 = (A_0, B_0; E_0)$, where $A_0 = \{c\}$, $B_0 = V$ corresponds to the vertex set of D, and $E_0 = \{cu : u \in B_0\}$. We call c the *center* of H_0. Thus H_0 is a star which spans the vertices of D from a new vertex c at the center. The bipartite graph $H = (A, B; E)$ that we construct next is obtained by adding a clause to H_0 for each edge of D. The clause for an edge $e = uv \in J$ is the following (see Figure 3). $f_1, f_2, \ldots, f_{d-2}$ are the new vertices of B and x is the new vertex of A in the clause (these vertices do not belong to other clauses). The edges of the clause are ux, vx and $f_i x, f_i c$ for $i \in \{1, 2, \ldots, d-2\}$ (so if $d = 2$ the only new edges are ux and vx). Let C'_e denote these edges of the clause. So $C'_e \cap C'_f = \emptyset$ for each pair of distinct edges $e, f \in J$. Let $C_{uv} := C'_{uv} + cu + cv$. Note that C_e is not d-sharp, since $|A(C_e)| = 2$ and $2d = |C_e| > d|A(C_e)| + |B(C_e)| - (d+1) = 2d - 1$. However, it is easy to check that removing any edge makes C_e d-sharp. We set $N = |E| - M$.

Lemma 7. *Let $Y \subseteq E$. Then $E - Y$ is d-sharp if and only if $C_e \cap Y \neq \emptyset$ for every $e \in J$.*

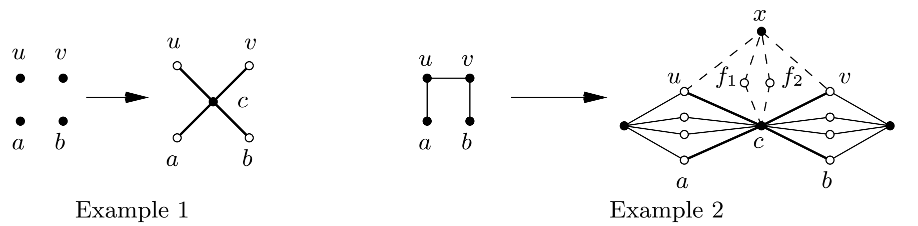

Fig. 3. Two examples of the reduction for $d = 4$. Empty circles are the vertices of B, filled circles are the vertices of A of H. The dotted lines are the edges of C'_{uv}. The thick lines are the edges of H_0.

Proof. The only if direction follows from the fact that C_e is not d-sharp for $e \in J$. To see the other direction suppose, for a contradiction, that $Z \subseteq E - Y$ is a set with $|A(Z)| \geq 2$ and $|Z| > d|A(Z)| + |B(Z)| - (d+1)$. We may assume that, subject to these properties, $A(Z) \cup B(Z)$ is minimal.

Minimality implies that if $|A(Z)| \geq 3$ then each vertex $w \in A(Z)$ is incident to at least $d+1$ edges of Z. Thus, since every vertex of $A - c$ has degree d in H, we must have $|A(Z)| = 2$. Minimality also implies that each vertex $f \in B(Z)$ is incident to at least two edges of Z. If $c \in A(Z)$ then $Z \subseteq C_e$, for some $e \in J$, since each vertex $f \in B(Z)$ has at least two edges from Z. But then $Y \cap C_e \neq \emptyset$ implies that Z is d-sharp. On the other hand if $c \notin A(Z)$ and $|A(Z)| = 2$ then $|Z| \leq 2$, and hence Z is d-sharp. This contradicts the choice of Z. Thus $E - Y$ is d-sharp. □

Lemma 8. *Let $Y \subseteq E$ and suppose that $E - Y$ is d-sharp. Then there is a set $Y' \subseteq E$ with $|Y'| \leq |Y|$ for which $E - Y'$ is d-sharp and $Y' \subseteq \{cu : u \in B\}$.*

Proof. Since $E - Y$ is d-sharp and C_e is not d-sharp we must have $C_e \cap Y \neq \emptyset$ for each $e \in J$. We obtain Y' by modifying Y with the following operations. If $|C_{uv} \cap Y| \geq 2$ for some $uv \in J$ then we replace $C_{uv} \cap Y$ by $\{cu, cv\}$. If $C_{uv} \cap Y = \{f\}$ and $f \notin \{cu, cv\}$ for some $uv \in J$ then we replace f by cu in Y. The new set Y' satisfies $|Y'| \leq |Y|$, and, by Lemma 7, $E - Y'$ is also d-sharp. □

We claim that H has a d-sharp edge set F with $|F| \geq N$ if and only if D has a vertex cover of size at most M. First suppose $F \subseteq E$ is d-sharp with $|F| \geq N$. Now $E - Y$ is d-sharp for $Y := E - F$, and hence, by Lemma 8, there is a set $Y' \subseteq E$ with $|Y'| \leq |Y| \leq M$ for which $E - Y'$ is d-sharp and $Y' \subseteq \{cu : u \in B\}$. Since $E - Y'$ is d-sharp, Lemma 7 implies that $X = \{u \in V : cu \in Y'\}$ is a vertex cover of D of size at most M.

Conversely, suppose that X is a vertex cover of D of size at most M. Let $Y = \{cu : u \in X\}$. Since X intersects every edge of D, we have $Y \cap C_{uv} \neq \emptyset$ for every $e \in J$. Thus, by Lemma 7, $F := E - Y$ is d-sharp, and $|F| \geq |E| - M = N$. Since our reduction is polynomial, this equivalence completes the proof of the theorem. □

Note that finding a maximum size d-sharp edge set is easy for $d = 1$, since an edge set F is 1-sharp if and only if each vertex of B is incident to at most one edge of F.

References

1. A. Berg and T. Jordán, A proof of Connelly's conjecture on 3-connected circuits of the rigidity matroid, *J. Combinatorial Theory, Ser. B.* Vol. 88, 77–97, 2003..
2. R. Connelly, On generic global rigidity, Applied geometry and discrete mathematics, 147–155, DIMACS Ser. Discrete Math. Theoret. Comput. Sci., 4, Amer. Math. Soc., Providence, RI, 1991.
3. A. Frank and A. Gyárfás, How to orient the edges of a graph, Combinatorics, (Keszthely), Coll. Math. Soc. J. Bolyai 18, 353–364, North-Holland, 1976.
4. H.N. Gabow and H.H. Westermann, Forests, frames and games: Algorithms for matroid sums and applications, *Algorithmica* 7, 465–497 (1992).
5. J. Graver, B. Servatius, and H. Servatius, Combinatorial Rigidity, AMS Graduate Studies in Mathematics Vol. 2, 1993.
6. B. Hendrickson, Conditions for unique graph realizations, *SIAM J. Comput.* **21** (1992), no. 1, 65-84.
7. B. Hendrickson and D. Jacobs, An algorithm for two-dimensional rigidity percolation: the pebble game, *J. Computational Physics* 137, 346-365 (1997).
8. J.E. Hopcroft and R.E. Tarjan, Dividing a graph into triconnected components, *SIAM J. Comput.* 2 (1973), 135–158.
9. H. Imai, On combinatorial structures of line drawings of polyhedra, *Discrete Appl. Math.* 10, 79 (1985).
10. B. Jackson and T. Jordán, Connected rigidity matroids and unique realizations of graphs, EGRES Technical Report 2002-12 (`www.cs.elte.hu/egres/`), submitted to J. Combin. Theory Ser. B.
11. D.J. Jacobs and M.F. Thorpe, Generic rigdity percolation: the pebble game, *Phys. Rev. Lett.* 75, 4051 (1995).
12. G. Laman, On graphs and rigidity of plane skeletal structures, *J. Engineering Math.* 4 (1970), 331–340.
13. L. Lovász and Y. Yemini, On generic rigidity in the plane, *SIAM J. Algebraic Discrete Methods* 3 (1982), no. 1, 91–98.
14. K. Sugihara, On some problems in the design of plane skeletal structures, SIAM J. Algebraic Discrete Methods 4 (1983), no. 3, 355–362.
15. K. Sugihara, Machine interpretation of line drawings, MIT Press, 1986.
16. W. Whiteley, Parallel redrawing of configurations in 3-space, preprint, Department of Mathematics and Statistics, York University, North York, Ontario, 1987.
17. W. Whiteley, A matroid on hypergraphs with applications in scene analysis and geometry, *Discrete Comput. Geometry* 4 (1989) 75–95.
18. W. Whiteley, Some matroids from discrete applied geometry. Matroid theory (Seattle, WA, 1995), 171–311, Contemp. Math., 197, Amer. Math. Soc., Providence, RI, 1996.