

ON AFFINE RIGIDITY

Steven J. Gortler,^{*} Craig Gotsman,[†] Ligang Liu,[‡] and Dylan P. Thurston,[§]

ABSTRACT. We define the notion of affine rigidity of a hypergraph and prove a variety of fundamental results for this notion. First, we show that affine rigidity can be determined by the rank of a specific matrix which implies that affine rigidity is a generic property of the hypergraph. Then we prove that if a graph is $(d + 1)$ -vertex-connected, then it must be “generically neighborhood affinely rigid” in d -dimensional space. This implies that if a graph is $(d + 1)$ -vertex-connected then any generic framework of its squared graph must be universally rigid. Our results, and affine rigidity more generally, have natural applications in point registration and localization, as well as connections to manifold learning.

1 Introduction

Suppose one has a number of overlapping “scans” of a set of points in some space, and that the corresponding points shared between scans have been identified. One naturally may want to register these scans and merge them together into a single configuration [26, *inter alia*]. Such a merging problem is called a *realization problem*. The study of the uniqueness of the solutions to such realization problems is known as *rigidity*.

We model the combinatorics of this problem using a hypergraph Θ , with vertices representing the points, and hyperedges representing the sets of points in each scan. The geometry of the problem is modeled with a configuration p , associating each vertex with a point in space.

One natural setting is the Euclidean setting, where the scans are known to be related by Euclidean transforms. In this case it is sufficient to study just the case of a graph, where we think of each edge as its own scan with only 2 points. Unfortunately, many of the Euclidean problems are NP-HARD [21]. In this paper, we study what happens when one relaxes the problem to the affine setting, that is, one assumes that the scans are known to be related by affine transforms. Under this relaxation, much of the analysis reduces to linear algebra, and uniqueness questions reduce to rank calculations. We prove a variety of fundamental results about this type of rigidity and also place it in the context of other rigidity classes such as global rigidity and universal rigidity.

^{*}Harvard University, sjg@cs.harvard.edu

[†]Technion, gotsman@cs.technion.ac.il, SJG and CG were partially supported by United States–Israel Binational Science Foundation grant 2006089.

[‡]University of Science and Technology of China, lgliu@ustc.edu.cn, LL was partially supported by the National Natural Science Foundation of China (61222206) and the One Hundred Talent Project of the Chinese Academy of Sciences.

[§]Indiana University, dpthurst@indiana.edu, DPT was supported by the Mathematical Sciences Research Institute and a Sloan Research Fellowship.

We also specifically investigate the case of hypergraphs Θ that arises by starting with an input graph Γ , and considering each one-ring (a vertex and its neighbors) in Γ as a hyperedge in Θ . We call such a hypergraph the neighborhood hypergraph of Γ . Such neighborhood hypergraphs naturally arise when studying molecules [16], when applying a divide and conquer approach to sensor network localization [23] and in machine learning [19].

1.1 Summary of Results

We start by describing how affine rigidity in \mathbb{R}^d is fully characterized by the kernel size of one of its associated “affinity matrices”. (This result was first shown by Zha and Zhang [30].) We show how this implies a number of interesting corollaries including the fact that affine rigidity is a generic property. That is, given a hypergraph Θ and dimension d , either all generic embeddings of Θ are affinely rigid in \mathbb{R}^d or all generic embeddings are affinely flexible in \mathbb{R}^d . The specific geometric positions of the vertices are irrelevant to this property, as long as they are in sufficiently general position. Thus we call such a hypergraph *generically affinely rigid* in \mathbb{R}^d .

Next we relate affine rigidity in \mathbb{R}^d to the related notion of universal Euclidean rigidity. A framework is universally Euclideanly rigid if it is rigid (in the Euclidean sense) in *any* dimension. In this context, we prove that affine rigidity in \mathbb{R}^d implies universal Euclidean rigidity.

We then prove the following sufficiency result: if a graph Γ is $d+1$ (vertex) connected, then its neighborhood hypergraph is generically affinely rigid in \mathbb{R}^d ; alternatively, we say that the graph Γ itself is generically neighborhood affinely rigid in \mathbb{R}^d . In particular we will show that almost every *non-symmetric equilibrium stress matrix* for any generic embedding of Γ in \mathbb{R}^d will have co-rank $d+1$ (i.e., rank $v-d-1$). Putting these two results together, we show that if a graph is $d+1$ connected, then any generic embedding of its square graph in \mathbb{R}^d is universally rigid. This result is interesting, as very few families of graphs have been proven to be generically universally rigid.

We give examples showing that many of the implications proved in this paper do not reverse. Finally we discuss some of the motivating applications.

The main properties of frameworks of graphs and their implications proven in this paper are summarized below.

GGR \rightarrow DP1C \rightarrow GNSESM \rightarrow GNAR \rightarrow GNUR \rightarrow GNGR	
Property	Graph ...
GGR	is generically globally rigid in \mathbb{R}^d
DP1C	is $d+1$ connected ([13])
GNSESM	generically has non-symmetric equilibrium stress matrix of rank $v-d-1$ (Proposition 5.8)
GNAR	is generically neighborhood affinely rigid in \mathbb{R}^d (Proposition 5.9)
GNUR	is generically neighborhood universally rigid in \mathbb{R}^d (Corollary 4.2)
GNGR	is generically neighborhood globally rigid in \mathbb{R}^d (by definition)




Figure 1: Left: A hypergraph with 6 vertices and 4 hyperedges. Each hyperedge is represented by a dotted ellipse enclosing a set of vertices. The hyperedges are: red $\{1, 2, 6\}$, green $\{2, 3, 5\}$, purple $\{5, 6\}$, blue $\{4, 5\}$. Right: The body graph of the hypergraph shown in the left.

2 A rigidity zoo

In this paper we will consider several different rigidity theories. They all fit in to a unifying framework, which we now explain.

Most generally, rigidity questions (of any type) ask if all of the geometric information about a set of points is determined by information from small subsets. In the usual Euclidean rigidity problem, we measure the distances between pairs of points. However, in other cases it is not enough to consider pairs of points for the small subsets; as such, we need to consider hypergraphs rather than just graphs.

See Figure 1 for an example.

Definition 2.1. A *hypergraph* Θ is a set of v vertices $\mathcal{V}(\Theta)$ and h hyperedges $\mathcal{E}(\Theta)$, where $\mathcal{E}(\Theta)$ is a set of subsets of $\mathcal{V}(\Theta)$. We will typically write just \mathcal{V} or \mathcal{E} , dropping the hypergraph Θ from the notation.

There are natural ways to pass from a hypergraph back and forth to a graph.

Definition 2.2. Given a hypergraph Θ , define its *body graph* $B(\Theta)$ as follows. For each vertex in Θ , we have a vertex in $B(\Theta)$. For each hyperedge h in Θ and each pair of vertices in h we have an edge in $B(\Theta)$.

Definition 2.3. Given a graph Γ , define its *neighborhood hypergraph*, written as $N(\Gamma)$ as follows. For each vertex in Γ , we have an associated vertex in $N(\Gamma)$. For each vertex in Γ we have a hyperedge in $N(\Gamma)$ consisting of that vertex and its neighbors in Γ .

Definition 2.4. Given a graph Γ , its *squared graph* Γ^2 is obtained by adding to Γ an edge between two vertices i and j if i and j share some neighbor vertex k .

Lemma 2.5. For any graph Γ , $B(N(\Gamma)) = \Gamma^2$.

Proof. Immediate from the definitions. (See Figure 2 for an example.) □




Figure 2: Left: A graph with 6 vertices and 7 edges. Middle: Its neighborhood hypergraph with 6 vertices and 6 hyperedges: red $\{1, 2, 6\}$, purple $\{1, 2, 3, 6\}$, light blue $\{2, 3, 5\}$, dark blue $\{4, 5\}$, orange $\{3, 4, 5, 6\}$, green $\{1, 2, 5, 6\}$. Right: The body graph of the hypergraph in the middle. It is also the squared graph of the graph in the left.

Definition 2.6. A k -hypergraph Θ is a hypergraph where each hyperedge has exactly k vertices. For any $k \in \mathbb{N}$ and hypergraph Θ , let $B_k(\Theta)$ be the k -hypergraph whose hyperedges are all the subsets S of size k of vertices that are contained together in at least one hyperedge of Θ :

$$\mathcal{E}(B_k(\Theta)) = \{ S \mid \exists h \in \mathcal{E}(\Theta), S \subset h, |S| = k \}.$$

For a vertex set S , the *complete k -hypergraph on S* , written $K_k(S)$, is the k -hypergraph whose hyperedges are all $\binom{|S|}{k}$ subsets of S of size k .

For instance, a 2-hypergraph is a graph, and $B_2(\Theta)$ is just the body graph $B(\Theta)$.

Definition 2.7. A *configuration* of the vertices $\mathcal{V}(\Theta)$ of a hypergraph in \mathbb{R}^d is a mapping p from $\mathcal{V}(\Theta)$ to \mathbb{R}^d . Let $C^d(\mathcal{V})$ be the space of configurations in \mathbb{R}^d . For $p \in C^d(\mathcal{V})$ and $u \in \mathcal{V}(\Theta)$, we write $p(u) \in \mathbb{R}^d$ for the image of u under p .

A *framework* $\rho = (p, \Theta)$ of a hypergraph is the pair of a hypergraph and a configuration of its vertices. $C^d(\Theta)$ is the space of frameworks (p, Θ) with hypergraph Θ and $p \in C^d(\mathcal{V}(\Theta))$. We may also write $\rho(u)$ for $p(u)$ where $\rho = (p, \Theta)$ is a framework of the configuration p .

A framework of a hypergraph has also been called a *body-and-joint framework* [29] and a *body-and-multipin framework* [16].

Definition 2.8. Let M be a monoid acting on \mathbb{R}^d , such as $\text{Eucl}(d)$, the Euclidean isometries of \mathbb{R}^d . (We study monoids instead of groups since we don't want to restrict ourselves to always having inverses. In particular in the case of affine rigidity, we wish to allow singular affine transforms as well). The framework $\rho \in C^d(\Theta)$ is *M -preequivalent* to the framework $\sigma \in C^d(\Theta)$ if for each hyperedge $h \in \mathcal{E}(\Theta)$, the positions of the vertices in ρ can be mapped to their positions in σ by an element of M depending only on h . That is, for each $h \in \mathcal{E}$, there is a $g_h \in M$ so that for each $u \in h$, we have $g_h(\rho(u)) = \sigma(u)$.

The configuration $p \in C^d(\mathcal{V})$ is *M -precongruent* to the configuration $q \in C^d(\mathcal{V})$ if the positions of all the vertices in p can be mapped to their positions in q by a single element

$g \in M$ (not depending on h). When the configuration p is M -precongruent to q , we also say that the framework (p, Θ) is M -precongruent to (q, Θ) .

A framework $\rho \in C^d(\Theta)$ is *globally M -rigid* if for any other framework $\sigma \in C^d(\Theta)$ to which ρ is M -preequivalent, we also have that ρ is M -precongruent to σ . Otherwise we say that ρ is globally M -flexible in \mathbb{R}^d .

Similarly, a framework $\rho \in C^d(\Theta)$ is *locally M -rigid* in \mathbb{R}^d if there is a small neighborhood U of ρ in $C^d(\Theta)$ so that for any $\sigma \in U$ to which ρ is M -preequivalent, we also have that ρ is M -precongruent to σ . Otherwise we say that ρ is locally M -flexible in \mathbb{R}^d .

Remark 2.9. When there are non-invertible elements of M , then neither M -preequivalence nor M -precongruence is a symmetric relation. When M is a group, then M -preequivalence is a symmetric relation and can be called M -equivalence, and likewise M -precongruence can be called M -congruence.

A related notion of group based rigidity has been explored in the computer aided design literature [22].

In this paper, we are mainly concerned with the cases when M is either $\text{Eucl}(d)$ or $\text{Aff}(d)$, the set of all (including singular) affine linear maps of \mathbb{R}^d , in which case we speak about *Euclidean* or *affine* rigidity, respectively. But there are other interesting possibilities, like projective transformations. Another interesting case is when M consists of the dilations and translations of \mathbb{R}^d (with no rotations); this gives the theory of parallel-line redrawings [29].

In this terminology, Euclidean rigidity is the default: if M is not specified, it is the Euclidean group. In much of the rigidity literature, local rigidity is the default, and the qualifier “local” is dropped. However, in this paper this distinction is important and we will write “local” or “global” when the distinction is meaningful.

Lemma 2.10. *A framework (p, Θ) is locally (resp. globally) Euclideanly rigid iff the body framework $(p, B(\Theta))$ is locally (resp. globally) Euclideanly rigid.*

Proof. This easily follows from the fact that, for each hyperedge $h \in \mathcal{E}(\Theta)$, the complete graph on $|h|$ vertices is globally rigid. \square

Thus we only need to consider Euclidean rigidity for frameworks of graphs, not hypergraphs.

In the next section (Corollary 3.5) we will see that a framework is locally affinely rigid iff it is globally affinely rigid. Thus we can drop the local/global distinction for affine rigidity.

In the following definition, for $d < d'$, we view $C^d(\mathcal{V})$ as contained in $C^{d'}(\mathcal{V})$ by the inclusion of \mathbb{R}^d as the first d coordinates of $\mathbb{R}^{d'}$.

Definition 2.11. Let M be a family of monoids M_d acting on \mathbb{R}^d , so that for $d < d'$, M_d is the submonoid of $M_{d'}$ that fixes \mathbb{R}^d as a subset of $\mathbb{R}^{d'}$. A framework $\rho \in C^d(\Theta)$ is *universally locally* (resp. *globally*) M -rigid if it is locally (resp. globally) M -rigid as a framework in $C^{d'}(\Theta)$ for all $d' \geq d$.

Note that universal rigidity of any sort implies rigidity of the same sort.

Lemma 2.12. *A framework $p \in C^d(\Theta)$ is universally globally Euclideanly rigid iff it is universally locally Euclideanly rigid.*

Proof. For any two equivalent frameworks ρ in $C^d(\Theta)$ and ρ' in $C^{d'}(\Theta)$, Bezdek and Connelly [5] show how to construct an explicit flex between ρ and ρ' in $C^{d+d'}(\Theta)$. Thus, if ρ is a d -dimensional framework with a equivalent but non-congruent framework in d' dimensions, then their constructed flex shows that ρ is not locally rigid in $\mathbb{R}^{d+d'}$. \square

Thus we can also drop the local/global distinction in the case of universal Euclidean rigidity.

Definition 2.13. A framework (p, Γ) of the graph Γ is *neighborhood rigid* (of any of the sorts above) if the corresponding framework $(p, N(\Gamma))$ of the neighborhood hypergraph is rigid (of the same sort).

For instance, Lemmas 2.10 and 2.5 tell us that neighborhood Euclidean rigidity of (ρ, Γ) is equivalent to the Euclidean rigidity of (ρ, Γ^2) .

Remark 2.14. Related to universal Euclidean rigidity is the notion of *dimensional rigidity* [1]. A framework (locally rigid or not) in \mathbb{R}^d is called dimensionally rigid if there is no (Euclidean) equivalent framework with an affine span of dimension strictly greater than d .

Another related notion is that of *d-realizability* [4]. A graph is *d-realizable* if any framework of the graph, in any dimension, has a (Euclidean) equivalent framework with an affine span of dimension d or less.

Presumably one could extend these notions to arbitrary monoids as well but we will not pursue these in this paper.

Definition 2.15. A configuration p in $C^d(\mathcal{V})$ is *generic* if the coordinates do not satisfy any non-zero algebraic equation with rational coefficients. A framework is generic if its configuration is generic.

A property is *generic* in \mathbb{R}^d if, for every (hyper)graph, either all generic frameworks in $C^d(\Theta)$ have the property or none do. For instance, local and global Euclidean rigidity in \mathbb{R}^d are both generic properties of graphs and therefore for hypergraphs as well [3, 11]. For any property P (generic or not) of frameworks, a (hyper)graph Θ is *generically P* in \mathbb{R}^d if every generic framework in $C^d(\Theta)$ has P . (For a non-generic property like universal Euclidean rigidity, there are (hyper)graphs that are neither generically P or generically not P .)

Thus, for any framework, we may talk about

(generic/ \emptyset) (universal/ \emptyset) (local/global) (Euclidean/affine) rigidity.

where by \emptyset we mean that this term has been dropped.

3 Affine Rigidity in \mathbb{R}^d

We now move on the main focus of this work, affine rigidity, as defined in the previous section. Though the definitions start from a different point of view, this notion of affine rigidity is, in fact, identical to the one defined by Zha and Zhang [30] and the concept is also informally mentioned by Brand [6]. Additionally, Theorem 1 below is essentially equivalent to [30, Theorem 5.2].

Our contribution here, described by the corollaries, is showing how affine rigidity fits in to the general scheme of rigidity problems.

Lemma 3.1. *Any framework of a complete $(d + 2)$ -hypergraph in \mathbb{R}^d is affinely rigid.*

Proof. Let $q \in C^d(\mathcal{V})$ be a configuration with such that (p, Θ) is affinely preequivalent in \mathbb{R}^d to (q, Θ) .

Let $c \leq d$ be the dimension of the affine span, S , of the configuration p . Select $c + 1$ vertices whose affine span in p is S . Let A_0 be an affine transform that maps these vertices from their positions in p to their positions in q . The action of A_0 on the space S (and thus all of the vertices p) is uniquely determined by these selected vertices.

For any vertex v , there must be a hyperedge h_v that includes v and the selected vertices. Let A_{h_v} be an affine transform that maps these vertices from their positions in p to q (which must exist by affine preequivalence). The action of A_{h_v} on the space S (and thus all of the vertices p) is uniquely determined by the selected vertices, and thus must agree with that of A_0 . Thus for all v , we see that their positions in q are obtained from the positions in p through A_0 . Thus p is affinely precongruent to q . \square

Proposition 3.2. *A framework (p, Θ) in general position is affinely locally (resp. globally) rigid iff the associated framework $(p, B_{d+2}(\Theta))$ is affinely locally (resp. globally) rigid.*

(Compare Lemma 2.10.)

Proof. First consider a hyperedge with less than $d + 2$ vertices in general position. Using a d -dimensional affine transform, we can move these vertices to any other configuration in \mathbb{R}^d . Therefore this hyperedge does not affect affine preequivalence and may be dropped without affecting affine rigidity. Next consider a hyperedge with k vertices, with $k > d + 2$. By Lemma 3.1, one can replace this hyperedge with $\binom{k}{d+2}$ hyperedges corresponding to all subsets of $d + 2$ vertices. The framework of the new hypergraph will be affinely rigid in \mathbb{R}^d iff the original one is. \square

Definition 3.3. An *affinity matrix* of a framework (p, Θ) in $C^d(\Theta)$ is a matrix with v columns such that each row encodes some affine relation between the coordinates of the vertices in a hyperedge of (p, Θ) as a homogeneous linear equation in the following sense. The only non-zero entries in a row correspond to vertices in some hyperedge, the sum of the entries in a row is 0, and each of the coordinates of p , thought of as a vector of length v , is in the kernel of the matrix.

An affinity matrix is *strong* if it encodes all of the affinely independent relations in every hyperedge of (p, Θ) .

Lemma 3.4. *If the framework (p, Θ) is affinely preequivalent to (q, Θ) then the coordinates of q are in the kernel of any affinity matrix for (p, Θ) . Additionally, the converse is true if the affinity matrix is strong.*

Proof. Clear from the definitions. □

The kernel of an affinity matrix of a framework $(p, \Theta) \in C^d(\Theta)$ always contains the subspace of \mathbb{R}^v spanned by the coordinates of p along each axis and the vector $\vec{1}$ of all 1's. This corresponds to the fact that any p is preequivalent to any of its affine images. If p is a proper d -dimensional configuration (with full d -dimensional affine span), these vectors are independent and span a $(d + 1)$ -dimensional space. In particular, a generic framework of a hypergraph with at least $d + 1$ vertices in \mathbb{R}^d is proper, so for such frameworks the corank of any of its affinity matrices must be no less than $d + 1$.

The rank of strong affinity matrices fully characterize affine rigidity.

Theorem 1. *Let Θ be a hypergraph with at least $d + 1$ vertices. Let (p, Θ) be any proper, d -dimensional framework and let M be any strong affinity matrix for (p, Θ) . Then (p, Θ) is affinely rigid in \mathbb{R}^d iff $\dim(\ker(M)) = d + 1$.*

Proof. By Lemma 3.4, for any other configuration q in $C^d(\mathcal{V})$ such that (p, Θ) is affinely preequivalent to (q, Θ) , the coordinates of q must be in the kernel of M . When $\dim(\ker(M)) = d + 1$, the kernel of M contains only one-dimensional projections of p and the all-ones vector. Thus when (p, Θ) is affinely preequivalent to (q, Θ) , we have that (p, Θ) must in fact be affinely precongruent to (q, Θ) .

Conversely, if the corank is higher, then the kernel must contain an “extra” vector that is not a one-dimensional projection of p . Adding any amount of this vector to one of the coordinates of p must, by Lemma 3.4, produce a q such that (p, Θ) is affinely preequivalent to (q, Θ) but not precongruent to it. □

It is easy now to prove the following corollaries.

Corollary 3.5. *If (p, Θ) is affinely globally flexible in \mathbb{R}^d then it is affinely locally flexible.*

Proof. From the proof of Theorem 1, when (p, Θ) is affinely globally flexible in \mathbb{R}^d there is an extra vector δ in the kernel of a strong affinity matrix, and we can add any multiple of δ to one of the coordinates of p to get a framework to which (p, Θ) is affinely preequivalent but not precongruent. □

Remark 3.6. In fact, if (p, Θ) is affinely preequivalent to (q, Θ) , there is a continuous path of frameworks in $C^d(\Theta)$ to which (p, Θ) is affinely preequivalent, namely $((1 - t)p + tq, \Theta)$.

Corollary 3.7. *A framework $(p, \Theta) \in C^d(\Theta)$ is affinely rigid in \mathbb{R}^d iff it is affinely rigid when considered as a (degenerate) framework in $\mathbb{R}^{d'}$ for $d' \geq d$.*

Proof. Follows from the proof of Theorem 1. □

Thus there is no distinct notion of “universal” affine rigidity.

Corollary 3.8. *Affine rigidity in \mathbb{R}^d is a generic property of a hypergraph.*

Proof. The condition that M is an affinity matrix for (p, Θ) is linear in the entries in M . The corollary then follows from Proposition A.1. □

Though we will not pursue the details here, one can use the concept of an affinity matrix to derive an efficient randomized discrete algorithm for testing generic affine rigidity of a hypergraph in \mathbb{R}^d . To do this, one needs to use integers of sufficiently many bits, and do the arithmetic modulo a suitably large prime. The details parallel those in the global rigidity case [11, Section 5].

Remark 3.9. There is also an strong connection between affine rigidity and a problem from polyhedral scene analysis [29]. This is most easily explained in two dimensions. Given a framework in \mathbb{R}^2 , one can interpret each hyperedge as a planar polygon drawn in \mathbb{R}^2 (the vertex order is not relevant here). We say that the framework is *sharp* if each vertex can be given a third coordinate, such that, in the resulting three dimensional drawing, each polygon remains planar, and the faces do not all lie in a single plane. This idea is easily generalized to arbitrary dimension.

As shown in [29, Proposition 2.1], a framework is sharp iff the rank of its strong affinity matrix is not maximal. Thus this notion of sharpness corresponds exactly to affine flexibility.

More deeply, due the combinatorial characterization of sharpness given by [29, Theorem 4.2], generic affine rigidity can be tested by an efficient deterministic algorithm.

4 Universal Euclidean Rigidity

We now turn to universal Euclidean rigidity. To begin, we need the following technical definition:

Definition 4.1. We say that the edge directions of a graph framework $(p, \Gamma) \in C^d(\Gamma)$ are *on a conic at infinity* if there exists a symmetric d -by- d matrix Q such that for all edges (u, v) of Γ , we have

$$[p(u) - p(v)]^t Q [p(u) - p(v)] = 0.$$

The edge directions of (p, Γ) are on a conic at infinity iff there is a continuous family of d -dimensional non-Euclidean affine transforms that preserve all of the edge lengths [9]. This is a very degenerate situation which is very easy to rule out. For example, if in a hypergraph framework (p, Θ) some hyperedge in Θ has vertices whose positions in p affinely span \mathbb{R}^d , then the edge directions $(p, B(\Theta))$ cannot be on a conic at infinity.

Theorem 2. *If a framework (p, Θ) of a hypergraph Θ with $p \in C^d(\mathcal{V})$ is affinely rigid in \mathbb{R}^d , and the edge directions of $(p, B(\Theta))$ are not on a conic at infinity, then (p, Θ) is universally Euclidean rigid.*




Figure 3: The framework in \mathbb{R}^2 of the hypergraph on the left is not affinely rigid as each hyperedge (shown as a dashed ellipse) has only 3 vertices. But this framework is universally Euclidean rigid, as its body graph (right) is a Cauchy polygon. (A Cauchy (bar) polygon on v vertices is a planar framework where the vertices $p(1), \dots, p(v)$, in order, form a strictly convex polygon in the plane, and the edge set consists of the edges $\{i, i + 1\}$, $i = 1, 2, \dots, v$, and $\beta, i + 2\}$, $i = 1, \dots, v - 2$ (indices modulo v). A Cauchy polygon is universally rigid [8]).

Proof. Let $q \in C^{d'}(\mathcal{V})$ be a configuration with $d' > d$ such that (p, Θ) is Euclidean equivalent in $\mathbb{R}^{d'}$ to (q, Θ) . Then (p, Θ) is affinely preequivalent in $\mathbb{R}^{d'}$ to (q, Θ) . Since (p, Θ) is affinely rigid in \mathbb{R}^d , from Corollary 3.7, we have that p is affinely precongruent to q in $\mathbb{R}^{d'}$ and the affine span of q must be of dimension no larger than d . Let $R(q)$ be a rotation of q down to \mathbb{R}^d . Then p is affine precongruent in \mathbb{R}^d to $R(q)$ and (p, Θ) is Euclidean equivalent in \mathbb{R}^d to $(R(q), \Theta)$.

Let A be an affine transform such that $A(p) = R(q)$ (which must exist due to affine precongruence). By Euclidean equivalence, all of edge lengths agree between $(p, B(\Theta))$ and $(A(p), B(\Theta))$. If A is not Euclidean, then this means that the edge directions of $(p, B(\Theta))$ are on a conic at infinity, which contradicts our assumption. Thus A is Euclidean making p and $R(q)$ Euclidean congruent in \mathbb{R}^d . Likewise p must be Euclidean congruent to q in $\mathbb{R}^{d'}$, and we can conclude that (p, Θ) is universally rigid. \square

Corollary 4.2. *Let Θ be a hypergraph with at least $d + 2$ vertices. If a generic framework (p, Θ) of a hypergraph Θ with $p \in C^d(\mathcal{V})$ is affinely rigid in \mathbb{R}^d then (p, Θ) is universally rigid.*

Proof. As in the proof of Proposition 3.2, any generic framework of a hypergraph Θ with at least $d + 2$ vertices that is affinely rigid in \mathbb{R}^d , must have at least one hyperedge h with at least $d + 2$ vertices, and these vertices in p have a d -dimensional affine span. Thus $(p, B(\Theta))$ must include a general position framework of a $(d + 1)$ -simplex and thus cannot have edge directions at a conic at infinity. Then Theorem 2 applies. \square

There can be frameworks that are universally rigid but not affinely rigid in \mathbb{R}^d . (See Figure 3.)

Corollary 4.2 can be generalized beyond the Euclidean case to apply to much larger set of groups and monoids.¹

¹Thanks to Louis Theran for suggesting we look at this generality.

Theorem 3. *Let M be a family of monoids M_d (as in Definition 2.11) with each M_d a submonoid of $\text{Aff}(d)$. Let $(p, \Theta) \in C^d(\Theta)$ be a framework with some hyperedge h_0 in Θ whose vertex positions in p affinely span \mathbb{R}^d . If (p, Θ) is affinely rigid in \mathbb{R}^d then (p, Θ) is universally M -rigid.*

Proof. Let $q \in C^{d'}(\mathcal{V})$ be a configuration with $d' > d$ such that (p, Θ) is M -preequivalent in $\mathbb{R}^{d'}$ to (q, Θ) . Then (p, Θ) is affinely preequivalent in $\mathbb{R}^{d'}$ to (q, Θ) . Since (q, Θ) is affinely rigid in \mathbb{R}^d , from Corollary 3.7, we have that p is affinely precongruent to q in $\mathbb{R}^{d'}$; there is an $A \in \text{Aff}(d')$ such that $A(p) = q$.

By the assumption of M -preequivalence, for each hyperedge h there is an element $g_h \in M_{d'}$ which maps the vertices of h from their positions in p to their positions in q . Since $M_{d'}$ is a subgroup of $\text{Aff}(d')$, and the specific hyperedge h_0 has $d + 1$ vertices in general position in the configuration p , the action of g_{h_0} on \mathbb{R}^d is fully determined by these vertices and must agree with the action of A on \mathbb{R}^d . Thus $g_{h_0}(p) = A(p) = q$, making p M -precongruent to q , and making (p, Θ) universally M -rigid. \square

5 Neighborhood affine rigidity

In this section we prove the following theorem about the generic neighborhood affine rigidity of a graph.

Theorem 4. *Let Γ be a graph. If Γ is $(d + 1)$ -vertex-connected, then Γ is generically neighborhood affinely rigid in \mathbb{R}^d .*

The strategy to prove this theorem is as follows. First we show, using a “rubber band” construction [8, 18, 27], that a sufficiently connected graph must have a framework with certain nice geometric properties. Moreover, these geometric properties are stable under generic perturbations of the configuration. Then we show that any such framework must have a “non-symmetric equilibrium stress matrix” of appropriate high rank. Since the perturbed framework is generic, then any generic framework must have such a matrix. This matrix serves as a certificate of neighborhood affine rigidity.

Definition 5.1. An *equilibrium stress matrix* of a framework (p, Γ) of a graph in $C^d(\Gamma)$ is a matrix Ω indexed by $\mathcal{V} \times \mathcal{V}$ so that

1. for all $u, w \in \mathcal{V}$, we have $\Omega(u, w) = \Omega(w, u)$;
2. for all $u, w \in \mathcal{V}$ with $u \neq w$ and $\{u, w\} \notin \mathcal{E}$, we have $\Omega(u, w) = 0$;
3. for all $u \in \mathcal{V}$, we have $\sum_{w \in \mathcal{V}} \Omega(u, w) = 0$; and
4. for all $u \in \mathcal{V}$, we have $\sum_{w \in \mathcal{V}} \Omega(u, w)p(w) = 0$.

A *non-symmetric equilibrium stress matrix* of a framework (p, Γ) is a matrix that satisfies properties (2)–(4) above.

Observe first that an equilibrium stress matrix (symmetric or not) Ω of (p, Γ) is an affinity matrix of $(p, N(\Gamma))$. From the properties of affinity matrices, the kernel of Ω always contains the subspace spanned by the coordinates of p along each axis and the vector $\vec{1}$ of all 1's.

Definition 5.2. We say that a framework of a graph in $C^d(\Gamma)$ has the *convex containment* property if

1. the configuration of each vertex along with its neighboring vertices has an affine span of dimension d , and
2. Almost every vertex in the framework is contained in the strict d -dimensional convex hull of its neighbors. There are may be up to $d + 1$, so-called, *exceptional vertices* which do not have this property.

Lemma 5.3. *Let Γ be a graph with at least $d + 1$ vertices. Suppose Γ is $(d + 1)$ -connected. Then there exists a generic framework (q, Γ) in $C^d(\Gamma)$ with the convex containment property.*

Proof. Pick any $d + 1$ vertices to be exceptional. Constrain the exceptional vertices to fixed generic positions in \mathbb{R}^d (at the vertices of a simplex). Associate generic positive weights ω_{ij} with each (undirected) edge ij . Find the “rubber band” configuration consistent with the constrained vertices and these weights. Namely, find a framework (r, Γ) so that each non-exceptional vertex is the weighted linear average of its neighbors:

$$\sum_{j \in N(i)} \omega_{ij}(r(i) - r(j)) = 0,$$

where $N(i)$ are the neighbors of vertex i . This involves solving d systems of linear equations, one for each component of r . Note that the resulting configuration r may not be generic.

From [18], we know that if Γ is $(d + 1)$ -connected and the constraints on the exceptional vertices and the edge weights ω are generic, then no set of $d + 1$ vertices in r will be contained in a $(d - 1)$ -dimensional affine plane, giving us the first condition.

By construction, any non-exceptional vertex in (r, Γ) must be contained in the convex hull of its neighbors. Again, from [18], the convex containment must be strict.

Finally, we perturb each vertex in \mathbb{R}^d to obtain a generic configuration in $q \in C^d(\mathcal{V})$. By the first convex containment condition, the convex hull of the neighbors of a vertex has non-empty interior, so a sufficiently small perturbation will maintain both conditions. \square

Definition 5.4. Suppose that (q, Γ) has the convex containment property and Ω is a non-symmetric equilibrium stress matrix for (q, Γ) . We call a row of Ω *non-exceptional* if its corresponding vertex is in the strict d -dimensional convex hull of its neighbors.

Lemma 5.5. *Let Γ be a graph with at least $d + 1$ vertices. Suppose Γ is a $(d + 1)$ -connected graph, and we have a framework (q, Γ) in $C^d(\Gamma)$ with the convex containment property. Then there is a non-symmetric equilibrium stress matrix Ω of (q, Γ) , such that for every non-exceptional row i , we have the following property: If there is an edge connecting vertex i and vertex j , then Ω_{ij} is positive.*

Proof. All vertices have $d + 1$ or more neighbors. For each vertex i , we can therefore find “barycentric coordinates”: non-zero edge weights ω_{ij} on the adjoining edges so that

$$\sum_{j \in N(i)} \omega_{ij}(q(j) - q(i)) = 0.$$

If i is a non-exceptional vertex, due to the convex containment property we can choose the ω_{ij} to be positive. We then choose $\Omega_{ij} = \omega_{ij}$ for $i \neq j$ and $\Omega_{ii} = -\sum_j \omega_{ij}$. \square

Remark 5.6. This lemma is false if we require the stress matrix to be symmetric, because this prevents us from choosing ω_{ij} and ω_{ji} independently.

Lemma 5.7. *Let Γ be a graph with at least $d + 1$ vertices. Suppose Γ is $(d + 1)$ -connected, and we have a framework (q, Γ) in $C^d(\Gamma)$ with the convex containment property. Then there is a non-symmetric equilibrium stress matrix Ω of (q, Γ) with co-rank $d + 1$.*

Proof. From Lemma 5.5 we find for (q, Γ) a non-symmetric equilibrium stress matrix Ω with the desired positive entries. We now show that Ω has the stated rank.

First remove the $d + 1$ rows and columns associated with the exceptional vertices to create a smaller matrix Ω' . Due to the sign pattern assumed in Ω , as well as property (3) of any equilibrium stress matrix, Ω' must be weakly diagonally dominant.

Let us call a vertex *EN* if it has an exceptional neighbor and refer to its corresponding row in Ω' as *EN*. Any *EN* row must be strictly diagonally dominant (since at least one non-zero off-diagonal entry of Ω have been removed from this row).

Since all entries corresponding to edges are non-zero, the irreducible components of Ω' correspond to vertex subsets that remain connected after the exceptional vertices have been removed. (An irreducible square matrix is one that is not similar via a permutation to a block upper triangular matrix. Any square matrix has a unique irreducible decomposition).

Each irreducible component of Ω' includes such an *EN* row, thus Ω' must be full rank. (See, e.g., [28, Theorem 1.21].)

Since Ω' has co-rank 0, the co-rank of Ω must be at most $d + 1$. It is no less since any equilibrium stress matrix must have a $(d + 1)$ -dimensional kernel spanned by the coordinates of q and the all-ones vector. \square

Proposition 5.8. *Let Γ be a graph with at least $d + 1$ vertices. Suppose Γ is $(d + 1)$ -connected, and p is generic in $C^d(\mathcal{V})$. Then there is a non-symmetric equilibrium stress matrix Ω of (p, Γ) with co-rank $d + 1$.*

Proof. From Lemma 5.3, there must exist a generic framework (q, Γ) in $C^d(\Gamma)$ that has the convex containment property. From Lemma 5.7, (q, Γ) must have a non-symmetric equilibrium stress matrix of co-rank $d + 1$. Thus from Proposition A.1, any generic framework (p, Γ) must have such a matrix as well. \square

See Figure 4 for an example showing that the converse of Proposition 5.8 is not true. Since the upper (and lower) vertex in the framework has 3 neighbors in affine general




Figure 4: This framework in \mathbb{R}^2 is not 3-connected but does have a non-symmetric stress matrix of high rank.




Figure 5: This framework in \mathbb{R}^2 does not have a non-symmetric equilibrium stress matrix of co-rank $d + 1 = 3$, but is (trivially) neighborhood affinely rigid.

position, its position can be written as an affine combination of these neighbors producing a non-zero, non-symmetric equilibrium stress matrix. Any non-zero, non-symmetric equilibrium stress matrix must have rank at least 1 and co-rank of no more than $3 = d + 1$. Thus, as an equilibrium stress matrix, it has co-rank of exactly $d + 1$. Meanwhile, this framework is not 3-connected.

Note that from the proof of Proposition A.1 it is clear that if Γ is $(d + 1)$ -connected, then *almost every* non-symmetric stress matrix for almost every (p, Γ) in $C^d(\Gamma)$ will have co-rank $d + 1$. Moreover, each row of such a non-symmetric stress matrix of p can be constructed independently from the other rows, and we still expect to find this minimal co-rank.

Proposition 5.9. *Let Γ be a graph with at least $d + 1$ vertices. Suppose (p, Γ) , a framework in $C^d(\Gamma)$, has a non-symmetric equilibrium stress matrix Ω that has co-rank $d + 1$. Then (p, Γ) is neighborhood affinely rigid in \mathbb{R}^d .*

Proof. Ω is a (not strong) affinity matrix of $(p, N(\Gamma))$ and so the proof follows that of the first direction of Theorem 1. \square

Proof of Theorem 4. The theorem now follows directly from Propositions 5.8 and 5.9. If Γ has less than $d + 1$ vertices and is $(d + 1)$ -connected, then it is a simplex and thus neighborhood affinely rigid for any configuration. \square

See Figure 5 for an example showing that the converse of Proposition 5.9 is not true. The framework is clearly neighborhood affinely rigid since the central vertex is adjacent to all of the other vertices. Meanwhile the outer 4 vertices have only one neighbor and hence must have all zeros in their corresponding rows of any non-symmetric equilibrium stress

matrix. Thus, this framework cannot have a non-symmetric equilibrium stress matrix of co-rank $d + 1 = 3$.

Remark 5.10. Generic global rigidity of a graph Γ in \mathbb{E}^d can be characterized either using the dimension of the kernel of a single symmetric stress matrix of a generic framework (p, Γ) or using the dimension of the *shared symmetric stress kernel* of a generic p : the intersection of the kernels of all stress matrices of p [11].

By contrast, the analogous statement is not true in the affine rigidity case. By “vertically concatenating” a sufficient number of non-symmetric equilibrium stress matrices of (p, Γ) , we can create a strong affinity matrix for $(p, N(\Gamma))$. The kernel of the vertical concatenation will be the shared non-symmetric stress kernel of (p, Γ) , and the dimension of this kernel characterizes affine rigidity. Since the converse of Proposition 5.9 is not true, we see that neighborhood affine rigidity cannot in general be characterized by the rank of one single (say, generic) non-symmetric equilibrium stress matrix for (p, Γ) .

Note that there is a different sufficiency condition for affine rigidity given by Zha and Zhang [30]. Their condition is complementary to our condition (neither strictly stronger or weaker), and (like trilateralization [10]) is greedy in nature. Their condition on generic frameworks of a hypergraph requires that for each pair of vertices s and t , one can find a sequence of hyperedges starting with some hyperedge containing s and ending with some hyperedge containing t , such that for each pair (i, j) of hyperedges in the sequence, h_i and h_j share at least $d + 1$ vertices. When translated to a neighborhood hypergraph $N(\Gamma)$, it states that one can walk between any two vertices along edges, such that for each pair (i, j) of vertices along the walk, the neighborhoods of these vertices share at least $d + 1$ vertices in Γ .

Figure 6 shows a graph which clearly fails Zha and Zhang’s condition, but is 3-connected, showing that their condition does not imply Theorem 4. It is not hard to construct examples in the opposite direction, as well.

We also have the following corollary of Theorem 4

Corollary 5.11. *Let Γ be a graph. If Γ is $(d + 1)$ -connected, then any generic framework of Γ^2 in \mathbb{R}^d is universally rigid.*

Proof. If Γ has at least $d + 2$ vertices, then we can directly apply Corollary 4.2. Any graph with fewer vertices that is $(d + 1)$ -connected must be a simplex and be universally rigid for all configurations. \square

Remark 5.12. This corollary can also be proven without reference to affine rigidity and Corollary 4.2. In particular, Proposition 5.8 guarantees a non symmetric maximal-rank stress Ω for (p, Γ) , and then $\Omega^t \Omega$ is a symmetric, positive semi-definite, maximal rank stress for (p, Γ^2) . Universal rigidity then follows by a theorem of Connelly [8]. (See also [12].)

A manuscript by Cheung and Whiteley [7] contains a variety of other results relating graph powers to global rigidity.

We wish to highlight this corollary since the only other known (to us) class of graphs that are universally rigid for all generic configurations in \mathbb{R}^d are graphs that can be realized




Figure 6: A drawing of a hexagonal lattice on the torus. (The vertices on the top edge should be identified with those on the bottom and similarly with the left and right, as indicated by the arrows.) This graph is 3-connected but its neighborhood hypergraph does not satisfy the sufficiency condition of Zha and Zhang [30], and its squared graph is not a 2-trilateralization graph.

greedily such as the d -trilateralization graphs (A d -trilateralization graph is one that can be obtained from a complete graph by successively adding vertices, each connected to at least $d + 1$ old vertices) and their generalizations (such as graphs formed by gluing together d -trilateralization graphs along $d + 1$ vertices).

See Figure 6 for an example of a framework whose square is universally rigid by this corollary but is not a trilateralization graph.

We also mention that a theorem of a related nature, showing a relationship between the connectivity of a graph and global rigidity in the squared graph, has been described by Anderson et al. [2].

6 Applications

6.1 Registration

There are many applications where one has multiple views of some underlying configuration, but it is not known how these views all fit together. We assume that these views share some points in common, and this correspondence is known. (Of course in practice, establishing such a correspondence could in itself be a very challenging problem.) For example, in computer vision, one may have multiple uncalibrated laser scans of overlapping parts of some three-dimensional object.

In our setting we model all of the points as vertices, and each of the views as a hyperedge. The geometry of the vertices in each hyperedge h is given up to some unknown

transform T_h from a relevant class. The goal in registration then is to *realize* the entire hypergraph up to the relevant congruence class.

Affine case: Suppose we wish to realize a framework (p, Θ) where we are given as input the geometry of each hyperedge h up to an affine transform A_h . Theorem 1 tells us that if (p, Θ) is affine rigid, then we can compute the realization just using linear algebra. In particular, we can use the data for each hyperedge to build its associated rows in a strong affinity matrix. Then we can solve for its kernel, giving us our answer p .

If our hypergraph Θ happens to be the neighborhood graph of an underlying graph Γ , then one could also construct a (smaller) non-symmetric equilibrium stress matrix Ω for (p, Γ) . This is not guaranteed to work; even when (p, Γ) is neighborhood affinely rigid in \mathbb{R}^d , the matrix Ω may have co-rank larger than $d + 1$. But Theorem 4 states that if Γ is $(d + 1)$ -connected, this method will indeed work for almost every p in \mathbb{R}^d (and, in fact, using almost any non-symmetric equilibrium stress matrix for (p, Γ)).

Euclidean case: The Euclidean framework registration problem is perhaps more natural and common.

When (p, Θ) is globally rigid in \mathbb{R}^d , this problem is well posed, but it is general hard to solve, as the graph case includes the graph realization problem which is strongly NP-HARD [21].

When (p, Θ) is, in fact, also universally rigid there is an efficient algorithm: we can solve the Euclidean registration problem using semi-definite programming. One simply sets up the program that looks for the Gram matrix of an embedding of the vertices in \mathbb{R}^v (a semi-definite constraint on a Gram matrix) subject to the length constraints (linear constraints on the Gram matrix) [17]. Due to universal rigidity, one does not need to explicitly enforce the (non-convex) constraint that the embedding have a d -dimensional affine span [25].

When (p, Θ) is, furthermore, affinely rigid then we can solve the Euclidean registration problem using linear algebra. We can simply reduce this problem to an affine registration problem above, and find p using the kernel vectors of an affinity matrix. This determines p up to some global affine transform. Moreover, for (p, Θ) that is generically globally rigid in \mathbb{R}^d , we can solve a second (least squares) linear system to remove the unwanted global affine transform, leaving just the unknown global Euclidean transform (see Appendix B). This approach is morally the same affine relaxation used in the initialization step of the registration method of Krishnan et al. [15] (though in their case, they think of the inter-patch transforms as the unknown variables instead of the point positions).

6.2 Global embeddings from edge lengths

Similar approaches have been applied to the (NP-HARD) problem of solving for the framework of a graph given its edge lengths. In these approaches one first attempts to find local d -dimensional embeddings for each one-ring (a vertex and its neighbors) of the framework

up to an unknown local Euclidean transform. This step alone is NP-HARD and can fail. But assuming this step is (approximately) successful one can reduce the rest of the problem to the Euclidean registration problem above.

In the As-Affine-As-Possible (AAAP) method [14, 31], this was done using what is essentially a strong affinity matrix. In the Locally-Rigid-Embedding (LRE) method [23] this was done using a non-symmetric equilibrium stress matrix. Both approaches then removed the global affine transform using the least squares linear system described in Appendix B.

6.3 Manifold learning

Many of the ideas of affine rigidity first appeared in the context of manifold learning. Suppose one has d -dimensional smooth manifold \mathcal{M} which is a topological d -ball embedded in a larger D -dimensional space \mathbb{R}^D . Also suppose that one has a set \mathcal{V} of sample vertices on the manifold. In manifold learning, one first connects nearby samples to form a proximity graph Γ . One then looks for a framework (p, Γ) of this graph in \mathbb{R}^d that in some way that preserves some of the geometric relations of the points in \mathbb{R}^D . This is used to represent a parametrization of \mathcal{M} .

To compute the coordinates p , the Locally Linear Embedding (LLE) method [19] builds a matrix Ω with structure similar to a non-symmetric equilibrium stress matrix. In particular, row i encodes one affine relation between vertex i and its neighbors in \mathbb{R}^D . Then (after ignoring the all ones vector) the smallest d eigenvectors of $\Omega^t\Omega$ are used to form the coordinates of p in \mathbb{R}^d . Unfortunately, since the original embedding is in \mathbb{R}^D , for a graph of high enough valence and assuming no noise, Ω must have a kernel of size at least $D+1$, which is much larger than $d+1$. Thus it is not clear how the numerically smallest $d+1$ eigenvectors will behave. The paper suggests to add an additional regularization term, possibly to address this issue. A follow up to the LLE paper [20] describes a PCA-LLE variant where a d -dimensional local PCA is computed to “flatten” each one-ring before calculating its corresponding row in the matrix Ω . Thus Ω is designed to represent d -dimensional affine relations between the points. The Local-Tangent-Space-Alignment (LTSA) method [32] is an interesting variant of PCA-LLE. In this method, a $v \times v$ matrix N is formed that is the Hessian of a quadratic energy. Thus this matrix plays the role of a strong affinity matrix. It is in this context that Zha and Zhang investigated the rank of this matrix and affine rigidity [30].

In all of these methods, an understanding of affine rigidity is important. In particular it tells us what the rank of the computed matrix would be if the original d -dimensional manifold was in fact embedded in \mathbb{R}^d . For example, if such a framework was not affinely rigid in \mathbb{R}^d , then the kernel would be too big, and we would not expect a manifold learning technique to succeed. However, in manifold learning the embedding has an affine span greater than d and the analysis becomes more difficult. The kernel of a strong D -dimensional affinity matrix is too large, while the kernel of a strong affinity matrix for the locally flattened configurations contains only the all-ones vector but it is hoped that the numerically next smallest d eigenvectors are somehow geometrically meaningful. For an analysis along these lines, see [24].

A Matrix rank

For completeness, we recall the necessary material for determining the generic matrix rank. This material is standard. For a more detailed treatment, see, e.g., [11, Section 5].

We will consider the general setting where there is a set of linear constraints that must be satisfied by a vector $m \in \mathbb{R}^n$. The entries of m are then arranged in some fixed manner as the entries of a matrix M , whose rank we wish to understand. The linear constraints are described by a constraint matrix with n columns: $C(p)$. Each of the entries of $C(p)$ is defined by some polynomial function, with coefficients in \mathbb{Q} , of the coordinates of an input configuration p . We wish to study the behavior of the rank of M as one changes p .

We apply this in the proof of Corollary 3.8 where the constraints C specify that the matrix M is an affinity matrix for (p, Θ) and in the proof of Proposition 5.8 where the constraints C specify that the matrix M is a non-symmetric equilibrium stress matrix for (p, Γ) .

Proposition A.1. *Suppose that for some generic p , there is a matrix M of rank s consistent with $C(p)m = 0$.*

Then for all generic p , there is some matrix M of rank $\geq s$ consistent with $C(p)m = 0$.

Proof. To prove this proposition we first need the following lemma.

Lemma A.2. *Let $M(\pi)$ be a matrix whose entries are polynomial functions with rational coefficients in the variables $\pi \in \mathbb{R}^n$. Let r be a rank achieved by some $M(\pi_0)$. Then $\text{rank}(M(\pi)) \geq r$ for all points π that are generic in \mathbb{R}^n .*

Proof. The rank of the $M(\pi)$ is less than r iff the determinants of all of the $r \times r$ submatrices vanish. Let $\pi_0 \in \mathbb{R}^n$ be a choice of parameters so $M(\pi_0)$ has rank r . Then there is an $r \times r$ submatrix $T(\pi_0)$ of $M(\pi_0)$ with non-zero determinant. Thus $\det(T(\pi))$ is a non-zero polynomial of π . For any π with $\text{rank}(M(\pi)) < r$, this determinant must vanish. Thus, any such π cannot be generic. \square

Next we recall that for a non-singular $n \times n$ matrix \hat{C} ,

$$\text{adj}(\hat{C}) = \det(\hat{C})\hat{C}^{-1}, \quad (1)$$

where $\text{adj} \hat{C}$ is the *adjugate matrix* of \hat{C} , the conjugate of the cofactor matrix of \hat{C} . In particular, $\text{adj} \hat{C}$ is a polynomial in \hat{C} .

For a given p , let $t(p)$ be the rank of $C(p)$. Let $t := \max_p t(p)$. By Lemma A.2 this maximum is obtained for generic p .

For each p we add a set H of $n - t$ additional rows to $C(p)$ to obtain a matrix $C(p, H)$, and determine m by solving the linear system $C(p, H)m = b$ where $b \in \mathbb{R}^n$ is a vector of all zeros except for a single 1 in one of the positions of a row in H (if there are any rows in H). This m is then converted to a matrix $M(p, H)$. $M(p, H)$ is well-defined iff this

linear system has a unique solution, i.e., iff $C(p, H)$ has rank n . Note that this happens for generic p and H .

Let p_0 be generic and have a compatible matrix M_0 with rank s , as in the hypotheses of the proposition. Find a set H_0 of additional rows so that $C(p_0, H_0)$ has rank n and $C(p_0, H_0)m_0 = b$. Let $\hat{C}(p, H)$ be an $n \times n$ submatrix of $C(p, H)$ so that $\hat{C}(p_0, H_0)$ is invertible. (\hat{C} necessarily uses t rows from $C(p)$ and all rows of H .) Define \hat{b} similarly, let $\tilde{m}(p, H) := \text{adj}(\hat{C})\hat{b}$, and let $\tilde{M}(p, H)$ be the associated matrix.

By Lemma A.2, the rank of $\hat{C}(p, H)$ is equal to its maximum value, n , at all points (p, H) that are not zeros of a polynomial $P_1(p, H) := \det \hat{C}(p, H)$. Moreover, when $P_1(p, H) \neq 0$, the linear equation defining $M(p, H)$ has a unique solution and the adjugate matrix $\tilde{M}(p, H)$ is a scalar multiple of $M(p, H)$. In particular we have assumed (p_0, H_0) is not a zero of P_1 and thus $\tilde{M}(p_0, H_0)$ has rank s . By Lemma A.2 again, the rank of $\tilde{M}(p, H)$ is less than s only at the zeros of a non-zero polynomial $P_2(p, H)$.

For any generic p , there must be some generic point (p, H) . Such a generic (p, H) cannot be a zero of P_1 or P_2 and thus $\tilde{M}(p, H)$ and $M(p, H)$ must have rank no less than s . \square

B Removing the Affine Transform

Suppose one has solved for q – a configuration in \mathbb{R}^d – up to an unknown global affine transform A of the true configuration p : $p = A(q)$. Given a set of edge lengths for p , it is possible to compute A up to an unknown global Euclidean transform. This approach was described by Singer [23].

In particular, let L be a $d \times d$ matrix representing the linear portion of A and let $G := L^T L$. Now consider the following set of linear equations (in the $\frac{d(d+1)}{2}$ unknowns of G): For each pair of vertices i, j , whose edge lengths are known, we require

$$(q(i) - q(j))^T G (q(i) - q(j)) = (p(i) - p(j))^T (p(i) - p(j)) \quad (2)$$

(Since we have more constraints than unknowns, for numerical purposes we would solve this as a least squares linear system in the unknown G .)

The only remaining concern is whether this system has more than one solution. The solution to Equation (2) will be unique as long as our set of edges with known lengths do not lie on a conic at infinity. Fortunately, we can conclude from Proposition 4.3 of [9] that if our known lengths form a graph $B(\Theta)$ with minimal valence at least d and p is generic, then the edges do not lie on a conic at infinity. This property holds for any hypergraph that is generically globally rigid in \mathbb{E}^d . Using Cholesky decomposition on G then yields L up to a global Euclidean transform.

References

- [1] ALFAKIH, A. Y. On the universal rigidity of generic bar frameworks. *Contributions to Discrete Mathematics* 5, 1 (2010).

- [2] ANDERSON, B., BELHUMEUR, P., EREN, T., GOLDENBERG, D., MORSE, A., WHITELEY, W., AND YANG, Y. Graphical properties of easily localizable sensor networks. *Wireless Networks* 15, 2 (2009), 177–191.
- [3] ASIMOW, L., AND ROTH, B. The rigidity of graphs. *Trans. Amer. Math. Soc.* 245 (1978), 279–289.
- [4] BELK, M., AND CONNELLY, R. Realizability of graphs. *Discrete & Computational Geometry* 37, 2 (2007), 125–137.
- [5] BEZDEK, K., AND CONNELLY, R. Pushing disks apart—the Kneser-Poulsen conjecture in the plane. *J. Reine Angew. Math.* 553 (2002), 221–236.
- [6] BRAND, M. From subspaces to submanifolds. In *Proceedings of the 15th British Machine Vision Conference, London, England* (2004).
- [7] CHEUNG, M., AND WHITELEY, W. Transfer of global rigidity results among dimensions: Graph powers and coning. Preprint, York University, 2005.
- [8] CONNELLY, R. Rigidity and energy. *Invent. Math.* 66, 1 (1982), 11–33.
- [9] CONNELLY, R. Generic global rigidity. *Discrete Comput. Geom* 33, 4 (2005), 549–563.
- [10] EREN, T., GOLDENBERG, O., WHITELEY, W., YANG, Y., MORSE, A., ANDERSON, B., AND BELHUMEUR, P. Rigidity, computation, and randomization in network localization. In *INFOCOM 2004. Twenty-third Annual Joint Conference of the IEEE Computer and Communications Societies* (2004), vol. 4.
- [11] GORTLER, S., HEALY, A., AND THURSTON, D. Characterizing generic global rigidity. *American Journal of Mathematics* 132, 4 (2010), 897–939, arXiv:0710.0926.
- [12] GORTLER, S., AND THURSTON, D. Characterizing the universal rigidity of generic frameworks, 2010, arXiv1001.0172. Preprint.
- [13] HENDRICKSON, B. Conditions for unique graph realizations. *SIAM J. Comput.* 21, 1 (February 1992), 65–84.
- [14] KOREN, Y., GOTSMAN, C., AND BEN-CHEN, M. PATCHWORK: Efficient localization for sensor networks by distributed global optimization. Technical Report, 2005. Available at <http://www.cs.technion.ac.il/~gotsman/patchwork-tr.pdf>.
- [15] KRISHNAN, S., LEE, P., MOORE, J., AND VENKATASUBRAMANIAN, S. Global registration of multiple 3D point sets via optimization-on-a-manifold. In *Proceedings of the third Eurographics symposium on Geometry processing* (2005), Eurographics Association, p. 187.
- [16] LEE, A. *Geometric Constraint Systems with Applications in CAD and Biology*. PhD thesis, University of Massachusetts Amherst, 2008.
- [17] LINIAL, N., LONDON, E., AND RABINOVICH, Y. The geometry of graphs and some of its algorithmic applications. *Combinatorica* 15, 2 (1995), 215–245.

- [18] LINIAL, N., LOVÁSZ, L., AND WIGDERSON, A. Rubber bands, convex embeddings and graph connectivity. *Combinatorica* 8, 1 (1988), 91–102.
- [19] ROWEIS, S., AND SAUL, L. Nonlinear dimensionality reduction by locally linear embedding. *Science* 290, 5500 (2000), 2323–2326.
- [20] SAUL, L., AND ROWEIS, S. Think globally, fit locally: unsupervised learning of low dimensional manifolds. *The Journal of Machine Learning Research* 4 (2003), 119–155.
- [21] SAXE, J. B. Embeddability of weighted graphs in k -space is strongly NP-hard. In *Proc. 17th Allerton Conf. in Communications, Control, and Computing* (1979), pp. 480–489.
- [22] SCHRECK, P., AND MATHIS, P. Geometrical constraint system decomposition: a multi-group approach. *Int. J. Comput. Geometry Appl.* 16, 5-6 (2006), 431–442.
- [23] SINGER, A. A remark on global positioning from local distances. *Proceedings of the National Academy of Sciences* 105, 28 (2008), 9507.
- [24] SMITH, A., HUO, X., AND ZHA, H. Convergence and rate of convergence of a manifold-based dimension reduction algorithm. In *NIPS* (2008), pp. 1529–1536.
- [25] SO, A., AND YE, Y. Theory of semidefinite programming for sensor network localization. *Mathematical Programming* 109, 2 (2007), 367–384.
- [26] TURK, G., AND LEVOY, M. Zippered polygon meshes from range images. In *Proceedings of the 21st annual conference on Computer graphics and interactive techniques* (1994), ACM, pp. 311–318.
- [27] TUTTE, W. T. How to draw a graph. *Proc. London Math. Soc* 13, 3 (1963), 743–768.
- [28] VARGA, R. *Matrix iterative analysis*. Springer, 2000.
- [29] WHITELEY, W. A matroid on hypergraphs, with applications in scene analysis and geometry. *Discrete and Computational Geometry* 4, 1 (1989), 75–95.
- [30] ZHA, H., AND ZHANG, Z. Spectral properties of the alignment matrices in manifold learning. *SIAM Review* 51, 3 (2009), 545–566.
- [31] ZHANG, L., LIU, L., GOTSMAN, C., AND GORTLER, S. An as-rigid-as-possible approach to sensor network localization. *ACM Transactions on Sensor Networks (TOSN)* 6, 4 (2010), 1–21.
- [32] ZHANG, Z., AND ZHA, H. Principal manifolds and nonlinear dimensionality reduction via tangent space alignment. *SIAM Journal of Scientific Computing* 26 (2004), 313–338.