

WHO NEEDS CROSSINGS?: NONCROSSING LINKAGES ARE UNIVERSAL, AND DECIDING (GLOBAL) RIGIDITY IS HARD*

Zachary Abel,[†] Erik D. Demaine,[‡] Martin L. Demaine,[‡] Sarah Eisenstat,[‡] Jayson Lynch,[‡] and Tao B. Schardl.[‡]

ABSTRACT. We exactly settle the complexity of graph realization, graph rigidity, and graph global rigidity as applied to three types of graphs: “globally noncrossing” graphs, which avoid crossings in all of their configurations; matchstick graphs, with unit-length edges and where only noncrossing configurations are considered; and unrestricted graphs (crossings allowed) with unit edge lengths (or in the global rigidity case, edge lengths in $\{1, 2\}$). We show that all nine of these questions are complete for the class $\exists\mathbb{R}$, defined by the Existential Theory of the Reals, or its complement $\forall\mathbb{R}$; in particular, each problem is (co)NP-hard.

One of these nine results—that realization of unit-distance graphs is $\exists\mathbb{R}$ -complete—was shown previously by Schaefer (2013), but the other eight are new. We strengthen several prior results. Matchstick graph realization was known to be NP-hard (Eades & Wormald 1990, or Cabello et al. 2007), but its membership in NP remained open; we show it is complete for the (possibly) larger class $\exists\mathbb{R}$. Global rigidity of graphs with edge lengths in $\{1, 2\}$ was known to be coNP-hard (Saxe 1979); we show it is $\forall\mathbb{R}$ -complete.

The majority of the paper is devoted to proving an analog of Kempe’s Universality Theorem—informally, “there is a linkage to sign your name”—for globally noncrossing linkages. In particular, we show that any polynomial curve $\varphi(x, y) = 0$ can be traced by a noncrossing linkage, settling an open problem from 2004. More generally, we show that the regions in the plane that may be traced by a noncrossing linkage are precisely the compact semialgebraic regions (plus the trivial case of the entire plane). Thus, no drawing power is lost by restricting to noncrossing linkages. We prove analogous results for matchstick linkages and unit-distance linkages as well.

1 Introduction

The rise of the steam engine in the mid-1700s led to an active study of *mechanical linkages*, typically made from rigid bars connected together at hinges. For example, steam engines need to convert the linear motion of a piston into the circular motion of a wheel, a problem

*A preliminary version of this paper appeared at the 32nd International Symposium on Computational Geometry, June 2016. Originally invited to the corresponding special issue, but the preparation of the final paper was delayed.

[†]MIT Department of Electrical Engineering and Computer Science, 50 Vassar St., Cambridge, MA 02139, USA, zabel@mit.edu. Partially supported by an NSF Graduate Research Fellowship.

[‡]MIT Computer Science and Artificial Intelligence Laboratory, 32 Vassar St., Cambridge, MA 02139, USA, {edemaine,mdemaine,seisenst,jaysonl,neboat}@mit.edu

Graph type	Realization	Rigidity	Global rigidity	Universality
General	$\exists\mathbb{R}$ -complete [22]	$\forall\mathbb{R}$ -complete [22]	$\forall\mathbb{R}$-complete (CoNP-hard [21])	Compact semialg. [15]
Globally noncrossing (no configs. cross)	$\exists\mathbb{R}$-complete	$\forall\mathbb{R}$-complete	$\forall\mathbb{R}$-complete	Compact semialg.
Matchstick graph (unit + noncrossing)	$\exists\mathbb{R}$-complete (NP-hard [8])	$\forall\mathbb{R}$-complete	$\forall\mathbb{R}$-complete	Bounded semialg.
Unit edge lengths (allowing crossings)	$\exists\mathbb{R}$ -complete [22]	$\forall\mathbb{R}$-complete	Open (do they even exist?)	Compact semialg.
Edge lengths in $\{1, 2\}$ (allowing crossings)	$\exists\mathbb{R}$ -complete [22]	$\forall\mathbb{R}$-complete	$\forall\mathbb{R}$-complete (CoNP-hard [21])	Compact semialg.

Table 1: Summary of our results (bold) compared with earlier results (cited). The rows give the special types of graphs considered. The middle three columns give complexity results for the three natural decision problems about graph embedding; all completeness results are strong. The rightmost column gives the exact characterization of drawable sets.

solved approximately by Watt’s parallel motion linkage (1784) and exactly by Peaucellier’s invensor (1864) [7, Section 3.1]. These and other linkages are featured in an 1877 book called *How to Draw a Straight Line* [13] by Alfred Bray Kempe—a barrister and amateur mathematician in London, perhaps most famous for his false “proof” of the Four-Color Theorem [14] that nonetheless introduced key ideas used in the correct proofs of today [3, 19].

Kempe’s Universality Theorem. Kempe investigated far beyond drawing a straight line by turning a circular crank. In 1876, he claimed a universality result, now known as Kempe’s Universality Theorem: every polynomial curve $\varphi(x, y) = 0$ can be traced by a vertex of a 2D linkage [12]. Unfortunately, his “proof” was again flawed: the linkage he constructs indeed traces the intended curve, but also traces finitely many unintended additional curves. Fortunately, his idea was spot on.

Many researchers have since solidified and/or strengthened Kempe’s Universality Theorem [10, 9, 15, 1, 22]. In particular, small modifications to Kempe’s gadgets lead to a working proof [1, 7, Section 3.2]. Furthermore, the regions of the plane drawable by a 2D linkage (other than the entire plane $\mathbb{R}^2$) are exactly compact semialgebraic regions¹ [15, 1]. By carefully constructing these linkages to have rational coordinates, Abbott, Barton, and Demaine [1] showed how to reduce the problem of testing isolatedness of a point in an algebraic set² to testing rigidity of a linkage. Isolatedness was proved coNP-hard [16] and then $\forall\mathbb{R}$ -complete³ [22]; thus linkage rigidity is $\forall\mathbb{R}$ -complete.

¹A compact planar region is *semialgebraic* if it can be obtained as the intersection and/or union of finitely many basic sets defined by polynomial inequalities $p(x, y) \geq 0$.

²A set $S \subseteq \mathbb{R}^d$ is *algebraic* if it can be written as the set of solutions to a polynomial equation $p(x_1, \dots, x_d) = 0$. (Algebraic sets are closed under intersection and finite union.)

³The class $\forall\mathbb{R} = \text{co-}\exists\mathbb{R}$ consists of decision problems whose complement (inverting yes/no instances) belong to $\exists\mathbb{R}$. The class $\exists\mathbb{R}$ refers to the problems (Karp) reducible to the *existential theory of the reals* ($\exists x_1 : \dots \exists x_n : \pi(x_1, \dots, x_n)$ for an arithmetic predicate $\pi : \mathbb{R}^n \rightarrow \{\text{true}, \text{false}\}$), which is somewhere between

Our results: no crossings. See Table 1 for a summary of our results in comparison to past results. Notably, all known linkage constructions for Kempe’s Universality Theorem (and its various strengthenings) critically need to allow the bars to cross each other. In practice, certain crossings can be made physically possible, by placing bars in multiple parallel planes and constructing vertices as vertical pins. Without extreme care, however, bars can still be blocked by other pins, and it seems difficult to guarantee crossing avoidance for complex linkages. Beyond these practical issues, it is natural to wonder whether allowing bars to cross is necessary to achieve linkage universality. Don Shimamoto first posed this problem in April 2004, and it was highlighted as a key open problem in the first chapter of *Geometric Folding Algorithms* [7].

We solve this open problem by strengthening most of the results mentioned above to work for *globally noncrossing graphs*, that is, graphs plus edge-length constraints that alone force all configurations to be (strictly) noncrossing.⁴ In particular, we prove the following universality and complexity results.

1. The planar regions drawable by globally noncrossing linkages are exactly $(\mathbb{R}^2$ and) the compact semialgebraic regions (Theorem 4.8), settling Shimamoto’s 2004 open problem.
2. Testing whether a globally noncrossing graph with constant-sized integer edge lengths has any valid configurations is $\exists\mathbb{R}$ -complete (Theorem 4.6).
3. Testing rigidity is strongly⁵ $\forall\mathbb{R}$ -complete even for globally noncrossing graphs with constant-sized integer edge lengths that are drawn with integer vertex coordinates (Theorem 4.7).
4. Testing global rigidity (uniqueness of a given embedding) is strongly $\forall\mathbb{R}$ -complete even for globally noncrossing graphs with constant-sized integer edge lengths that are drawn with integer vertex coordinates (Theorem 4.7).

Our techniques are quite general and give us results for two other restricted forms of graphs as well. First, *matchstick graphs* are graphs with *unit* edge-length constraints, and where only (strictly) noncrossing configurations are considered valid. We prove in Section 6 the following universality and complexity results:

5. The planar regions drawable by matchstick graphs are exactly $(\mathbb{R}^2$ and) the bounded semialgebraic regions (Theorem 7.25). Notably, unlike all other models considered, matchstick graphs enable the representation of open boundaries in addition to closed (compact) boundaries.

NP and PSPACE (by [6]). The classic example of an $\exists\mathbb{R}$ -complete problem is pseudoline stretchability [18]. The classes $\exists\mathbb{R}$ and $\forall\mathbb{R}$ are discussed more thoroughly in Section 3.2.

⁴Thus, the noncrossing constraint can be thought of as being “required” or not of a configuration; in either case, the configurations (even those reachable by discontinuous motions) will be noncrossing.

⁵A problem is said to be *strongly* hard if it remains hard when all integers in the input are specified in unary rather than binary. In other words, all integers are polynomially bounded in the size of the input rather than exponentially bounded. All of our hardness results show strong hardness when applicable.

6. Deciding whether an abstract graph can be drawn as a matchstick graph is $\exists\mathbb{R}$ -complete (Theorem 6.5). This result strengthens a 25-year-old NP-hardness result [8, 5], and settles an open question of [22].
7. Testing rigidity or global rigidity of a matchstick graph is strongly $\forall\mathbb{R}$ -complete (Theorems 6.6 and 6.7).

Second, we consider restrictions on edge lengths to be either all equal (unit) or all in $\{1, 2\}$, but at the price of allowing crossing configurations. Recognizing unit-distance graphs is already known to be $\exists\mathbb{R}$ -complete [22]. We prove in Section 5 the following additional universality and complexity results:

8. The planar regions drawable by unit-edge-length linkages are exactly the compact semialgebraic regions (and $\mathbb{R}^2$) (Theorem 5.9), proving a conjecture of Schaefer [22].
9. Testing rigidity of unit-edge-length graphs is strongly $\forall\mathbb{R}$ -complete (Theorem 5.7), proving a conjecture of Schaefer [22].
10. Testing global rigidity of graphs with edge lengths in $\{1, 2\}$ is strongly $\forall\mathbb{R}$ -complete (Theorem 5.8). This result strengthens a 35-year-old strong-coNP-hardness result for the same scenario [21]. While it would be nice to strengthen this result to unit edge lengths, we have been unable to find even a single globally rigid equilateral linkage larger than a triangle.

We introduce several techniques to make noncrossing linkages manageable in this setting. In Section 7.1 we define *extended linkages* to allow additional joint types, in particular, requiring angles between pairs of bars to stay within specified intervals. Section 7.2 then shows how to draw a polynomial curve and obtain Kempe’s Universality Theorem with these powerful linkages while avoiding crossings, by following the spirit of Kempe’s original construction but with specially designed modular gadgets to guarantee no crossings between (or within) the gadgets. We simulate extended linkage with linkages that have chosen subgraphs marked as rigid. In turn, in Sections 4–6, we simulate these “partially rigidified” linkages with the three desired linkage types: globally noncrossing, unit-distance or $\{1, 2\}$ -distance, and matchstick.

2 Description of the Main Construction

The heart of this paper is a single, somewhat intricate linkage construction. In this section, we describe and discuss the properties of this construction in detail, after building up the necessary terminology.

2.1 Definitions: Linkages and Graphs

Unless otherwise specified, all graphs $G = (V(G), E(G), \ell)$ in this text are connected, edge-weighted with positive edge lengths $\ell(e) > 0$, and contain no self-loops. We first recall and establish notation for our primary objects of study, linkages.

Definition 2.1 (Linkages). An *abstract linkage*, or simply a *linkage*, is a triple $\mathcal{L} = (G, W, P)$ consisting of a weighted graph $G = (V(G), E(G), \ell)$ together with a choice of *pin locations* $P(w) \in \mathbb{R}^2$ for vertices w in a chosen subset $W \subset V(G)$ of *pinned vertices*.

Definition 2.2 (Linkage Configurations). A *configuration* of a linkage $\mathcal{L} = (G, W, P)$ is an assignment of vertex locations $C : V(G) \rightarrow \mathbb{R}^2$ respecting the edge-length and pin assignments: $|C(u) - C(v)| = \ell(uv)$ for each edge uv , and $C(w) = P(w)$ for each pinned vertex $w \in W$. Two configurations are *congruent* if they differ only by a Euclidean transformation, i.e., a translation, rotation, and possibly reflection.

The *configuration space* $\text{Conf}(\mathcal{L}) \subseteq (\mathbb{R}^2)^{|V(G)|}$ is the set of all configurations; it is a closed, algebraic subset of $(\mathbb{R}^2)^{|V(G)|}$. (Be warned, however, that some modified types of linkages to follow, notably *NX-constrained* linkages or *extended linkages*, will have only *semialgebraic* configuration spaces.)

An abstract linkage is called *configurable* or *realizable* if its configuration space is nonempty, and a *configured linkage* is a linkage together with a chosen *initial configuration* $C_0 \in \text{Conf}(\mathcal{L})$. Configured linkage $(\mathcal{L}, C_0)$ is *rigid* if there is no nontrivial continuous deformation of $\mathcal{L}$, i.e., C_0 is isolated in $\text{Conf}(\mathcal{L})$. Similarly, $(\mathcal{L}, C_0)$ is *globally rigid* if C_0 is the *only* configuration in $\text{Conf}(\mathcal{L})$.

Convention 2.3. As a convenient abuse of notation, we often write v instead of $C(v)$ when configuration C is understood from context.

We consider *abstract graphs* and *configured graphs* as abstract or configured linkages without pins ($P = \emptyset$), with one key difference: rigidity and global rigidity for graphs are more liberally defined to allow Euclidean motions.

Definition 2.4 (Graph Rigidity and Global Rigidity). A configured graph (G, C_0) is *rigid* if the only continuous deformations of (G, C_0) are rigid Euclidean motions, i.e., if there is a neighborhood of C_0 in $\text{Conf}(G)$ consisting only of configurations congruent to C_0 . Likewise, (G, C_0) is *globally rigid* if all configurations are congruent to C_0 .

As defined above, configurations may have coincident vertices, vertices in the middle of edges, and/or properly crossing edges. The following notion forbids these undesirable features:

Definition 2.5 (Noncrossing Configurations). A configuration C of a linkage $\mathcal{L}$ is *non-crossing* if distinct edges intersect only at common endpoints: for any pair of incident edges $uv \neq uv' \in E(G)$ sharing a vertex u , the segments $C(u)C(v)$ and $C(u)C(v')$ intersect only at $C(u)$ in $\mathbb{R}^2$, and for any pair $uv, u'v'$ of disjoint edges in G , segments $C(u)C(v)$ and $C(u')C(v')$ are disjoint in $\mathbb{R}^2$.

Let $\text{NXConf}(\mathcal{L}) \subseteq \text{Conf}(\mathcal{L})$ denote the subset of noncrossing configurations. Then we say $\mathcal{L}$ is *globally noncrossing* if all of its configurations are noncrossing, i.e., $\text{NXConf}(\mathcal{L}) = \text{Conf}(\mathcal{L})$. If C is a noncrossing configuration, the *minimum feature size* of C is the shortest distance from a segment $C(u)C(v)$ to a point $C(w)$ for $uv \in E(G)$ and $w \in V(G) \setminus \{u, v\}$. If linkage $\mathcal{L}$ is globally noncrossing, its *global minimum feature size* is defined as the infimum of the minimum feature size of its configurations. If $\text{Conf}(\mathcal{L})$ is compact, this infimum is

achieved and is strictly greater than 0. (As a special case, if $\mathcal{L}$ is not realizable, then it is vacuously globally noncrossing and its global minimum feature size is $+\infty$.)

Definition 2.6 (Combinatorial Embeddings, Corners). A *combinatorial embedding* σ of a graph G consists of a cyclic ordering σ_v of v 's incident edges for each vertex $v \in V(G)$. A noncrossing configuration C of G *agrees with* σ if the counterclockwise cyclic ordering of segments $C(v)C(w)$ around $C(v)$ matches σ_v for each vertex v .

Whenever edge vu is followed by vw in σ_v , the two-edge path $\Lambda = uvw$ is a *corner* of σ at v ; in the special case where v is incident to exactly one edge vu , there is a single corner, uvu , at v . If C is a noncrossing configuration of G agreeing with σ , the *angles* of C at v are the angles $\angle C(\Lambda) := \angle C(u)C(v)C(w)$ for each corner $\Lambda = uvw$. When there is only one corner uvu at v , we define $\angle C(u)C(v)C(u)$ to have measure 360° . With this convention, the angles of C at v add to 360° no matter the degree of v . (Recall that all graphs are connected, so $\deg(v) \geq 1$.)

2.2 Constrained Linkages

We will make use of a number of special-purpose “constraints” or “annotations” that may be attached to linkages to artificially modify their behavior, such as “rigid constraints” that “rigidify” a subgraph into a chosen configuration while allowing the rest of the linkage to move freely. These annotations do not affect the linkage itself; instead, they merely indicate which configurations of the linkage they consider acceptable. The language of constraints allows us to separate a desired *effect* from the *implementation* or *construction* that enforces that effect. For example, Sections 4 through 6 develop three different techniques to force subgraphs to remain rigid—three different “implementations” of the rigidifying constraint—allowing a majority of the work, namely the Main Theorem (Theorem 2.14), to be reused in all three contexts. We now define constraints in general, and the rigid constraint in particular, more formally.

Definition 2.7. A *constraint* Con on an abstract linkage $\mathcal{L}$ is specified by a subset of the configuration space, $\text{Con} \subseteq \text{Conf}(\mathcal{L})$, and we say the configurations $C \in \text{Con}$ *satisfy* constraint Con . A *constrained linkage* $\mathcal{L}$ is an abstract linkage $\mathcal{L}_0$ together with a finite set K of constraints on $\mathcal{L}_0$, and the *constrained configuration space* is defined as $\text{Conf}(\mathcal{L}) := \text{Conf}(\mathcal{L}_0) \cap \bigcap_{\text{Con} \in K} \text{Con}$. In other words, constrained linkage $\mathcal{L}$ simply ignores any configurations of $\mathcal{L}_0$ that don't satisfy all of its constraints.

All terms discussed in Section 2.1—realizability, rigidity, global rigidity, etc.—apply equally well to constrained linkages via their *constrained* configuration space.

Definition 2.8. A *rigid constraint* $\text{RigidCon}_{\mathcal{L}}(H, C_H)$ on a linkage $\mathcal{L} = (G, W, P)$ is specified by a connected subgraph $H \subseteq G$ (whose edge lengths match those of G) together with a configuration C_H of H . A configuration $C \in \text{Conf}(\mathcal{L})$ satisfies the rigid constraint when C induces a configuration $C|_H$ on H that is congruent to the given C_H , i.e., differs only by a (possibly orientation-reversing) Euclidean transformation. When a constrained linkage $\mathcal{M}$ possesses constraint $\text{RigidCon}_{\mathcal{M}}(H, C_H)$, we say (H, C_H) is a *rigidified subgraph* of $\mathcal{M}$. A constrained linkage all of whose constraints are rigid constraints is called a *partially rigidified linkage*.

Other constraint types—the noncrossing constraint, angle constraint, and sliceform constraint—will be introduced as they are needed, in Sections 6 and 7.1.

2.3 Drawing with (Constrained) Linkages

Definition 2.9 (Linkage Trace and Drawing). For a (possibly constrained) linkage $\mathcal{L}$ and a tuple $X = (v_1, \dots, v_k)$ of distinct vertices of $\mathcal{L}$, the *trace* of X is defined as the image $\pi_X(\text{Conf}(\mathcal{L})) \subset (\mathbb{R}^2)^k$, where π_X is the projection map sending a configuration $C \in \text{Conf}(\mathcal{L})$ to $\pi_X(C) := (C(v_1), \dots, C(v_k))$. This trace is semialgebraic, and if $\text{Conf}(\mathcal{L})$ is compact, the trace is also compact. A linkage $(\mathcal{L}, X)$ is said to *draw* its trace, and a set $R \subseteq (\mathbb{R}^2)^k$ is *drawable (by a linkage)* if it can be expressed as the trace of some k vertices of a linkage.

The term “draw” is somewhat misleading, because a linkage might not be able to “draw” its entire trace through a single continuous motion. For example, some linkages have a disconnected trace, such as the linkage $\mathcal{T}$ formed by an equilateral triangle abc pinned at a and c , so that $X = \{b\}$ draws two separate points corresponding to the two possible triangle orientations. Even some linkages with connected traces cannot draw their trace with a continuous motion: for example, if we modify $\mathcal{T}$ by adding a longer edge bd , then d draws two intersecting circles, even though it can only draw one circle at a time in a continuous motion (depending on the orientation of triangle abc).

Because of the complicated relationship between a linkage and its trace, some types of drawings have been singled out as particularly nice:

Definition 2.10 (Liftable Drawing). Say $(\mathcal{L}, X)$ draws its trace *liftably* if the map π_X has the *path lifting property*: for any configuration $C \in \text{Conf}(\mathcal{L})$ and path $\gamma : [0, 1] \rightarrow \pi_X(\text{Conf}(\mathcal{L}))$ in the trace starting at $\gamma(0) = \pi_X(C)$, there is a path $\bar{\gamma} : [0, 1] \rightarrow \text{Conf}(\mathcal{L})$ starting at $\bar{\gamma}(0) = C$ and lifting γ , meaning $\gamma = \pi_X \circ \bar{\gamma}$.

For example, node d of linkage $\mathcal{T}$ above does not draw its trace *liftably*, because d cannot move from one circle to the other through a continuous motion of $\mathcal{T}$. By contrast, node c *does* draw liftably (vacuously, as there are no nontrivial paths in c ’s trace), even though c ’s trace is disconnected.

Liftable drawing was introduced in [1] (under the name “continuous” drawing) for its usefulness in studying linkage rigidity. Indeed, if $(\mathcal{L}, X)$ draws liftably and a point $p \in \pi_X(\text{Conf}(\mathcal{L}))$ is *not* isolated in the trace, then any configuration C with $\pi_X(C) = p$ is *not* rigid, because a nontrivial continuous path beginning at p can be lifted to a nontrivial motion beginning at C .

Though we have changed the name from “continuous” to “liftable” for greater clarity between *continuous* motions and *liftable* drawings, the two notions are closely related: indeed, liftability implies that each *connected component* of the trace can be drawn with a continuous motion. Note, however, that the implication does not go the other way. For example, the central vertex v in Watt’s famous linkage traces a lemniscate, and can do so with a continuous motion. However, a path through the trace that “turns a corner” at the central intersection point cannot be lifted to a motion of the full linkage, so v does not draw its lemniscate liftably.

The complementary notion is *rigid* drawing:

Definition 2.11 (Rigid Drawing). Say $(\mathcal{L}, X)$ draws its trace *rigidly* if the map π_X has finite fibers, i.e., for any $p \in \pi_X(\text{Conf}(\mathcal{L}))$, there are only finitely many configurations $C \in \text{Conf}(\mathcal{L})$ with $\pi_X(C) = p$.

For example, node d of linkage $\mathcal{T}$ draws its trace rigidly, because there are at most two configurations corresponding to each possible position of d . By contrast, node c does *not* draw rigidly, because infinitely many configurations exist for each position of c .

Rigid drawing was also introduced in [1], for similar reasons: if $(\mathcal{L}, X)$ draws *rigidly* and if p is isolated in $\pi_X(\text{Conf}(\mathcal{L}))$, then any configuration C with $\pi_X(C) = p$ is rigid, because the discrete set $\pi_X^{-1}(p)$ contains no nonconstant continuous paths.

As in [1], we are especially interested in cases where $(\mathcal{L}, X)$ draws both liftably and rigidly, as that creates a strong correlation between properties of the *trace* and the linkage’s own rigidity or flexibility. We also introduce a stronger notion:

Definition 2.12 (Perfect Drawing). If the map π_X is a *homeomorphism* between $\text{Conf}(\mathcal{L})$ and the trace, we say $(\mathcal{L}, X)$ draws its trace *perfectly*.

Perfect drawing easily implies liftably and rigid drawing, but it is even more restrictive. It is especially useful for *parametrizing* a linkage’s configuration space: if $(\mathcal{L}, X)$ draws perfectly, each configuration of $\mathcal{L}$ is uniquely and continuously determined solely from the locations of vertices in X .

Finally, we will often wish to create a linkage that “behaves like” a related linkage, which we formalize as follows:

Definition 2.13 (Linkage Simulation). When $(\mathcal{L}, X)$ draws precisely the full configuration space $\text{Conf}(\mathcal{M})$ of another linkage $\mathcal{M}$, we say that $(\mathcal{L}, X)$ *simulates* $\mathcal{M}$. It may *liftably, rigidly, or perfectly simulate* $\mathcal{M}$ if it draws $\text{Conf}(\mathcal{M})$ in the corresponding manner.

2.4 Specification of the Main Theorem

For a collection $F = \{f_1, \dots, f_s\}$ of polynomials in $\mathbb{R}[x_1, y_1, \dots, x_m, y_m] = \mathbb{R}[\overrightarrow{xy}]$, the *algebraic set* (or *algebraic variety*) defined by F is the set of common zeros,

$$Z(F) := \{\overrightarrow{xy} \in \mathbb{R}^{2m} \mid f_1(\overrightarrow{xy}) = \dots = f_s(\overrightarrow{xy}) = 0\}.$$

The primary technical construction in this paper builds a globally noncrossing, partially rigidified linkage $\mathcal{L}(F)$ that draws the algebraic set $Z(f_1, \dots, f_s) \subseteq \mathbb{R}^{2m}$, or at least a large enough piece thereof, up to a translation of $\mathbb{R}^{2m}$. This translation is necessary: without it, some algebraic sets would require the drawing vertices in X to collocate in some configurations, precluding the possibility of global noncrossing. For example, two distinct vertices that draw the torus $S^1 \times S^1$, where $S^1 = \{(x, y) \in \mathbb{R}^2 \mid x^2 + y^2 = 1\}$, would be forced to intersect some of the time (e.g., at $(1, 0)$). Allowing translation, this locus is easily drawable by two unit-length edges, each pinned by one endpoint at $(0, 0)$ and $(3, 0)$, respectively.

We are now prepared to precisely specify the properties of this construction, from which the results listed in Table 1 follow as corollaries. We thoroughly detail these properties here, so that the corollaries may be derived solely from Theorem 2.14’s statement without referring to the specifics of its proof (with one small exception, discussed in Section 7.9). This also allows for maximal reuse: the commonalities in our arguments for our three linkage contexts—unconstrained globally noncrossing linkages in Section 4, unit-distance linkages in Section 5, and matchstick linkages in Section 6—have been unified and generalized into Theorem 2.14, so only features unique to each context need to be discussed in Sections 4–6.

The Main Theorem is divided into three parts because it must be used in subtly different ways by the four types of results we seek. Hardness of realizability requires a polynomial-time construction of an abstract linkage that draws $Z(f_1, \dots, f_s)$ *without knowing whether the resulting configuration space is empty*, whereas proving hardness of rigidity and global rigidity requires the polynomial-time construction of a linkage *together with a known configuration*. We thus separate these into different Parts of Theorem 2.14 with slightly different assumptions about the input polynomials f_j (Part II for realizability, Part III for rigidity and global rigidity). When proving universality, we must prove *existence* of a linkage to draw any compact semialgebraic set, but the coefficients of the polynomials defining this set may be irrational or non-algebraic, as might the edge lengths and coordinates of the resulting linkage, so we isolate this in Part I, away from algorithmic and efficiency concerns.

If the input set of polynomials F has $|F| = s$ real polynomials, each of total degree d in the $2m$ variables $(x_1, y_1, \dots, x_m, y_m)$, then the number of coefficients is $s \cdot \binom{2m+d}{d}$ (in dense representation). The number of vertices and edges of the resulting linkage will be bounded by a polynomial in the related quantities m^d , d^d , and s . We do not attempt bounds that are finer tuned depending on whether F is sparse or somehow otherwise simpler than parameters m, d, s might indicate. Indeed, sparse polynomial input does not lead to noticeable efficiency gains with our algorithm, due to a change of coordinates (detailed in Section 7.2) that can turn sparse polynomials into dense ones.

Similarly, numerator bounds on edge lengths and coordinates are written in terms of M , an upper bound on the absolute value of coefficients of the input polynomials. We wish to emphasize that M bounds the number of *unary* digits of coefficients, not binary digits. Nowhere in this paper do we measure the number of binary digits of an integer: the *magnitude*, *absolute value*, or *size* of an integer always refers to its length in unary.

Theorem 2.14. Part I. *Take as input a collection of polynomials $F = \{f_1, \dots, f_s\}$, each in $\mathbb{R}[x_1, y_1, \dots, x_m, y_m]$ with total degree at most d . Then we may construct a partially rigidified linkage $\mathcal{L} = \mathcal{L}(F)$ that draws, up to translation, a bounded portion of the algebraic set $Z(F)$: specifically, there is a translation T on $\mathbb{R}^{2m}$ and a subset X of m vertices of $\mathcal{L}$ such that*

$$T(Z(F) \cap [-1, 1]^{2m}) \subseteq \pi_X(\text{Conf}(\mathcal{L})) \subseteq T(Z(F)).$$

Furthermore, there is a constant integer D depending only on s, m, d (but independent of the coefficients in input F) such that:

1. *Vertex list X draws this trace liftablely and rigidly.*
2. *The number of vertices and edges in $\mathcal{L}$ is $O(\text{poly}(m^d, d^d, s))$.*

3. Each edge of $\mathcal{L}$ has length at least $1/D$, and $\mathcal{L}$ is globally noncrossing with global minimum feature size at least $1/D$.
4. For each constraint $\text{RigidCon}_{\mathcal{L}}(H, C_H)$ on $\mathcal{L}$, H is a tree that connects to $G \setminus H$ only at leaves of H , and configuration C_H has all edges parallel to the x - or y -axes. Tree H has at least three noncollinear vertices, and it does not flip: for each configuration C of $\mathcal{L}$, configuration C_H and the induced configuration $C|_H$ have the same orientation. Each edge of G is contained in at most one rigidified subtree (H, C_H) .
5. There is a combinatorial embedding σ of G such that every configuration $C \in \text{Conf}(\mathcal{L})$ is noncrossing and agrees with σ . Furthermore, for any vertex v of degree at least 2 that is not internal to any rigidified tree, each corner Λ at v has angle $\angle C(\Lambda) \in (60^\circ, 240^\circ)$.
6. Linkage $\mathcal{L}$ has precisely $|P| = 3$ pinned vertices, which belong to a single rigidified tree (H, C_H) and are not collinear in C_H .

Part II. Suppose polynomials f_j have integer coefficients, each bounded in absolute value by M . Then we may bound the complexity of $\mathcal{L}$ as follows:

7. All edge lengths in $\mathcal{L}$ belong to $\frac{1}{D} \cdot \mathbb{Z}$ and have size $O(\text{poly}(m^d, d^d, s, M))$. Edges not contained in a rigidified subtree have lengths at most D .
8. Constrained linkage $\mathcal{L}$, the set X of vertices, translation T , and combinatorial embedding σ may be constructed from F deterministically in time $O(\text{poly}(m^d, d^d, s, M))$.

Part III. Finally, if the polynomials f_j each satisfy $f_j(\vec{0}) = 0$, we may additionally compute an initial configuration C_0 satisfying:

9. All coordinates of C_0 belong to $\frac{1}{D} \cdot \mathbb{Z}$ and have absolute value $O(\text{poly}(m^d, d^d, s, M))$.
10. C_0 is the only configuration of $\mathcal{L}$ that projects to $T(\vec{0}) \in \pi_X(\text{Conf}(\mathcal{L}))$.
11. For each rigidified subtree (H, C_H) , C_0 induces a configuration of H in which all edges are parallel to the x - or y -axes. (Edges not in any rigidified subtree need not be axis-aligned.)
12. C_0 may also be computed deterministically in time $O(\text{poly}(m^d, d^d, s, M))$.

Convention 2.15 (Orientation). Property 4 above refers to the *orientation* of a configured tree. In this paper, the word *orientation* always refers to handedness, not to the angle or direction of an object. So a reflection reverses orientation; a rotation changes angle but preserves orientation.

2.5 Roadmap

The rest of the paper is organized as follows: After gathering a few preliminary facts about real (semi)algebraic sets and the class $\exists\mathbb{R}$ in Section 3, we use the Main Theorem (Theorem 2.14) to prove that graph realizability, rigidity, and global rigidity are $\exists\mathbb{R}$ -complete or $\forall\mathbb{R}$ -complete, and that linkages are universal at drawing semialgebraic sets, in each of three separate contexts: for globally noncrossing graphs/linkages in Section 4, for unit-distance (or $\{1, 2\}$ -distance) graphs/linkages in Section 5, and for matchstick graphs/linkages in Section 6. After all of that, in Section 7, we finally provide the gory details of the Main Construction itself.

3 Preliminaries on Semialgebraic Sets and $\exists\mathbb{R}$

3.1 Semialgebraic Sets as Projections of Algebraic Sets

Here we prove some elementary facts about representing certain semialgebraic sets as projections of (semi)algebraic sets having specific forms, which will be used in proofs of linkage universality (not in complexity results). For a general introduction to real semialgebraic geometry, see [4].

A *basic semialgebraic set* in $\mathbb{R}^k$ is a set of the form $\{\vec{x} \in \mathbb{R}^k \mid f_1(\vec{x}) = \dots = f_s(\vec{x}) = 0, g_1(\vec{x}) > 0, \dots, g_r(\vec{x}) > 0\}$ for polynomials f_i and g_j . Any *semialgebraic set* can be written as a finite union of basic semialgebraic sets [4, Prop. 2.1.8]; in fact, this is one of several equivalent ways to define semialgebraic sets.

Lemma 3.1. *Any bounded semialgebraic set $R \subset \mathbb{R}^k$ can be expressed as the projection of some bounded basic semialgebraic set onto the first k coordinates.*

Proof. We may assume R is nonempty. Any semialgebraic set can be written as a finite union of basic semialgebraic sets [4, Prop. 2.1.8], so write $R = \bigcup_{j=1}^t R_j$, where each R_j is a nonempty basic semialgebraic set. Let $\vec{x}, \vec{x}_1, \dots, \vec{x}_t$ each denote a variable point in $\mathbb{R}^k$, and define $R' \subset (\mathbb{R}^k)^{t+1}$ as the *basic* semialgebraic set defined by the conditions $\vec{x}_j \in R_j$ for $1 \leq j \leq t$, as well as

$$f(\vec{x}, \vec{x}_1, \dots, \vec{x}_t) := \prod_{j=1}^t |\vec{x} - \vec{x}_j|^2 = 0.$$

This last equation exactly stipulates that $\vec{x} = \vec{x}_j$ for some $1 \leq j \leq t$, so $\pi(R') = \bigcup_{j=1}^t R_j = R$, where π denotes the projection onto the coordinates of $\vec{x}$. Set R' is bounded because each R_j is bounded. $\square$

Lemma 3.2. *Any bounded semialgebraic set $R \subset \mathbb{R}^k$ can be expressed as the projection of a bounded set of the form $\{\vec{x} \in \mathbb{R}^m \mid f_1(\vec{x}) = \dots = f_s(\vec{x}) = 0, g_1(\vec{x}) \neq 0, \dots, g_r(\vec{x}) \neq 0\}$, for polynomials f_i and g_j .*

Proof. By Lemma 3.1, we may assume R is a basic semialgebraic set, $R = \{\vec{x} \in \mathbb{R}^k \mid p_1(\vec{x}) = \dots = p_s(\vec{x}), q_1(\vec{x}) > 0, \dots, q_r(\vec{x}) > 0\}$. Because R is bounded, by scaling the polynomials if necessary, we may assume $|p_i(\vec{x})| \leq 1$ and $|q_j(\vec{x})| \leq 1$ for all $\vec{x} \in R$. Now introduce new real variables $\vec{a} = (a_1, \dots, a_r)$ and define

$$R' = \{(\vec{x}, \vec{a}) \in \mathbb{R}^{k+r} \mid p_1(\vec{x}) = \dots = p_s(\vec{x}) = 0, \\ q_1(\vec{x}) = a_1^2, a_1 \neq 0, \dots, q_r(\vec{x}) = a_r^2, a_r \neq 0\}.$$

This set is contained in $R \times [-1, 1]^r$ and is therefore bounded, and R is the projection of R' onto the coordinates of $\vec{x}$. $\square$

A *basic closed semialgebraic set* has the form $\{\vec{x} \in \mathbb{R}^k \mid f_1(\vec{x}) \geq 0, \dots, f_s(\vec{x}) \geq 0\}$; if this set is also bounded, we call it a *basic compact semialgebraic set*.

Lemma 3.3. *Any compact semialgebraic set $R \subset \mathbb{R}^k$ can be expressed as the projection of some compact algebraic set onto the first k coordinates.*

Proof. Compact semialgebraic set R can be written as a finite union of basic compact semialgebraic sets [4, Thm. 2.7.2], and the same proof used in Lemma 3.1 shows that R can be written as a coordinate projection of a basic compact semialgebraic set. (The condition $f = 0$ can be expressed as $f \geq 0$ and $-f \geq 0$.) So it suffices to show that a compact *basic* semialgebraic set is a coordinate projection of a compact algebraic set.

Assume R is a compact basic semialgebraic set,

$$R = \{\vec{x} \in \mathbb{R}^n \mid f_1(\vec{x}) \geq 0, \dots, f_s(\vec{x}) \geq 0\}.$$

Choose new variables $\vec{y} \in \mathbb{R}^s$, and define $R' = \{(\vec{x}, \vec{y}) \mid f_i(\vec{x}) = y_i^2 \text{ for } 1 \leq i \leq s\}$. If π is the projection onto coordinates of $\vec{x}$, then one may check that $\pi(R') = R$. Region R' is bounded because R itself is bounded, so this completes the proof. $\square$

Semialgebraic sets are closed under projection [4, Thm. 2.2.1] (a form of the Tarski–Seidenberg principle), as are compact semialgebraic sets. The lemmas above imply that algebraic sets and basic semialgebraic sets are not closed under projection.

3.2 Existential Theory of the Reals

The class $\exists\mathbb{R}$ is the complexity class consisting of all problems polynomially (Karp) reducible to the Existential Theory of the Reals (ETR), which is the language consisting of true formulas of the form $(\exists x_1, \dots, x_n \in \mathbb{R})\varphi(x_1, \dots, x_n)$, where φ is a quantifier-free predicate over real variables $x_1, \dots, x_n$ using arithmetic symbols $+$, $-$, $\times$, 0 , and 1 , logical predicates $<$, $\leq$, and $=$, and boolean operators $\wedge$, $\vee$ and $\neg$. It is known that $\text{NP} \subseteq \exists\mathbb{R} \subseteq \text{PSPACE}$, though neither inclusion is known to be strict. The first inclusion is simple, as SAT may be encoded with polynomial constraints, so any problem complete for $\exists\mathbb{R}$ is automatically NP-hard. The second inclusion is a nontrivial theorem of Canny [6] and is the tightest known upper bound on the hardness of $\exists\mathbb{R}$.

We now prove the $\exists\mathbb{R}$ hardness of two problems used in the forthcoming reductions. The problem H_2N (probably short for *Hilbert’s homogeneous Nullstellensatz* [22]) asks whether homogeneous polynomials $f_1, \dots, f_s \in \mathbb{Z}[x_1, \dots, x_k]$ have a nonzero common root in $\mathbb{R}^k$. It was introduced in [16] and was shown to be $\exists\mathbb{R}$ -complete in [22], even when all polynomials have degree 4. We wish to reduce from a stronger version of this problem where, additionally, all coefficients are in $\{0, \pm 1, \pm 2\}$. To prove hardness of this problem, only a slight modification of Schaefer’s original argument is necessary. Our starting point is a strengthening of the well-known **CommonZero** problem:

Lemma 3.4. *The **CommonZero** problem—determining whether integer-coefficient polynomials $f_1, \dots, f_s \in \mathbb{Z}[x_1, \dots, x_n]$ have a common solution in $\mathbb{R}^n$ —is $\exists\mathbb{R}$ -complete, even when the polynomials must have total degree at most 2 and coefficients in $\{-1, 0, 1\}$.*

Proof. The problem is a subproblem of ETR, so it is certainly in $\exists\mathbb{R}$. For hardness, we reduce from the **PointConfiguration** problem: given a choice of orientation (clockwise, counterclockwise, or collinear) for each triple $1 \leq i < j < k \leq n$, determine whether there exists a configuration of points $p_1, \dots, p_n$ in the plane so that each triangle $p_i p_j p_k$ has the specified configuration. This is equivalent (in fact, projectively dual) to the more well-known **Stretchability** problem, shown to be $\exists\mathbb{R}$ -complete by Mnëv [18]; see also [23].

If we write $p_i = (p_{i,1}, p_{i,2})$ in coordinates, then the triangle's signed area is given by

$$2 \cdot \text{area}(p_i p_j p_k) = \det(p_i, p_j, p_k) := \det \begin{pmatrix} p_{i,1} & p_{i,2} & 1 \\ p_{j,1} & p_{j,2} & 1 \\ p_{k,1} & p_{k,2} & 1 \end{pmatrix}, \quad (1)$$

which is a homogeneous, degree 2 polynomial in the coordinate variables with coefficients ± 1 . For the triples with collinear configuration, we simply include the polynomial $\det(p_i, p_j, p_k) = 0$. For those with counterclockwise orientation, include the polynomials $\det(p_i, p_j, p_k) = a_{i,j,k}^2$ and $a_{i,j,k} b_{i,j,k} = 1$, for new variables $a_{i,j,k}$ and $b_{i,j,k}$. Similarly, for clockwise triples, include $\det(p_i, p_j, p_k) = -a_{i,j,k}^2$ and $a_{i,j,k} b_{i,j,k} = 1$. It is clear that the point configuration is realizable if and only if the resulting system of polynomials with coefficients in $\{-1, 0, 1\}$ and degree at most 2 has a real solution. $\square$

Theorem 3.5. *The problem H_2N remains $\exists\mathbb{R}$ -complete even when the input polynomials have degree 4 and all coefficients lie in $\{0, \pm 1, \pm 2\}$.*

Proof. The proof by Schaefer [22, Lem. 3.9, Cor. 3.10] for the hardness of H_2N with degree 4 starts with a collection of polynomials of degree at most 2 and transforms them into a collection of homogeneous degree 4 polynomials that have a nontrivial solution if and only if the original collection has any solution. These transformations, when applied without modification to polynomials with coefficients in $\{0, \pm 1\}$, return a collection of homogeneous, degree 4 polynomials whose coefficients lie in $\{0, \pm 1, \pm 2\}$. So this result follows from Schaefer's proof paired with Lemma 3.4. $\square$

We can also provide another strengthening of the **CommonZero** problem:

Theorem 3.6. *The **CommonZero** problem is $\exists\mathbb{R}$ -complete, even when the given polynomials $f_1, \dots, f_s \in \mathbb{Z}[x_1, \dots, x_m]$ have total degree at most 4, all coefficients are in $\{0, \pm 1, \pm 2\}$, and all common zeros, if any, are promised to lie in $[-1, 1]^m$.*

Proof. We reduce from H_2N , as strengthened in Theorem 3.5: we are given an instance $F = \{f_1, \dots, f_s\}$ of H_2N with degree 4 polynomials whose coefficients lie in $\{0, \pm 1, \pm 2\}$. Return the instance $F' = F \cup \{g\}$ of **CommonZero**, where $g(x_1, \dots, x_m) = x_1^2 + \dots + x_m^2 - 1$. Polynomial g guarantees that $Z(F') \subset \overline{B}(\vec{0}, 1) \subset [-1, 1]^m$, and all coefficients in F' are still in $\{0, \pm 1, \pm 2\}$, so F' has the required form.

If $Z(F)$ contains only $\vec{0}$, then $Z(F')$ is empty. On the other hand, if $\vec{a} \in Z(F) \setminus \{\vec{0}\}$, then $\vec{a}/|\vec{a}| \in Z(F')$, so the reduction is correct. $\square$

4 Globally Noncrossing Graphs and Linkages

Using the statement of the Main Theorem (Theorem 2.14) and the hardness results from Section 3.2, we show here that deciding realizability, rigidity, or global rigidity of a globally noncrossing graph/linkage is $\exists\mathbb{R}$ -complete or $\forall\mathbb{R}$ -complete, and that globally noncrossing linkages can draw any compact semialgebraic set in the plane.

4.1 Rigidifying Polygons

To simulate a rigidified tree, we will construct a globally rigid graph in the shape of a slight thickening of the tree. To that end, we provide here a general method that constructs a globally rigid triangulation, with Steiner points, of any simple polygon.

Lemma 4.1. *Any simple quadrilateral $A = A_1A_2A_3A_4$ has a globally rigid triangulation with four triangles and one Steiner point as in Figure 1a.*


Proof. If any of the four angles of A is 180° or greater, relabel so that this angle is at A_3 . This means any point in the interior of A within some distance $d > 0$ from A_1 is visible to all four of A 's vertices inside the quadrilateral.

Let P be the point on side A_1A_2 at distance $d/2$ from A_1 , so that the three angles $\angle A_2PA_3$, $\angle A_3PA_4$ and $\angle A_4PA_1$ lie strictly between 0° and 180° . By continuity, there is a point Q near P in the interior of A such that $\angle A_1QA_2$ is close to 180° and the three angles $\angle A_2QA_3$, $\angle A_3QA_4$ and $\angle A_4QA_1$ lie strictly between $180^\circ - \angle A_1QA_2$ and 180° . We may also assume Q is close to A_1 and is therefore visible to all four vertices of A .

We claim that the triangulation T using Steiner point Q is globally rigid. If not, then by Kawasaki's criterion [11] for flat-foldable single vertex crease patterns, it must be the case that some subset of the four angles at Q sum to 180° . But Q was chosen to ensure that this condition is false, since adding $\angle A_1QA_2$ to any of the other three angles results in more than 180° . $\square$

Lemma 4.2. *If A is a simple n -sided polygon, and points $P_1, \dots, P_m$ interior to A are specified, we may construct a triangulation of A with $2n + 6m - 5$ Steiner points that is globally rigid as a configured graph and has a Steiner point at each of $P_1, \dots, P_m$.*

Proof. Let T_1 be a triangulation of A with Steiner points at $P_1, \dots, P_m$; there are $n + 2m - 2$ triangles and $n + 3m - 3$ interior edges in this decomposition. Subdivide each triangle at its centroid to obtain triangulation T_2 , and delete the original $n + 3m - 3$ interior edges from T_2 to obtain a subdivision T_3 of polygon A into triangles and quadrilaterals, as illustrated in Figure 1b (for the special case $m = 0$). Finally, apply Lemma 4.1 to each quadrilateral in T_3 to obtain the final triangulation T . The three triangles or quadrilaterals meeting at each centroid are individually globally rigid, so the union of these three pieces is globally rigid again by Kawasaki's criterion: three creases around a vertex, no two collinear, are insufficient for a nontrivial single-vertex flat folding. Applying this reasoning around each centroid shows that the entire triangulation T is globally rigid, as desired. $\square$



(a) All points in the gray region are visible to the four vertices of quadrilateral A . Any point Q close enough to P renders this 4-triangle triangulation of A globally rigid.

(b) Rigidifying each quadrilateral in this decomposition using Lemma 4.1 rigidifies the entire graph, as shown in Lemma 4.2.

Figure 1: Rigidifying a simple quadrilateral (left) and a general simple polygon (right) by a triangulation with Steiner points, as in Lemmas 4.1 and 4.2.

This Lemma will be used when simulating rigidified subtrees in Section 4.3.

4.2 Rigidifying Polyominoes

When the partially rigidified tree (H, C_H) that we wish to simulate has *integer* coordinates, we will use a more refined grid which uses only rational coordinates and small rational edge lengths. To accomplish this, we show in this section that any polyomino, after scaling up by a factor of 1440, can be rigidified with only integer coordinates and constant-sized integer edge lengths. (A *polyomino* is a polygon with connected interior formed as the union of a finite set of squares in the standard unit-square tiling of $\mathbb{R}^2$.) We begin by rigidifying a single square:


Lemma 4.3. *The configured graph G_{cell} shown in Figure 2a, which has the shape of a 1440×1440 square with small *indents* on the edges, is globally rigid. The vertex coordinates and edge lengths are all integers.*

Proof. To specify G_{cell} graph in more detail, the labeled vertices have coordinates


$$\begin{aligned} A = (0, 0), \quad B = (720, 0), \quad C = (720, 720), \quad D = (0, 720), \quad E = (360, 224), \\ F = (615, 360), \quad G = (360, 496), \quad H = (105, 360), \quad I = (360, 825). \end{aligned}$$

In particular, pentagon $HFCID$ and quadrilateral $ABFH$ have a vertical line of symmetry, and AEB and DGC are reflections of each other through HF . The edge lengths are

$$AB = 720, \quad AE = 424, \quad AH = 375, \quad EF = 289, \quad GI = 329, \quad HF = 510.$$



(a) A globally rigid graph G_{cell} with integer coordinates and integer edge lengths in the shape of a 1440×1440 square with indents.



(b) Multiple copies of G_{cell} can be joined into a globally rigid polyomino.

Figure 2: Any polyomino made of 1440×1440 squares can be turned into a globally rigid graph with integer coordinates and integer edge lengths.

The rest of the coordinates may be computed by the 90° -degree rotational symmetry around C , and all distinct edge lengths are listed above. Note that A, E, F are not quite collinear, and similarly for H, G, C .

To show that G_{cell} is globally rigid, we again make repeated use of Kawasaki's criterion. First, the five triangles forming pentagon $HFCID$ (with Steiner point G) form a globally rigid subgraph: indeed it may be checked that $\angle HGF + \angle FGC = 180^\circ + \arcsin \frac{60}{901} > 180^\circ$ and $\angle HGF + \angle CGI = 180^\circ + \arcsin \frac{525}{15317} > 180^\circ$, so any subset of the five angles at G that includes $\angle HGF$ cannot add to exactly 180° . The four triangles forming quadrilateral $ABFH$ form a globally rigid subgraph as well, because the crease pattern around E is congruent to a subset of the crease pattern around G . The rotationally symmetric copies of pentagon $HFCID$ and $ABFG$ are likewise globally rigid. Finally, the quadrilateral and two pentagons meeting at F are globally rigid together because degree-3 crease patterns have no nontrivial flat foldings, and applying this reasoning four times around the square shows that all of G_{cell} is globally rigid. $\square$

Lemma 4.4. *Any polyomino P made of 1440×1440 squares can be triangulated (allowing Steiner points and edge subdivision) into a globally rigid triangulation that has integer coordinates and constant-sized integer edge lengths.*

Proof. Place a copy of G_{cell} in each cell of polyomino P , swapping orientation for every other cell so that adjacent cells have aligned indents. A short edge of length 210 within the indents—such as edge AB in Figure 2b—renders each pair of adjacent cells globally rigid. Because P is a polyomino which is connected via its edges, the whole assembly is thus globally rigid. Finally, indents along P 's boundary can be covered with edges of length 360, as in Figure 2b. $\square$

We show in the next section how to use this method to simulate integer-length

rigidified subtrees.

4.3 Simulation with Globally Noncrossing Linkages

Construction 4.5. Use notation as in Theorem 2.14: we are given a collection of polynomials $F = \{f_1, \dots, f_s\}$ in $2m$ variables, each of total degree at most d , from which Theorem 2.14 constructs a partially rigidified linkage $\mathcal{L} = \mathcal{L}(F)$.

Under Part I of Theorem 2.14 (i.e., with no additional assumptions about F), we may construct a globally noncrossing linkage $\mathcal{M} = \mathcal{M}(F)$ without constraints that perfectly simulates the constrained linkage $\mathcal{L}$.

If the Part II assumption holds, meaning the f_j have integer coefficients with absolute value at most M , then $\mathcal{M}$ may be constructed from $\mathcal{L}$ in time $O(\text{poly}(m^d, d^d, s, M))$, and furthermore, each edge length of $\mathcal{M}$ is rational, with size at most D and denominator dividing $28800D$.

Finally, if the Part III assumption also holds, i.e., each f_j satisfies $f_j(\vec{0}) = 0$ resulting in a configuration C_0 of $\mathcal{L}$, then the corresponding configuration of $\mathcal{M}$ has rational coordinates of magnitude $O(\text{poly}(m^d, d^d, s, M))$ with denominators dividing $28800D$.




Figure 3: Thickening a rigidified subtree (H, C_H) into a globally rigid polygon.

Proof. Consider one of the rigid constraints $\text{RigidCon}(H, C_H)$ on $\mathcal{L}$, and draw a polygon P that thickens (H, C_H) by distance $g := 1/(4D)$ in each direction, with angled wedges smaller than 60° at each leaf vertex, as shown in Figure 3a. Now apply the construction of Lemma 4.2 to this polygon P to obtain a globally rigid triangulation T that has a Steiner point at each interior vertex of C_H , and replace subgraph H of $\mathcal{L}$ with this triangulation T . Because tree (H, C_H) has at least three noncollinear vertices (Property 4), globally rigid triangulation T perfectly simulates the constraint $\text{RigidCon}(H, C_H)$: configurations of T are in bijection with configurations of rigidified tree (H, C_H) when considered in isolation. (The Steiner points chosen for T are there to satisfy the definition of *simulate*: T needs a vertex

corresponding to each vertex of H .) Let $\mathcal{M}$ be the linkage that results after performing this replacement for each rigid constraint on $\mathcal{L}$. One rigidified tree was originally pinned in $\mathcal{L}$ (by Property 6); keep these three pins in the corresponding locations in the triangulation T built from this tree. Because each tree is perfectly simulated, $\mathcal{M}$ perfectly simulates $\mathcal{L}$.

To see that $\mathcal{M}$ is globally noncrossing, note that Property 5 of Theorem 2.14 guarantees that the 60° wedges at leaf nodes do not intersect locally around their shared vertices, and because g is less than half the minimum feature size of any configuration of $\mathcal{L}$ (by Property 3), the thickened trees do not intersect elsewhere.

If Part II holds, we construct globally rigid graph T from each rigidified tree (H, C_H) more carefully, by building a rigidified polyomino with Lemma 4.4, as follows. Each edge length in $\mathcal{L}$ is an integer multiple of g , by Property 7. Thicken each rigidified tree (H, C_H) by distance g into a polygon P as above, where at each leaf of (H, C_H) , P forms a wedge with angle $2 \arcsin \frac{5}{13} < 60^\circ$. Now we can fill polygon P with a polyomino made from cells of side-length $u = g/5$, where the wedge at each leaf is attached with three edges of length $13u$, $11u$, and $13u$ respectively, as shown in Figure 2b. By rigidifying the polyomino with Lemma 4.4, the resulting graph T is a globally rigid thickening of tree (H, C_H) that has edge lengths in $\frac{1}{t}\mathbb{Z} \cap (0, \frac{1}{2}]$ and rational coordinates in $\frac{1}{t}\mathbb{Z}$ of magnitude at most $O(\text{poly}(m^d, d^d, s, M))$, where $t := 1440 \cdot 5 \cdot 4D = 28800D$. All edges of $\mathcal{L}$ not contained in a rigidified tree, which appear unmodified in $\mathcal{M}$, already have length at most D by Property 7, proving the Part II claim.

Finally, for Part III, first consider the nodes of $\mathcal{L}$ that are not interior to a rigidified subtree, i.e., are either a leaf of one or more rigidified trees or are not incident to any such tree (by Property 4). These nodes of $\mathcal{L}$ have corresponding nodes in $\mathcal{M}$ configured in the same locations, so these nodes of $\mathcal{M}$ are configured to have rational coordinates in $\frac{1}{D}\mathbb{Z}$ by Property 9. Pairing this fact with Property 11 implies that all remaining nodes of $\mathcal{M}$, which are interior to rigidified thickened trees built from polyominoes, are indeed configured to have rational coordinates in $\frac{1}{t}\mathbb{Z}$, proving the claim. $\square$

4.4 Hardness and Universality of Globally Noncrossing Linkages

We may finally prove the desired hardness and universality results about globally noncrossing graphs and linkages.

Theorem 4.6 (Hardness of Globally Noncrossing Realizability). *Deciding whether a given abstract weighted graph $\mathcal{G}$ is realizable, even when $\mathcal{G}$ is promised to be globally noncrossing and to have constant-sized integer edge lengths, is $\exists\mathbb{R}$ -complete.*

Proof. Membership in $\exists\mathbb{R}$ is evident. For hardness, we reduce from **CommonZero** as strengthened in Theorem 3.6. Given an instance $F = \{f_1, \dots, f_s\}$, which we may assume have $d = 4$ and $M = 2$, apply Construction 4.5 (Part II) to obtain a globally noncrossing linkage $\mathcal{M}$ whose edge lengths are rational with size at most D and denominators dividing $t := 28800D$. Let $\mathcal{G}$ be the weighted graph that results by unpinning $\mathcal{M}$'s three pins; the output of this reduction is the scaled graph $t \cdot \mathcal{G}$, whose edge lengths are *integers* bounded by $t \cdot D = O(1)$.

It remains to show why $t \cdot \mathcal{G}$ is realizable precisely when $Z(F)$ is nonempty. The linkage $\mathcal{M}$ from Construction 4.5 liftably and rigidly draws a translation of $Z(F)$ (because $Z(F) \subseteq [-1, 1]^{2m}$), so F has a common root if and only if $\text{Conf}(\mathcal{M})$ is nonempty, i.e., if and only if $\mathcal{M}$ is realizable. By Property 6, $\mathcal{M}$'s pins serve only to prevent rigid transformations and do not affect realizability, so $\mathcal{M}$ is realizable precisely when $\mathcal{G}$ is realizable. Finally, scaling by t does not affect realizability, proving the result. $\square$

Theorem 4.7 (Hardness of Noncrossing Rigidity and Global Rigidity). *Deciding whether a given configured weighted graph $(\mathcal{G}, C_0)$ is rigid, when $\mathcal{G}$ is promised to be globally noncrossing (so in particular, C_0 is noncrossing) and C_0 has integer coordinates and constant-sized integer edge lengths, is $\forall\mathbb{R}$ -complete. It remains $\forall\mathbb{R}$ -complete if “rigid” is replaced by “globally rigid”.*

Proof. Schaefer [22] has shown that the general linkage rigidity problem is $\forall\mathbb{R}$ complete and therefore belongs to $\forall\mathbb{R}$. Linkage global rigidity likewise belongs to $\forall\mathbb{R}$: it may be expressed in the form *for all valid configurations C , $|C(u) - C(v)| = |C_0(u) - C_0(v)|$ for all pairs of (not necessarily adjacent) vertices u and v .*

For hardness, we reduce from H_2N as strengthened in Theorem 3.5, so suppose we are given a family of homogeneous polynomials $F = \{f_1, \dots, f_s\}$ of degree $d = 4$ in variables $\{x_1, y_1, \dots, x_m, y_m\}$ with constant-sized integer coefficients. We may use Construction 4.5 (Part III) to build a globally noncrossing configured linkage $(\mathcal{M}, C_0)$ that liftably and rigidly draws a trace $\pi_X(\text{Conf}(\mathcal{M}))$ satisfying $T(Z(F) \cap [-1, 1]^{2m}) \subseteq \pi_X(\text{Conf}(\mathcal{M})) \subseteq T(Z(F))$ for some translation T . The result of this reduction will be the configured graph $\mathcal{G}$ formed by scaling $(\mathcal{M}, C_0)$ by $28800D$ and removing the three pins. Configured graph $\mathcal{G}$ indeed has polynomially-bounded integer coordinates and constant integer lengths.

To verify the validity of this reduction, suppose first that $Z(F)$ contains some nonzero point $\vec{a}$. Then $Z(F)$ contains the entire path $p \mapsto p \cdot \vec{a}$ starting at $\vec{0}$, so $\vec{0}$ is not isolated in $Z(F)$, i.e., $T(\vec{0})$ is not isolated in $\pi_X(\text{Conf}(\mathcal{M}))$. Because $\mathcal{M}$ draws liftably, $(\mathcal{M}, C_0)$ is not rigid (as a linkage), and therefore $\mathcal{G}$ is not rigid (as a graph). On the other hand, if $Z(F) = \{\vec{0}\}$ then $\pi_X(\text{Conf}(\mathcal{M}))$ contains only the single point $\pi_X(C_0) = T(\vec{0})$, and by Property 10 of Theorem 2.14 (uniqueness), it follows that $\text{Conf}(\mathcal{M}) = \{C_0\}$, i.e., $\mathcal{M}$ is both rigid and globally rigid as a linkage. By Property 6, all configurations of $\mathcal{G}$ are Euclidean transformations of configurations of $28800D \cdot \mathcal{M}$, and so $\mathcal{G}$ is rigid and globally rigid as a graph. $\square$

Theorem 4.8 (Universality of Globally Noncrossing Linkages). *The proper subsets $R \subsetneq \mathbb{R}^2$ that may be drawn by a globally noncrossing linkage are precisely the compact semialgebraic sets.*

Proof. Given any unconstrained linkage $\mathcal{L}$ and vertex v , the trace of v is either all of $\mathbb{R}^2$ (if $\mathcal{L}$ has no pins) or is the projection of compact algebraic set $\text{Conf}(\mathcal{L})$ and is therefore compact and semialgebraic. (Recall that $\mathcal{L}$ is assumed to be connected.)

Conversely, by Lemma 3.3, any compact semialgebraic region $R \subset \mathbb{R}^2$ may be written as the projection of some basic compact set $R' = Z(f_1, \dots, f_s) \subset \mathbb{R}^{2m}$ onto the first two variables, so it suffices (by ignoring all drawing vertices except the first) to show that some

translation of R' may be drawn by a globally noncrossing linkage. By scaling as necessary, we may further assume that the compact set R' lies in the box $[-1, 1]^{2m}$. But this now follows directly from Construction 4.5, Part I. $\square$

5 Unit-Distance and $\{1, 2\}$ -Distance Graphs and Linkages

Definition 5.1 (Unit Distance Graphs/Linkages). Define an *abstract unit-distance graph (or linkage)* as an abstract weighted graph (or abstract linkage) where all edges have weight 1; a *configured unit-distance graph (or linkage)* additionally comes with such a configuration.

With our terminology, an *abstract unit distance graph/linkage* does not necessarily have any valid configurations, in contrast to the more common usage of the term “unit-distance graph”. To mitigate confusion with our overloading of this term, we will always refer to a unit-distance graph/linkage as “abstract” or “configured”.

In this section we prove the strong $\exists\mathbb{R}$ -completeness or $\forall\mathbb{R}$ -completeness of realizability, rigidity, and global rigidity for unit-distance (or in the case of global rigidity, $\{1, 2\}$ -distance) graphs that allow crossings. We also show universality: any compact semialgebraic set in $\mathbb{R}^2$ can be drawn by a unit-distance linkage. (Unit-distance graphs that do *not* allow crossings, i.e., *matchstick graphs*, are the topic of Section 6.)

There are two noteworthy obstacles in these arguments that were not present in the previous section. First, the universality proof involves a new complication, namely, non-algebraic numbers. To illustrate, the circle $C = \{(x, y) \mid x^2 + y^2 = e^2\}$ (where e is Euler’s constant) can be drawn easily by a linkage (using a single edge of length e), but simulating such an edge with a unit-distance graph is impossible because e is transcendental. As a workaround, we instead rely on *pins* to introduce non-algebraic values. Indeed, we may slightly generalize curve C by introducing new variables (a, b) and considering the modified curve

$$C' = \{((x, y), (a, b)) \in \mathbb{R}^4 \mid x^2 + y^2 = a^2\}.$$

As C' is now defined by polynomials with *integer* coefficients, the Main Theorem (Part II) applies, and the resulting linkage may be simulated by a unit distance linkage using techniques to be presented below. Finally, with one pin, we may fix the values $a = e$ and $b = 0$, which recovers the desired circle C . Suitably generalized, this argument can be made to work for arbitrary compact semialgebraic sets; see Theorem 5.9 for details.

For the second obstacle, we were not able to prove $\forall\mathbb{R}$ -completeness of detecting global rigidity of unit-distance graphs. Indeed, we are not aware of the existence of *any* globally rigid unit-distance graphs larger than a triangle!

Open Question. *Are there any globally rigid unit-distance graphs with more than 3 edges?*

If such a graph is found, it is likely that the methods of this paper can turn it into a proof of hardness.

As a consolation prize, we demonstrate $\forall\mathbb{R}$ -completeness of global rigidity for graphs with edge lengths in $\{1, 2\}$, an appropriate strengthening of Saxe’s result [21] that global rigidity is coNP-hard for graphs with edge lengths in $\{1, 2\}$.

5.1 Simulation with Unit- and $\{1, 2\}$ -Distance Linkages

Here we show how to simulate Theorem 2.14 using $\{1, 2\}$ -distance linkages.

Lemma 5.2 (Reinforced Segment). *A single edge of integer length n is perfectly simulated by a reinforced bar graph formed by adjoining $n - 1$ degenerate $\{1, 1, 2\}$ -sided triangles along unit edges.*

Proof. This is a simple extension of a tool used in [21, Cor. 4.3]. $\square$

To rigidify orthogonal trees with $\{1, 2\}$ -graphs, it suffices to rigidify entire lattice grids:

Lemma 5.3 (Reinforced Grid). *Let (G, C) be the configured graph whose vertices lie at all integer points in $[0, n] \times [0, n]$ and whose unit-length edges connect vertically and horizontally adjacent vertices in this grid. Then the rigidified graph $\mathcal{G} := ((G, C), \text{RigidCon}_G(G, C))$, where configuration C is rigidified in its entirety, can be perfectly simulated by an unconstrained $\{1, 2\}$ -distance graph, called a reinforced grid.*

Proof. We may assume $n \geq 4$; otherwise, we apply the construction for $n = 4$ and restrict attention to the smaller subgrid. For each $0 \leq j, k \leq n$, let $v_{j,k}$ be the vertex of G with $C(v_{j,k}) = (j, k)$. By Lemma 5.2, we may add length-2 edges $v_{j,k}v_{j+2,k}$ and $v_{j,k}v_{j,k+2}$ to force each row and column of vertices in G to remain straight. Now add one more rigidified bar of length 5 connecting $v_{4,0}$ and $v_{0,3}$, which constrains row $k = 0$ and row $j = 0$ to remain at 90° from each other. In fact, this resulting graph G' is the desired graph. Indeed, suppose we have a configuration of G' ; by a Euclidean motion, we may assume $v_{0,0}$, $v_{n,0}$, and $v_{0,n}$ are configured at $(0, 0)$, $(n, 0)$, and $(0, n)$ respectively. Because $|v_{n,0} - v_{n,n}| = n$ and $|v_{0,n} - v_{n,n}| = n$ in any configuration, $v_{n,n}$ must rest at (n, n) or $(0, 0)$. In the latter case, $v_{n,1}$ rests at $(n - 1, 0)$, which is not distance n away from $v_{0,1}$ at $(0, 1)$, contradicting Lemma 5.2. So $v_{n,n}$ must indeed lie at (n, n) , and the rest of the vertices' locations are then fixed. $\square$

Construction 5.4. *Use notation as in Theorem 2.14: we are given a collection of polynomials $F = \{f_1, \dots, f_s\}$ in $2m$ variables, each of total degree at most d , from which Theorem 2.14 constructs a partially rigidified linkage $\mathcal{L} = \mathcal{L}(F)$. We make no claims under Part I alone.*

If the Part II assumption holds, meaning the f_j have integer coefficients with absolute value at most M , we may construct, in $O(\text{poly}(m^d, d^d, s, M))$ time, an abstract linkage $\mathcal{M} = \mathcal{M}(F)$ with edge lengths in $\{1, 2\}$ that perfectly simulates the scaled linkage $D \cdot \mathcal{L}(F)$. In particular, there is a translation T on $\mathbb{R}^{2m}$ and a subset X of m vertices such that

$$T(D \cdot (Z(F) \cap [-1, 1]^{2m})) \subseteq \pi_X(\text{Conf}(\mathcal{M})) \subseteq T(D \cdot Z(F)).$$

If the Part III assumption also holds, i.e., each f_j satisfies $f_j(\vec{0}) = 0$ which gives rise to a configuration C_0 of $\mathcal{L}$, then the configuration of $\mathcal{M}$ corresponding to $C_0 \in \text{Conf}(\mathcal{L})$ has rational coordinates whose numerators have magnitude $O(\text{poly}(m^d, d^d, s, M))$ and whose denominators are at most D^2 .

Proof. By Property 9, all edge lengths of $D \cdot \mathcal{L}$ are integers. By Properties 2, 4 and 7, each rigidified subtree (H, C_H) in $D \cdot \mathcal{L}$ has integer coordinates whose sizes are bounded by $O(\text{poly}(m^d, d^d, s, M))$.

Each edge of $D \cdot \mathcal{L}$ not belonging to any rigidified tree gets replaced with a reinforced segment of appropriate length as in Lemma 5.2. For each rigidified tree (H, C_H) of $D \cdot \mathcal{L}$, build a reinforced grid of unit squares as in Lemma 5.3 large enough to include the coordinates of the vertices of (H, C_H) , and replace the tree by this grid; neighboring edges or trees are attached at the corresponding grid point. The three pins of $D \cdot \mathcal{L}$ are likewise transferred to their corresponding grid points. Call the resulting linkage $\mathcal{M} = \mathcal{M}(F)$. Each reinforced segment of $\mathcal{M}$ perfectly simulates its edge by Lemma 5.2. Likewise, because each rigidified tree in $D \cdot \mathcal{L}$ has at least three noncollinear vertices (Property 4 of the Main Theorem) each reinforced grid perfectly simulates its rigidified tree by Lemma 5.3. So $\mathcal{M}$ perfectly simulates $\mathcal{L}$.

If Part III holds, then scaled configuration $D \cdot C_0$ of $D \cdot \mathcal{L}$ has integer coordinates bounded by $O(\text{poly}(m^d, d^d, s, M))$ in magnitude. By Property 11 and Lemma 5.3 (Reinforced Grid), all nodes of $\mathcal{M}$ in reinforced grids are configured with coordinates in $\frac{1}{5} \cdot \mathbb{Z}$. It remains to look at the coordinates of the remaining nodes: those on reinforced segments. Consider an edge e of $D \cdot \mathcal{L}$ *not* contained in a rigidified tree. Edge e has integer length $q \leq D^2$, and its endpoints are initially configured at integer coordinates (a_1, b_1) and (a_2, b_2) (by Property 9). Then the reinforced bar corresponding to e in $\mathcal{M}$ has nodes configured at $(a_1, b_1) + \frac{h}{q}(a_2 - a_1, b_2 - b_1)$ for integers $0 \leq h \leq q$, and these coordinates are rationals of the required form. $\square$

We rely on the full strength of *perfect* simulation in the proof of global rigidity below, but for the other hardness results, liftable and rigid simulation is sufficient. The latter may be achieved with only unit-length edges:

Lemma 5.5. *A single edge of length 2 can be liftable and rigidly simulated by a unit-distance graph with 19 edges, formed by joining two copies of Moser’s Spindle along a common equilateral triangle. If the edge has its endpoints configured at $(0, 0)$ and $(2, 0)$, then the configured unit-distance graph simulating this edge has coordinates of the form $a + b\sqrt{3} + c\sqrt{11} + d\sqrt{33}$ for rational numbers a, b, c, d .*

Proof. The first claim follows from [22, Lemma 3.4], while the second may be verified by direct computation. $\square$

5.2 Hardness and Universality of Unit-Distance and $\{1, 2\}$ -Distance Graphs

Theorem 5.6 (Hardness of Unit-Distance Realization). *Deciding whether a given abstract unit-distance graph is realizable is $\exists\mathbb{R}$ -complete.*

Note. This was shown by Schaefer [22] with a simpler, specialized construction, but we include it here for completeness.

Proof. Membership in $\exists\mathbb{R}$ is evident. Hardness follows by reduction from **CommonZero** exactly as in the proof of Theorem 4.6, using Construction 5.4 in place of Construction 4.5. $\square$

Theorem 5.7 (Hardness of Unit-Distance Rigidity). *The problem of determining whether a configured unit-distance graph with coordinates in $\mathbb{Q}[\sqrt{3}, \sqrt{11}]$ is rigid is $\forall\mathbb{R}$ -complete, even when all coordinates have the form $(a + b\sqrt{3} + c\sqrt{11} + d\sqrt{33})/n$ for integers a, b, c, d, n where b, c, d , and n have size $O(1)$ (but a need not have constant size).*

Proof. As in Theorem 4.7, this problem lies in $\exists\mathbb{R}$. Hardness follows by reduction from the complement of $\mathbf{H}_2\mathbf{N}$: given an instance $F = \{f_1, \dots, f_s\}$ of this problem (which we may assume consists of polynomials of degree 4 with constant-sized coefficients by Theorem 3.5), use Construction 5.4 and Lemma 5.5 to build a configured unit-distance linkage $\mathcal{M}$ that liftably and rigidly draws a translation of $D \cdot (Z(F) \cap [-1, 1]^{2m})$. Then this linkage is rigid if and only if $\vec{0}$ is the only common zero of F , i.e., F is a “no” instance of $\mathbf{H}_2\mathbf{N}$. As in the proof of Theorem 4.7, removing the three pins of $\mathcal{M}$ results in a unit-distance graph that is rigid if and only if F is a “no” instance of $\mathbf{H}_2\mathbf{N}$.

We must also show that the coordinates of $\mathcal{M}$ have the required form. Construction 5.4 guarantees that Lemma 5.5 is applied only to length-2 edges whose endpoints have rational coordinates with denominators bounded by D^2 . For each such edge, say with endpoints at $p = (p_1, p_2) \in \mathbb{Q}^2$ and $q = (q_1, q_2) \in \mathbb{Q}^2$, the gadget drawn to connect p and q may be computed by starting with the gadget of Lemma 5.5 (which connects $(0, 0)$ to $(2, 0)$), applying the rotation matrix

$$\frac{1}{2} \begin{pmatrix} q_1 - p_1 & -(q_2 - p_2) \\ q_2 - p_2 & q_1 - p_1 \end{pmatrix},$$

and then translating by (p_1, p_2) . The entries of the rotation matrix are rationals with denominators at most $2D^2$ and magnitudes at most 1 (because $|q - p| = 2$), while p_1 and p_2 have denominators bounded by D^2 , so the result follows. $\square$

Theorem 5.8 (Hardness of $\{1, 2\}$ -Distance Global Rigidity). *The problem of deciding whether a given configured $\{1, 2\}$ -distance graph with coordinates in $\mathbb{Q}$ is globally rigid is $\forall\mathbb{R}$ -complete, even when all coordinates have denominators of size $O(1)$.*

Proof. Membership in $\forall\mathbb{R}$ follows as in Theorem 4.7. Hardness follows by reduction from the complement of $\mathbf{H}_2\mathbf{N}$ just as in the proof of Theorem 4.7, using Construction 5.4 instead of Construction 4.5. As in the proof of Theorem 4.7, this makes essential use of the fact that $\mathcal{M}$ from Construction 5.4 simulates the linkage $\mathcal{L}$ from the Main Theorem *perfectly*, not just liftably and rigidly. $\square$

As discussed at the start of Section 5, Construction 5.4 requires the input polynomials to have integer coefficients, but some compact semialgebraic sets cannot be expressed in this way. We now formalize the workaround described there to prove universality of unit-distance linkages.

Theorem 5.9 (Universality of Unit-Distance Linkages). *Any compact semialgebraic set $R \subset \mathbb{R}^2$ may be drawn by a unit-distance linkage.*

Proof. As in the proof of Theorem 4.8 we may write R as the projection onto coordinates x_1, y_1 of some compact basic algebraic set

$$R' = Z(f_1, \dots, f_s), f_j \in \mathbb{R}[x_1, y_1, \dots, x_m, y_m],$$

and it suffices to show that some translation of R' may be drawn with a unit-distance linkage. In fact, it suffices to show that the scaled set $\frac{1}{n} \cdot R'$ may be drawn (up to translation) by a unit-distance linkage $\mathcal{L}$, for some $n \in \mathbb{N}$: indeed, if a unit-distance linkage $\mathcal{L}$ draws a translation of $\frac{1}{n} \cdot R'$, then $n \cdot \mathcal{L}$ draws a translation of R' and has integer edge lengths, so by Lemmas 5.2 and 5.5 it may be simulated by a unit-distance linkage, as required. By this reasoning, we may replace the compact set R' by some small-enough $\frac{1}{n} \cdot R'$ and thereby assume that R' lies in the box $[-1, 1]^{2m}$.

Write the scaled set $\frac{1}{D} \cdot R'$ as $Z(g_1, \dots, g_s)$, where $g_j(\vec{xy}) = f_j(D \cdot \vec{xy})$. We wish to apply Construction 5.4 to polynomials $g_1, \dots, g_s$ so the resulting linkage draws precisely R' (up to translation), but these coefficients may not be integers (or even algebraic numbers, as described above), so we will temporarily replace these coefficients with variables, as follows. By scaling the coefficients of each g_j (which does not affect $Z(g_1, \dots, g_s)$), we may assume that the coefficients are in $[-1, 1]$. For each nonzero monomial $c_{j,J} \vec{xy}^J$ in each g_j (where J is a vector of exponents), create new variables $a_{j,J}$ and $b_{j,J}$, gather all of these new variables into a vector $\vec{ab}$ with length $2r$, and define the new polynomials

$$h_j(\vec{xy}, \vec{ab}) := \sum_{J \text{ such that } c_{j,J} \neq 0} a_{j,J} \vec{xy}^J,$$

for $1 \leq j \leq s$. Polynomials h_j now have integer coefficients: in fact, all coefficients equal 1. It remains to implement the equations $h_j(\vec{xy}, \vec{ab}) = 0$ and $(a_{j,J}, b_{j,J}) = (c_{j,J}, 0)$, for each $1 \leq j \leq s$ and $c_{j,J} \neq 0$, which exactly recover the solutions $\vec{xy} \in Z(g_1, \dots, g_s)$.

Applying Construction 5.4 and Lemma 5.5 to the polynomials h_j , we may construct some unit-distance linkage $\mathcal{L}$, a set of drawing vertices

$$X = \{v_j \mid 1 \leq j \leq s\} \cup \{v_{j,J} \mid 1 \leq j \leq s \text{ and } c_{j,J} \neq 0\}$$

corresponding to variables (x_j, y_j) and $(a_{j,J}, b_{j,J})$ respectively, and some translation T on $\mathbb{R}^{2m+2r}$ such that X draws a set between $T(D \cdot (Z(h_1, \dots, h_s) \cap [-1, 1]^{2m+2r}))$ and $T(D \cdot Z(h_1, \dots, h_s))$. Finally, we pin all the vertices $v_{j,J}$ in the plane to force variables $(a_{j,J}, b_{j,J})$ to take the values $(c_{j,J}, 0)$: specifically, if $T_{j,J}$ denotes the translation T restricted to coordinates $(a_{j,J}, b_{j,J})$, we pin $v_{j,J}$ to the point $T_{j,J}(c_{j,J}, 0)$. The trace of vertices $v_1, \dots, v_s$ in this pinned linkage is a translation of $D \cdot Z(g_1, \dots, g_s) = R'$, as required. $\square$

6 Matchstick Graphs and Linkages

In Section 4 we discussed globally noncrossing linkages, which have no crossing configurations because they are required to be carefully designed to enforce this stringent property.

By contrast, in this section we look at *NX-constrained* linkages (short for non-crossing-constrained), which have no crossing configurations for a very different reason: we simply declare that crossing configurations should be ignored. *Any* linkage can be made into an NX-constrained linkage by simply attaching a constraint (cf. Definition 2.7), specifically an *NX-constraint*, that redefines the set of valid configurations as follows:

Definition 6.1 (NX-Constrained Linkages and Matchstick Linkages). If $\mathcal{L}$ is a linkage, we define the *NX-constraint* on $\mathcal{L}$, denoted $\text{NX}_{\mathcal{L}}$, as the set of configurations of $\mathcal{L}$ that do not cross. In other words, if $\mathcal{L}' := (\mathcal{L}, \text{NX}_{\mathcal{L}})$ is an *NX-constrained linkage*, then its configuration space is defined by $\text{Conf}(\mathcal{L}') := \text{NXConf}(\mathcal{L})$.

If $\mathcal{L}$ is an unconstrained linkage all of whose edges have length 1, then the NX-constrained linkage $\mathcal{L}' := (\mathcal{L}, \text{NX}_{\mathcal{L}})$ is called a *matchstick linkage*.

In this section we prove analogous hardness results about matchstick linkages and graphs: matchstick graph realization is $\exists\mathbb{R}$ -complete, and matchstick graph rigidity and global rigidity are $\forall\mathbb{R}$ -complete. An analogous “universality” result—that matchstick linkages can draw all compact semialgebraic sets—can also be proved with arguments similar to the prior universality results. However, such a theorem would be incomplete, because matchstick linkages (more generally, NX-constrained linkages) can draw more than just compact semialgebraic sets! As a simple example, the NX-constrained two-bar linkage $\mathcal{A}$ of Figure 4 draws the half-open annulus $\{(x, y) \in \mathbb{R}^2 \mid 1 < x^2 + y^2 \leq 9\}$, because configurations of (the unconstrained linkage underlying) $\mathcal{A}$ with v on circle $x^2 + y^2 = 1$ are crossing and are thus considered invalid by the constraint. In general, the traces of NX-constrained linkages are *bounded and semialgebraic* sets (assuming at least one pin). NX-constrained linkages—in fact, matchstick linkages—can indeed draw all such sets, but our proof of this stronger result (Theorem 7.25) subtly breaks the abstraction barrier set up by the Main Theorem, so we postpone this proof until Section 7.9.




Figure 4: The trace of an NX-constrained linkage need not be closed.

6.1 Simulation with Matchstick Linkages

For an integer $a \geq 5$, we will simulate length- a edges by *edge polyiamonds* as shown in Figure 5. Note that any polyiamond, considered as a matchstick graph, is globally rigid. As proven in the following lemma, two simple but effective *wing edges* suffice to fix the relative orientation and position of two adjacent edge polyiamonds (Figure 5). We may more easily

describe the *relative* positions of edge polyamonds P and Q by temporarily pinning P in place:

Lemma 6.2 (Wing Edges). *Consider the matchstick linkage $\mathcal{L}$ drawn in Figure 5, with two edge polyamonds P and Q sharing a vertex, three pins in P , and two extra wing edges attached at vertices a and b as shown. Then in every configuration of this linkage, Q has the same orientation as P , and its central axis is rotated from P 's axis by an angle $60^\circ < \theta < 240^\circ$. Every such θ corresponds to a unique configuration of $\mathcal{L}$.*

Proof. Polyamond Q must be drawn with the same orientation as P , or else the wing edges will create crossings. Likewise, a , o , b , and c must form a rhombus to prevent these four edges from intersecting each other. So Q 's and c 's location are determined by the single angle θ , and it may be seen that no crossings occur precisely when $60^\circ < \theta < 240^\circ$. $\square$




Figure 5: Left: Edge polyamonds used to simulate edges of integer length. Right: Edge polyamonds braced at 90° based on a 5-12-13 right triangle.

Lemma 6.3 (Orthogonal Braces). *By modifying the edge polyamonds and adding a hypotenuse polyamond as shown in Figure 5, the resulting matchstick linkage $\mathcal{M}$ is globally rigid, and its unique configuration has $\theta = 90^\circ$.*

Proof. The previous lemma and the 5-12-13 right triangle force Q to be drawn at a $+90^\circ$ angle from P and with the same orientation. The hypotenuse polyamond then has only one crossing-free position. $\square$

Construction 6.4. *Use notation as in Theorem 2.14: we are given a collection of polynomials $F = \{f_1, \dots, f_s\}$ in $2m$ variables, each of total degree at most d , from which Theorem 2.14 constructs a partially rigidified linkage $\mathcal{L} = \mathcal{L}(F)$.*

If the hypotheses of Part II hold, so polynomials f_j have integer coordinates with absolute values at most M , we may construct, in $O(\text{poly}(m^d, d^d, s, M))$ time, an abstract matchstick linkage $\mathcal{M}$ that perfectly simulates the scaled linkage $D \cdot \mathcal{L}$.

If the Part III assumption also holds, i.e., each f_j satisfies $f_j(\vec{0}) = 0$ giving rise to configuration C_0 of $\mathcal{L}$, then the corresponding configuration of $\mathcal{M}$ has coordinates of the

form $(a + b\sqrt{3})/c$, where a is an integer of magnitude $O(\text{poly}(m^d, d^d, s, M))$, and b and c are integers with magnitude at most $O(1)$.

Proof. We may assume that in each rigidified subtree (H, C_H) of $\mathcal{L}$, no internal node has degree exactly 2 with edges at 180° from each other in C_H , because such a node can be erased by merging its two edges together. For the remainder of this proof, replace $\mathcal{L}$ with $40D \cdot \mathcal{L}$: each rigidified tree has integer coordinates, each edge has integer length at least 40, and each configuration has feature size at least 40, by Theorem 2.14.

We create a matchstick linkage $\mathcal{M}$ simulating $\mathcal{L}$ using Lemmas 6.2 and 6.3, as follows. First, $\mathcal{M}$ has a vertex corresponding to each vertex of $\mathcal{L}$. Each edge of $\mathcal{L}$ is replaced by an edge polyiamond of the appropriate length connecting the corresponding vertices.

At any vertex v internal to some (necessarily unique, by Property 4) rigidified tree (H, C_H) , brace any right angles at v in (H, C_H) as in Lemma 6.3, by modifying the corresponding edge polyiamonds and adding wings and a hypotenuse. Any 180° or 270° angles at v in (H, C_H) can be left alone. At any vertex v of $\mathcal{L}$ that is *not* internal to any rigidified tree (and has degree at least 2), simply attach wings to the incident edge polyiamonds as in Lemma 6.2 according to the cyclic order σ_v at v . Finally, $\mathcal{L}$ has exactly three pins in one of the rigidified trees (H, C_H) ; in $\mathcal{M}$, we instead place three non-collinear pins in the edge polyiamond corresponding to *one* of H 's edges.

We must show that $\mathcal{M}$ perfectly simulates $\mathcal{L}$. First, we claim no edge polyiamond may reverse its orientation: they all have the same orientation as in Figure 5 in every configuration of $\mathcal{M}$. This is certainly true for the pinned edge polyiamond. Lemma 6.2 shows that edge polyiamonds connected by wing edges must maintain their relative orientations. At each vertex v of $\mathcal{L}$, *all* of the edge polyiamonds of $\mathcal{M}$ incident to v are connected to each other by wing edges: if v is internal to a rigidified tree then at most one of its corners is missing wing edges (by the assumption above), and otherwise all of v 's corners have wing edges. So by connectivity, all edge polyiamonds maintain the same orientation, as claimed. This shows that each configuration C of $\mathcal{L}$ induces *at most* one configuration of $\mathcal{M}$: the edge polyiamond built for edge uv must be rotated and translated (not reflected) to connect points $C(u)$ and $C(v)$ in the plane.

Similarly, each vertex interior to a rigidified tree (H, C_H) has hypotenuse polyiamonds on all or all but one of its corners, which enforce the constraint $\text{RigidCon}(H, C_H)$ by Lemma 6.3. So each configuration of $\mathcal{M}$ comes from a configuration of $\mathcal{L}$, i.e., $\mathcal{M}$ draws a *subset* of $\text{Conf}(\mathcal{L})$.

It remains to show that each configuration C of $\mathcal{L}$ induces a valid (i.e., noncrossing) configuration of $\mathcal{M}$. For each rigidified tree (H, C_H) , the edge and hypotenuse polyiamonds simulating (H, C_H) in $\mathcal{M}$ may be configured to match C without intersecting each other because the induced configuration $C|_H$ has the same orientation as C_H , by Property 4 of the Main Theorem. For each vertex v of $\mathcal{L}$ not interior to a rigidified tree, the angles of v 's corners in configuration C lie strictly between 60° and 240° by Property 5 of the Main Theorem, so locally, wing edges around v do not create crossings by Lemma 6.2. Globally, configuration C of $\mathcal{L}$ has feature size at least 40, while the linkage $\mathcal{M}$ extends less than 20 units away from the vertices and edges it simulates, so C indeed induces a (noncrossing)

configuration of $\mathcal{M}$, as desired.

Finally, if Part III applies, let C'_0 be the configuration of $\mathcal{M}$ induced by initial configuration C_0 (after scaling C_0 by $40D$ as above), and note that (the scaled) C_0 has integer coordinates and integer edge lengths, both of polynomial magnitude. We must show that the coordinates of C'_0 have the required form $(a + b\sqrt{3})/c$ for integers a, b, c bounded as claimed. For any edge e of C_0 with integer endpoints (a_1, b_1) and (a_2, b_2) and integer length q , the vertices along the central axis of the corresponding edge polyiamond have rational coordinates of the form $p = (a_1, b_1) + \frac{h}{q}(a_2 - a_1, b_2 - b_1)$ for integers $0 \leq h \leq q$. If e is part of a rigidified tree then this edge is axis-aligned (by Property 11) and so p has integer coordinates; otherwise, e has length $q \leq 40D^2$ by Property 7, so the denominators of p have magnitude $q = O(1)$. The rest of the vertices in the polyiamond are offset from these coordinates by the unit vector $\vec{w} = (a_2 - a_1, b_2 - b_1)/q$ rotated by some multiple of 60° ; for the same reasons, denominators in w 's entries have constant size, so the same holds for the vertices in the edge polyiamond. The same computations holds for hypotenuse polyiamonds. The only vertices not yet accounted for are the wing vertices, and these have the form $A + B - O$ where A, B , and O are rational-coordinate points along edge-polyiamond axes as described above. $\square$

Note. A single edge on its own is *not* perfectly simulated by an edge polyiamond, because the latter may reflect across its central axis. In the proof above, we were careful to argue that all edge polyiamonds maintain a fixed orientation (due to pins and wing edges) so that this ambiguity is impossible.

6.2 Hardness of Matchstick Linkages

Theorem 6.5 (Hardness of Matchstick Graph Realizability). *Deciding whether a given abstract matchstick graph is realizable is $\forall\mathbb{R}$ -complete.*

Theorem 6.6 (Hardness of Matchstick Graph Rigidity). *Deciding whether a given configured matchstick graph with coordinates in $\mathbb{Q}[\sqrt{3}]$ is rigid is $\forall\mathbb{R}$ -complete, even when all coordinates have the form $(a + b\sqrt{3})/c$ for integers a, b, c where b and c have size $O(1)$.*

Theorem 6.7 (Hardness of Matchstick Graph Global Rigidity). *Deciding whether a given configured matchstick graph with coordinates in $\mathbb{Q}[\sqrt{3}]$ is globally rigid is $\forall\mathbb{R}$ -complete, even when all coordinates have the form $(a + b\sqrt{3})/c$ for integers a, b, c where b and c have size $O(1)$.*

These hardness proofs are perfectly analogous to those in Section 5, using Construction 6.4 instead of Construction 5.4 and Lemma 5.5, so we omit their proofs.

As described at the beginning of this section, *universality* of drawing with matchstick linkages is more subtle than prior universality results, because matchstick linkages can draw traces that are bounded semialgebraic sets which need not be closed. We will prove that matchstick linkages can indeed draw *all* such sets, but our argument utilizes details from the *proof* of the Main Theorem. The universality statement and proof may therefore be found at Theorem 7.25, at the end of Section 7.

7 Extended Linkages and the Main Construction

In this section we finally prove the Main Theorem. We first define *extended linkages*, a special type of constrained linkage designed specifically for this proof, in Section 7.1. Section 7.2 then presents a detailed outline of our proof strategy, and Section 7.3 details our choices of parameters. The remainder of this section then covers the proof in full.

7.1 Defining Extended Linkages

We define and use *extended linkages*, which are constrained linkages whose constraints are tailored for the specifics of our construction. (More general constraints are possible but not explored in this paper.)

The first of these constraints, the angle constraint, specifies a preferred arrangement of edges around each vertex, and specifies that the angle at each corner may not change very much. Recall that a *combinatorial embedding* consists of a cyclic counterclockwise ordering σ_v of the edges incident to each vertex v (Definition 2.6).

Definition 7.1 (Angle Constraint). For a linkage $\mathcal{L}$ and a combinatorial embedding σ of $\mathcal{L}$, an *angle constraint*, $\text{AngleCon}_{\mathcal{L}}(\sigma, A, \Delta)$, is specified by an assignment of an angle $0 \leq A(\Lambda) \leq 2\pi$ and an angle tolerance $\Delta(\Lambda) \geq 0$ to each corner Λ of σ , with the condition that A assigns a total of 2π to the corners around each vertex.

A configuration $C \in \text{Conf}(\mathcal{L})$ satisfies the angle constraint $\text{AngleCon}_{\mathcal{L}}(A, \Delta)$ if, for each corner Λ , angle $\angle C(\Lambda)$ lies in the closed interval

$$[A(\Lambda) - \Delta(\Lambda), A(\Lambda) + \Delta(\Lambda)].$$

In particular, any corner with $\Delta(\Lambda) = 0$ is rigid: its angle in C must be exactly $A(\Lambda)$. We call such corners *frozen*.

Note that edges incident to a vertex v might not be configured in the cyclic order that σ prefers, if tolerances $\Delta(\Lambda)$ are large enough to allow otherwise. But even if each vertex locally agrees with σ 's cyclic ordering in some configuration C , we need C to be noncrossing before we can say that it *agrees with* σ as in Definition 2.6. In an angle constraint, σ has only one role: to identify which pairs of edges should be considered corners.

The other constraint allows for sliceforms:

Definition 7.2 (Sliceform Constraint). For a constrained abstract linkage $\mathcal{L}$ and a combinatorial embedding σ , a *Sliceform Constraint*, $\text{SliceCon}_{\mathcal{L}}(\sigma, S)$, is specified by a subset $S \subset V(G)$ of (some or all of the) vertices of degree 4. A configuration $C \in \text{Conf}(\mathcal{L})$ satisfies the sliceform constraint $\text{SliceCon}_{\mathcal{L}}(\sigma, S)$ if, for each *sliceform vertex* $v \in S$ with neighboring vertices w, x, y, z in cyclic order according to σ , points $C(w), C(v), C(y)$ are collinear in this order, and $C(x), C(v), C(z)$ are collinear in this order.

Sliceforms permit a limited form of “nonplanar” interaction while still being simulatable without crossings (cf. Lemma 7.22 on page 62), so they are our primary tool in circumventing the difficulties of planarity.

An extended linkage is simply a linkage with each type of constraint listed above, with a few convenient restrictions:

Definition 7.3 (Extended Linkage). An (ε, δ) -*extended linkage* where $0 < \delta < \varepsilon < \pi/4$ is defined as a constrained linkage $\mathcal{L}$ whose constraints K have the form

$$K = \{\text{AngleCon}_{\mathcal{L}}(\sigma, A, \Delta), \text{SliceCon}_{\mathcal{L}}(\sigma, S)\}$$

with the same combinatorial embedding σ , where at each corner Λ of $\mathcal{L}$, $A(\Lambda)$ is chosen from $\{90^\circ, 180^\circ, 270^\circ, 360^\circ\}$, and $\Delta(\Lambda)$ lies in $\{0, \delta, \varepsilon\}$. We will call $\mathcal{L}$ simply an *extended linkage* when ε and δ are the global constants given later by Equations 6–7.

Convention 7.4 (Drawing Extended Linkages). In drawings of extended linkages, all corners Λ are drawn at their “base” angle, $A(\Lambda)$. Corners marked with a solid gray sector are frozen ($\Delta(\Lambda) = 0$), and the rest have $\Delta(\Lambda) = \varepsilon$ unless otherwise specified. Vertices surrounded by small squares are pinned, and those marked with an “x” are sliceform vertices. See Figure 11 for examples.

It will often be useful to describe extended linkages not by the positions of their vertices, as with the projection π_X (cf. Definition 2.9 of *trace*), but by the angles of a chosen set of corners. We therefore define a function *Offset* that measures how these angles differ from their “neutral” values given by A in the angle constraint $\text{AngleCon}(\sigma, A, \Delta)$:

Definition 7.5 (The Offset Function). If $\mathcal{L}$ is an extended linkage and $Y = (\Lambda_1, \dots, \Lambda_k)$ is a tuple of (some or all of) the corners of $\mathcal{L}$, define the function $\text{Offset}_Y : \text{Conf}(\mathcal{L}) \rightarrow \mathbb{R}^{|X|}$ by

$$\text{Offset}_Y(C) = (\angle C(\Lambda_1) - A(\Lambda_1), \dots, \angle C(\Lambda_k) - A(\Lambda_k)).$$

We conclude this section with some useful facts about extended linkages.

Lemma 7.6. *If $\mathcal{L}$ is any (ε, δ) -extended linkage, then each vertex of $\mathcal{L}$ has degree at most 4. If a configuration $C \in \text{Conf}(\mathcal{L})$ has no crossings, then C agrees with combinatorial embedding σ .*

Proof. The $A(\Lambda)$ values around each vertex must sum to 2π , and each $A(\Lambda)$ is at least $\pi/2$, so each vertex has degree at most 4, verifying the first claim.

For the second claim, pick a vertex v , which we may assume has degree at least 2, for otherwise there is nothing to prove at v . Let $\Lambda_1, \dots, \Lambda_{\deg(v)}$ be the corners around v , and let C be *any* configuration of $\mathcal{L}$. Because each corner begins where the previous corner ends, the angles $\angle C(\Lambda_j)$ around $C(v)$ will always add to an integer multiple of 2π ; the edges surrounding v are configured in the cyclic order specified by σ precisely when each $\angle C(\Lambda_j)$ is positive and they all add to exactly 2π .

For each Λ_j ($1 \leq j \leq \deg(v)$), we have $0 < \angle C(\Lambda_j) < A(\Lambda_j) + \pi/4$ by our assumptions on A and Δ . Thus,

$$0 < \sum_{j=1}^{\deg(v)} \angle C(\Lambda_j) < \deg(v) \cdot \frac{\pi}{4} + \sum_{j=1}^{\deg(v)} A(\Lambda_j) = \deg(v) \cdot \frac{\pi}{4} + 2\pi < 4\pi,$$

so the angles $\angle C(\Lambda_j)$ must indeed add to 2π , as claimed. Thus, if C has no crossings, it agrees with σ . $\square$

Lemma 7.7. *For any (ε, δ) -extended linkage $\mathcal{L}$, the configuration space $\text{Conf}(\mathcal{L})$ is closed and semialgebraic. If $\mathcal{L}$ is connected and has at least one pin, then $\text{Conf}(\mathcal{L})$ is compact.*

Proof. Let $\mathcal{L}'$ be the underlying, unconstrained linkage. Then $\text{Conf}(\mathcal{L}')$ is closed and semi-algebraic (in fact algebraic), so it is enough to show that each of the constraints specifies a closed, algebraic subset of $\text{Conf}(\mathcal{L}')$.

For the angle constraint, suppose $\Lambda = uvw$ is a corner with $A(\Lambda) = \alpha$ and $\Delta(\Lambda) = \theta$. If $u = w$ then v has degree 1 and there is no constraint, so we may assume u and w are distinct vertices. We claim the angle constraint at Λ can be expressed by the dot product inequality $p' \cdot q \geq \cos \theta$, where $p = \frac{u-v}{\ell(uv)}$, $p' = \text{Rot}_\alpha p$ is the rotation of p by angle α around the origin, and $q = \frac{w-v}{\ell(wv)}$. To see this, first observe that in any configuration of $\mathcal{L}'$, p and q (and hence also p') will be unit vectors. The dot product then computes the cosine of the angle between p' and q . The inequality bounds this angle to the interval $[-\theta, \theta]$, as desired. So the angle constraint is indeed a closed, semialgebraic subset of $\text{Conf}(\mathcal{L}')$.

The sliceform constraint is simpler. If v is a sliceform vertex with neighboring vertices w, x, y, z in cyclic order, then the constraint may be expressed by the vector equalities

$$\frac{w-v}{\ell(vw)} = -\frac{y-v}{\ell(vy)} \quad \text{and} \quad \frac{x-v}{\ell(vx)} = -\frac{z-v}{\ell(vz)},$$

which is an algebraic (and therefore closed semialgebraic) constraint on $\text{Conf}(\mathcal{L}')$. The configuration space $\text{Conf}(\mathcal{L})$ is therefore closed and semialgebraic, as claimed.

If $\mathcal{L}$ is connected and has at least one pin, then $\text{Conf}(\mathcal{L})$ is closed and *bounded* and is therefore compact. $\square$

7.2 Detailed Overview of Strategy

Our Main Construction is primarily concerned with showing, for a finite set F of polynomials in $\mathbb{R}[x_1, y_1, \dots, x_m, y_m]$, how to construct an extended linkage that draws a bounded portion of the common zero set $Z(F)$, i.e., a set between $Z(F) \cap [-1, 1]^{2m}$ and $Z(F)$, up to a translation. This subsection outlines the key points of our approach, with full details to follow in subsequent subsections.

Convention 7.8. For convenience, points in $\mathbb{R}^2$ will be specified by cartesian coordinates or by complex numbers interchangeably, so $(3 \cos \theta, 3 \sin \theta)$ and $3 \exp(i\theta)$ are identical.

Our construction will be able to more easily manipulate *angles* than *vectors*, so wherever possible we will encode a point using two angles α and β , which relate to the point's Cartesian coordinates as follows:

$$\text{Rect}(\alpha, \beta) := (\cos \alpha, \sin \alpha) + (-\sin \beta, \cos \beta) - (1, 1) = e^{i\alpha} + i \cdot e^{i\beta} - (1 + i). \quad (2)$$

Usually α and β will be small angles, so geometrically, $\text{Rect}(\alpha, \beta)$ is the result of starting at $(-1, -1)$, stepping one unit at angle α (nearly horizontal), and then one unit at angle $\beta + \pi/2$ (nearly vertical). In particular, $\text{Rect}(0, 0) = (0, 0)$. We refer to these angles α and β as *angular coordinates* for the corresponding point. Similar strategies are employed in [12, 1].

Lemma 7.9. *For any angle $0 < \theta < \pi/4$, the function Rect is a homeomorphism from the region $[-\theta, \theta]^2$ onto its image, which is a compact region (specifically, the Minkowski sum of two circular arcs) containing the box $[-\theta/2, \theta/2]^2$.*

Proof. The region $M := \text{Rect}([-\theta, \theta]^2)$ is the Minkowski sum of two orthogonal circular arcs, each with central angle 2θ , as shown in Figure 6.

Rect is a local homeomorphism on $[-\theta, \theta]^2$, because the Jacobian determinant at (α, β) is $\cos(\alpha - \beta) \neq 0$. It may also be seen geometrically that Rect is injective on the boundary of $[-\theta, \theta]^2$, since the four circular arcs bounding M in Figure 6 are disjoint away from their endpoints. This is enough to ensure that Rect is bijective and in fact homeomorphic with its image, by [17, Corollary 2.6].

Region M contains the square $[-\sin \theta, \sin \theta + \cos \theta - 1]^2$, as shown in Figure 6. If we could show that $p(\theta) := \sin \theta - \theta/2$ and $q(\theta) := (\sin \theta + \cos \theta - 1) - \theta/2$ are both nonnegative, we would conclude that this square, and hence M itself, contains $[-\theta/2, \theta/2]^2$. Both p and q are concave down functions of $\theta \in [0, \pi/4]$, so it is enough to verify nonnegativity only at the endpoints 0 and $\pi/4$, which is straightforward: $p(0) = q(0) = 0$, $p(\pi/4) = (4\sqrt{2} - \pi)/8 > 0.314 > 0$, and $q(\pi/4) = \sqrt{2} - 1 - \pi/8 > 0.0215 > 0$. $\square$




Figure 6: The region M described in Lemma 7.9 (gray) is bounded by four circular arcs and contains the square $[-\theta/2, \theta/2]^2$.

For the constructions below, in place of rectangular coordinates (x_j, y_j) we will use angles (α_j, β_j) related by $(x_j, y_j) = 2r \cdot \text{Rect}(\alpha_j, \beta_j)$, where the radius $2r$ will be carefully

chosen later. We may write this equivalently as

$$\begin{aligned} x_j &= r \left(e^{i\alpha_j} + e^{-i\alpha_j} + ie^{i\beta_j} - ie^{-i\beta_j} - 2 \right), \\ y_j &= r \left(-ie^{i\alpha_j} + ie^{-i\alpha_j} + e^{i\beta_j} + e^{-i\beta_j} - 2 \right). \end{aligned} \quad (3)$$

Applying (3) for each $1 \leq j \leq m$, write $\vec{xy} := (x_1, y_1, \dots, x_m, y_m)$ as a function $\vec{xy}(\vec{\alpha\beta})$ where $\vec{\alpha\beta} := (\alpha_1, \beta_1, \dots, \alpha_m, \beta_m)$. Make this substitution into each polynomial $f \in F$, expand fully, and combine like terms, resulting in an expression of the form

$$f(\vec{xy}(\vec{\alpha\beta})) = \sum_{u=0}^3 \sum_{I \in \text{Coeffs}(2m, d)} i^u \cdot d_{u, I} \cdot e^{i \cdot (I \cdot \vec{\alpha\beta})},$$

where each $d_{u, I}$ is a nonnegative real number, and

$$\text{Coeffs}(2m, d) := \{(a_1, b_1, \dots, a_m, b_m) \in \mathbb{Z}^{2m} \mid |a_1| + |b_1| + \dots + |a_m| + |b_m| \leq d\}. \quad (4)$$

(The factors of r from (3) have been incorporated into coefficients $d_{u, I}$.) Only $d_{u, I} - d_{u+2, I}$ (with indices taken modulo 4) affects the total, so by further cancellation we may assume one or both of these coefficients is 0 (for each u, I pair). It will prove useful to change the form of this expression a bit: replace each term $e^{i \cdot (I \cdot \vec{\alpha\beta})}$ from (4) with $e^{i \cdot (I \cdot \vec{\alpha\beta})} - 1$, and add a constant term to compensate for this change. The result is

$$f(\vec{xy}(\vec{\alpha\beta})) = f(\vec{0}) + \sum_{u=0}^3 \sum_{I \in \text{Coeffs}(2m, d)} i^u \cdot d_{u, I} \cdot \left(e^{i \cdot (I \cdot \vec{\alpha\beta})} - 1 \right); \quad (5)$$

by evaluating at $\vec{\alpha\beta} = \vec{0}$ (equivalently, $\vec{xy} = 0$), it may be verified that the required constant term is in fact $f(\vec{0})$. Note that the four coefficients $d_{u, I}$ with $I = \vec{0}$ can be discarded, as these terms in (5) are identically 0.

Even though $f(\vec{xy}(\vec{\alpha\beta}))$ computes real values, these complex representations of f will prove more useful for turning polynomials into linkages. In particular, the modified form (5) has the advantage that each term other than the $f(\vec{0})$ term evaluates to $0 + 0i$ at $\vec{\alpha\beta} = \vec{0}$, which is crucial for ensuring that our linkages come with nice, rational coordinates when $f(\vec{0}) = 0$.

Lemma 7.10. *If f has $2m$ variables of total degree at most d and coefficients with magnitude at most M , and $r > 0$, then in representation (5) there are at most $2(2m)^d(2d+1)^d$ nonzero coefficients $d_{u, I}$, and the absolute values of all of these coefficients add to at most $6^d \cdot r^d \cdot M \cdot \binom{2m+d}{d}$. If f has integer coefficients and r is a positive integer, then entries $d_{u, I}$ are also integers, and they may be computed from f in deterministic time $O(\text{poly}(m^d, d^d, r^d, M))$.*

Proof. We may bound the size of $\text{Coeffs}(2m, d)$ by overcounting as follows: construct an integer vector $I \in \mathbb{Z}^{2m}$ by choosing d entries (or choosing all of them if $d > 2m$), assigning each chosen entry a value from $\{-d, -d+1, \dots, d\}$, and assigning 0 for the remaining entries. The number of such I is $\max(\binom{2m}{d}, 1) \cdot (2d+1)^d \leq (2m)^d(2d+1)^d$. There are therefore at most $2 \cdot (2m)^d(2d+1)^d$ nonzero coefficients $d_{u, I}$.

For each monomial $c \cdot \overrightarrow{xy}^J$ in f of total degree $k \leq d$, substitute each x_j and y_j as in Equation (3) and expand fully *without collecting like terms*. There are at most 6^k terms in this expansion (interpreting 2 as two separate terms $1 + 1$), each of the form $c \cdot r^k \cdot i^u \cdot e^{I \cdot \overrightarrow{\alpha\beta}}$, so the sum of the magnitudes of these coefficients is at most $6^k \cdot r^k \cdot c \leq 6^d \cdot r^d \cdot M$. Adding this quantity across all monomials in f , we obtain a sum no larger than $6^d \cdot r^d \cdot M \cdot \binom{2m+d}{d}$. Finally, collecting like terms (including cancelling $d_{u,I}$ and $d_{u+2,I}$ pairs as much as possible) can only decrease the sum of the magnitudes of the coefficients by the triangle inequality. The coefficients $d_{u,I}$ do not change when modifying Equation (4) to Equation (5) except that the coefficients with $I = \vec{0}$ can be discarded, as the corresponding terms are identically zero.

If the coefficients of f are integers, then every term in the above expansion also has integer coefficients. This expansion may be performed in time $O(\text{poly}(m^d, d^d, r^d, M))$ with straightforward methods. $\square$

We will use this angular representation as a template to compute each polynomial f in the linkage. Indeed, much like in Kempe’s original strategy [12] and especially in Abbott et al.’s correction thereof [1], we provide gadgets for the following tasks.

- The Start Gadget (Figure 15a) converts from rectangular position (x_j, y_j) to angles (α_j, β_j) .
- The Angle Sum Gadget (Figure 16b), built from the Angle Average Gadget (Figure 13), allows adding and subtracting angles to construct all of the $I \cdot \overrightarrow{\alpha\beta}$ values.
- The Vector Creation Gadget (Figure 15b) and Vector Rotation Gadget (Figure 11d) construct the vectors $i^u \cdot d_{u,I} \cdot e^{i(I \cdot \overrightarrow{\alpha\beta})}$.
- The Vector Sum Gadget (Figure 16c), built from the Vector Average Gadget (Figure 16a), allows adding vectors to construct the values $f(\overrightarrow{xy}(\overrightarrow{\alpha\beta})) - f(\vec{0})$ for each $f \in F$.
- The End Gadget (Figure 15c) constrains these values to equal $-f(\vec{0})$.

We employ several new ideas to ensure the resulting extended linkage $\mathcal{E}(F)$ is noncrossing. First, we construct a rigid grid of large square cells. Each gadget is isolated in one or $O(1)$ of these cells, and information is passed between gadgets/cells only using sliceform vertices along grid edges. In this way, these modular gadgets may be analyzed individually, as there is no way for distinct gadgets to intersect each other. Second, the linkage $\mathcal{E}(F)$ is an (ε, δ) -extended linkage, where ε and δ (the angle tolerances in $\text{AngleCon}_{\mathcal{E}(F)}(\sigma, A, \Delta)$) are global constants that restrict movement enough to prevent crossings within each separate gadget. The values of constants ε and δ are discussed in the next subsection. We rely on the Copy Gadget (Figure 11a) to copy angles and propagate them along paths of cells to distant gadgets in the grid. The Crossover Gadget (Figure 11b) allows these paths to cross, so we are not restricted to planar communication between gadgets. All of our gadgets make frequent use of the Parallel Gadget in Figure 10, which (with pins removed) keeps segments

parallel without otherwise restricting motion (within some neighborhood). Figure 17 shows an example of the gadgets working together.

Finally, we simulate extended linkage $\mathcal{E}(F)$ with a partially rigidified linkage $\mathcal{L}(F)$, in two steps. First, by replacing a vicinity of each sliceform vertex in $\mathcal{E}(F)$ with the Sliceform Gadget (Figure 18), we construct an extended linkage $\mathcal{E}'(F)$ that perfectly simulates $\mathcal{E}(F)$ but has no sliceforms. Then, we replace each edge of $\mathcal{E}'(F)$ with a rigidified orthogonal tree, connected to neighboring edges with the Angle Restrictor Gadget (Figure 19), which enforces the precise bounds in the angle constraints.

7.3 Constants and Parameters

Our construction uses a few carefully chosen and interconnected parameters, which we explain and collect here for ease of reference.

Linkage $\mathcal{E}(F)$ is an (ε, δ) -extended linkage, where the angle tolerances ε and δ are global constants tuned for specific purposes. The parameter ε has a single, simple goal: to constrain bar movement enough to ensure uniqueness (but not necessarily existence) of configurations. For example, Kempe’s fatal flaw was failing to prevent parallelograms from “flipping” into contraparallelograms, and vice versa. Abbott et al. [1] resolved this problem by using an additional *brace* along the midline of the contraparallelogram. Our angle constraints provide an alternative solution⁶: if a rectangular linkage is given angle constraints with tolerance $\varepsilon < \pi/2$, then it cannot flip into a contraparallelogram, since one of its angles, constrained to the interval $\pi/2 \pm \varepsilon$, would be forced to lie in the disjoint interval $3\pi/2 \pm \varepsilon$.

The smaller parameter δ has a double purpose. First, δ further restricts movement enough to prevent intersections and to control the global minimum feature size. But simply upper bounding how much each gadget can move is not enough: at some point, we must be able to say that our linkage can move *at least* enough to draw the loci we are interested in. This is δ ’s second purpose: in each of our gadgets, we ensure that every corner with tolerance δ can in fact realize any angle offset in the entire interval $[-\delta, \delta]$. In other words, while ε ensures uniqueness, δ measures *existence*.

The precise values we have chosen are

$$\varepsilon := \cos^{-1} \left(1 - \frac{3}{10} \cdot \frac{2n_\varepsilon}{n_\varepsilon^2 + 1} \right), \quad \text{where } n_\varepsilon = 5 \cdot 10^3; \quad \text{and} \quad (6)$$

$$\delta := \cos^{-1} \left(1 - \frac{3}{10} \cdot \frac{2n_\delta}{n_\delta^2 + 1} \right), \quad \text{where } n_\delta = 4 \cdot 10^{14}. \quad (7)$$

Numerically, these come out to approximately $\varepsilon \approx 0.0155$ and $\delta \approx 5.48 \cdot 10^{-8}$. Most uses of ε and δ rely only on the facts that $\varepsilon \leq \pi/16$ and $\delta \cdot 2^{18} \leq \varepsilon$, which may be numerically verified for the choices above. Only our final gadget, the Angle Restrictor Gadget in Section 7.7, places additional constraints on these values, in terms of both size

⁶Angle constraints *per se* only apply to *constrained* linkages. It is more correct to say that we resolve parallelogram flipping via angle constraints *combined* with a method of *simulating* these constraints using classical linkages.

and form: the trigonometric forms allow the Angle Restrictor Gadget to have rational coordinates.

Henceforth, all (ε, δ) -extended linkages will use these global constants for ε and δ , so we will refer to them simply as “extended linkages”.

Recall that the input F is a set of polynomials in variables $\vec{xy} = (x_1, y_1, \dots, x_m, y_m)$ each with total degree at most d and with integer coefficients bounded in magnitude by M . There are $|F| = s$ polynomials. The recurring parameters used in our construction are as follows:

- $r = \lceil d/\delta \rceil$ is half the radius used in the angular change of coordinates $(x, y) = 2r \cdot \text{Rect}(\alpha, \beta)$ applied to each polynomial.
- $Q = 40 \lceil 6^d r^d M \binom{2m+d}{d} / (6\delta) \rceil$ is the side-length of each cell in the background grid, chosen large enough to accommodate the necessary constructions. It is a positive integer divisible by 40.
- $R = 3Q/10$ is the radius used in angular coordinate representations in most of the gadgets. It is a positive integer, and $R > 4r$ may be readily verified.

It may seem unnecessary to use two different radii ($2r$ and R) for angular coordinates, but there is a good reason to do so. We evaluate the angular representation (5) of each polynomial term-by-term, constructing partial sums along the way. The size of these partial sums grows as a function of the radius $2r$ used in the polynomial change of coordinates, and they may be much larger than $2r$. We thus make r as small as possible to control the size of the partial sums, while R must be large enough to be able to compute the partial sums. As consolation, choosing a constant ratio R/Q allows most gadgets to scale uniformly with Q , making them simpler to analyze.

We have not attempted to optimize the parameters $\varepsilon, \delta, r, Q, R$ or the gadgets they depend on, instead choosing the most expedient values that allow the proofs to succeed. In fact, for convenience, we round many constant factors to powers of 2. Significant improvements are likely possible.

7.4 Hook Linkage and Parallel Linkage

To streamline the analysis of the gadgets in Sections 7.5 and 7.7, we collect here a few tools to provide explicit bounds on the displacement and rotation of certain linkages under small perturbation. We show that the *Parallel Gadget* in Figure 10, which is used in every gadget in Section 7.5, constrains two edges ab and ef to remain parallel without otherwise restricting their movement, so long as neither tries to move or rotate too far from how it started (Lemma 7.17). Along the way we analyze the *hook linkage*, i.e., a path of two edges ab and bc : if vertices a and c are each displaced by a small amount, we provide explicit bounds on how far b must move to preserve the lengths of edges ab and bc (Lemma 7.15). This hook linkage bound is helpful for the parallel linkage, but it is also directly invoked by a few of the gadgets in Sections 7.5 and 7.7.

Lemma 7.11. For $0 < t < 1/2$, we have $\sin t > 0.9t$ and $\sin^{-1} t < 1.1t$.

Proof. For $0 < t < 1/2$ we may compute

$$\sin t - \frac{9}{10}t > \left(t - \frac{1}{6}t^3\right) - \frac{9}{10}t = \frac{1}{30}t(3 - 5t^2) > 0$$

and

$$\frac{11}{10}t - \sin^{-1} t \geq \frac{11}{10}t - \left(t + \frac{1}{6}t^3\right) = \frac{1}{30}t(3 - 5t^2) > 0. \quad \square$$

Lemma 7.12. If A and B are complex numbers with $|B| \leq \frac{1}{2}|A|$, and if $\theta = \arg \frac{A+B}{A}$, then $|\theta| \leq \sin^{-1} \frac{|B|}{|A|} \leq 1.1 \frac{|B|}{|A|}$.

Proof. By replacing (A, B) with $(1, B/A)$, we may assume $A = 1$. Let $|B| = t$. Geometrically, $1 + B$ lies on a circle with center 1 and radius t , and the points p on this circle where $\arg p$ is maximized or minimized are the points of contact of the lines through the origin that are *tangent* to this circle. These angles are $\sin^{-1} t = \sin^{-1} \frac{|B|}{|A|}$, as claimed. Lemma 7.11 shows that $\sin^{-1} t \leq 1.1t$ so long as $t \leq 1/2$. $\square$

Lemma 7.13. If an edge of length ℓ rotates around an endpoint by angle θ , the other endpoint lands no farther than $\ell \cdot |\theta|$ from its initial position. More generally, if a path of edges $v_1v_2, v_2v_3, \dots, v_{k-1}v_k$ of total length ℓ reconfigures itself so that v_1 stays fixed while each edge changes its angle by at most $\pm\theta$, then vertex v_k is displaced by at most $\ell \cdot |\theta|$.

Proof. In the case of a single edge of length ℓ , the distance in question is

$$\ell \cdot \left| e^{i\theta} - 1 \right| = \ell \cdot \left| 2 \sin \frac{\theta}{2} \right| \leq \ell \cdot \left| 2 \cdot \frac{\theta}{2} \right| = \ell \cdot |\theta|,$$

as claimed. The more general claim follows by applying the previous fact to each edge separately, noting that the displacements simply add together, and applying the triangle inequality to their sum. $\square$

Lemma 7.14. Let C be a noncrossing configuration of a (possibly constrained) linkage $\mathcal{L}$ that is noncrossing with minimum feature size f . If C' is another configuration such that $|C'(v) - C(v)| \leq h$ for each vertex v , and if $h < f/2$, then C' is noncrossing with minimum feature size $f - 2h$.

Proof. Let p be a general point on $\mathcal{L}$, by which we mean either a vertex v or a fixed ratio k along an edge vw ; call the latter an *edge point*. For an edge point p , $C(p) := (1 - k) \cdot C(v) + k \cdot C(w)$, and likewise for $C'(p)$. We claim that $|C'(p) - C(p)| \leq h$ for any general point p . The result is immediate if p is a vertex, so consider the edge point case:

$$|C'(p) - C(p)| \leq (1 - k) \cdot |C'(v) - C(v)| + k \cdot |C'(w) - C(w)| \leq (1 - k) \cdot h + k \cdot h = h,$$

as claimed. It follows that if p and q are two general points on $\mathcal{L}$ that start with $|C(p) - C(q)| \geq f$, then $|C'(p) - C'(q)| \geq f - 2h$ by the triangle inequality. This means C' has minimum feature size $f - 2h$, so long as we can show that C' is noncrossing.

So suppose C' has a crossing. This means either a vertex $p = u$ coincides with a point q on edge vw not incident with u , or two edge points p and q from non-incident edges u_1u_2 and v_1v_2 coincide. Then $C'(p) = C'(q)$. On the other hand, $|C(p) - C(q)| \geq f$ because the smallest distance between non-incident edges is lower bounded by the smallest distance between a vertex and a non-incident edge, which is the minimum feature size f (Definition 2.5). As noted above, $|C'(p) - C'(q)| \geq f - 2h$, which is strictly positive because $h < f/2$, contradicting that $C'(p) = C'(q)$. For completeness, we include a proof of the second case. Consider triangle $C(p)C(v_1)C(v_2)$. If one of the triangle's angles at $C(v_1)$ or $C(v_2)$ is at least $\pi/2$, then $C(p)C(q)$ is at least as long as one of the features $C(p)C(v_j)$. So we may assume both these angles are $\leq \pi/2$, meaning the projection of $C(p)$ onto line $C(v_1)C(v_2)$ lands on segment $C(v_1)C(v_2)$. Redefine q as this projection, which does not lengthen $C(p)C(q)$. Now slide p and q while keeping $C(p)C(q)$ perpendicular to $C(v_1)C(v_2)$; there is some direction of sliding that does not lengthen $C(p)C(q)$, so continue until one of p or q hits an endpoint of its segment. Thus we obtain a distance between a vertex and a non-incident edge that is no longer than the original $C(p)C(q)$, as desired. $\square$

Lemma 7.15 (Hook Linkage). *Define $\mathcal{L}_{\text{hook}} = \mathcal{L}_{\text{hook}}(\ell_1, \ell_2, \theta)$ as the unconstrained linkage shown in Figure 7, initially configured at $a_0 := C_0(a) = 0$, $b_0 := C_0(b) = \ell_1$, and $c_0 := C_0(c) = \ell_1 + \ell_2 e^{i\theta}$, where $\theta \in [\pi/3, 2\pi/3]$.*

1. *For any point c' in the plane, there is at most one configuration with vertex c resting at point c' such that $\angle cba \in (0, \pi)$. If $h \leq \min(\ell_1, \ell_2)/2^4$ and $|c' - c_0| \leq h$, this configuration indeed exists, and it has the additional properties that b has been displaced at most $|b - b_0| \leq 1.5h$ from its initial position, the angle of bar ab has changed by at most $\pm 1.5h/\ell_1$, and likewise the angle of bc has changed by at most $\pm 1.5h/\ell_2$.*
2. *If the pin in $\mathcal{L}_{\text{hook}}$ is removed and any points a' and c' in the plane are chosen, there is at most one configuration placing vertex a at a' and vertex c at c' having $\angle cba \in (0, \pi)$. If $h \leq \min(\ell_1, \ell_2)/2^5$ and $|a' - a_0|, |c' - c_0| \leq h$, then this configuration exists, and it has the additional properties that b has moved at most $|b - b_0| \leq 4h$ from its initial position, the angle of bar ab has changed by at most $\pm 4h/\ell_1$, and the angle of bc has changed by at most $\pm 4h/\ell_2$.*

Proof. For Part 1, we first prove uniqueness. Suppose C some configuration has $C(a) = a_0$ and $C(c) = c'$. Observe that $C(b)$ must lie on the circle centered at a_0 with radius ℓ_1 , as well as the circle centered at c' with radius ℓ_2 . These circles cannot be concentric: indeed, $C(a) = C(c)$ would force $\angle C(c)C(b)C(a) = 0$, which is outside the allowed bounds. It follows that there are at most two positions for b , namely at $C(b)$ or at the reflection of $C(b)$ across $C(a)C(c)$. If C' is the configuration that uses the latter option, then $\angle C'(c)C'(b)C'(a) = 2\pi - \angle C(c)C(b)C(a) \notin (0, \pi)$. So C is indeed unique, as claimed.

We now prove existence, so suppose $|c' - c_0| \leq h$ as above. If bars ab and bc rotate by angles α and β respectively, point c ends up at

$$H(\alpha, \beta) := \ell_1 e^{i\alpha} + \ell_2 e^{i(\theta+\beta)}.$$

(This function H is similar to the Rect function defined in Section 7.2, except that (1) $H(0, 0)$ has not been shifted to the origin, and (2) H will not be used beyond this proof.)

Let $U = [-1.5h/\ell_1, 1.5h/\ell_1] \times [-1.5h/\ell_2, 1.5h/\ell_2]$. We will first show that there are some angles (α, β) in U with $H(\alpha, \beta) = c'$. To that end, we define a related function

$$G(\alpha, \beta) := \ell_1(1 + i \sin \alpha) + e^{i\theta} \cdot \ell_2(1 + i \sin \beta).$$

When $|\alpha| \leq 1.5h/\ell_1$ we may verify that

$$\begin{aligned} |\ell_1(1 + i \sin \alpha) - \ell_1 e^{i\alpha}| &= \ell_1 \cdot |1 - \cos \alpha| = \ell_1 \cdot 2 \sin^2 \frac{\alpha}{2} \\ &\leq \ell_1 \cdot \frac{1}{2} \alpha^2 \leq \ell_1 \cdot \frac{1}{2} \left(\frac{3h}{2\ell_1} \right)^2 = \frac{9h^2}{8\ell_1} \leq \frac{h}{12}, \end{aligned}$$

where the last inequality follows from $h/\ell_1 \leq 1/16$. By similar comparison between the other two terms of H and G , we conclude that when $(\alpha, \beta) \in U$, $|H(\alpha, \beta) - G(\alpha, \beta)| \leq h/6$.




Figure 7: The Hook Linkage from Lemma 7.15 consists of two incident bars.




Figure 8: The region $G(U)$ in Lemma 7.15 is a parallelogram centered at c_0 with dimensions as shown.

The linear map G is easy to describe: region $G(U)$ is a parallelogram centered on c_0 with dimensions as shown in Figure 8. The horizontal distance from the center to (the line containing) either vertical edge is

$$\ell_1 \cdot \sin \frac{3h}{2\ell_1} \cdot \sin \theta \geq \ell_1 \cdot \frac{9}{10} \frac{3h}{2\ell_1} \cdot \frac{\sqrt{3}}{2} > \frac{7}{6}h,$$

where our use of Lemma 7.11 is valid because $1.5h/\ell_1 \leq 3/32 < 1/2$. The distance to the (line containing) each diagonal edge is likewise greater than $7h/6$, so the interior of $G(U)$ contains the disk centered at c_0 of radius $7h/6$. Restricting attention to the boundary of U , we find that the loop $G(\partial U)$ surrounds this disk and is disjoint from it. In particular, this loop has winding number 1 around c' and does not come within distance $h/6$ of c' .

By our comparison of H and G above, the straight-line homotopy between loops $G(\partial U)$ and $H(\partial U)$ (which moves corresponding points along straight line segments) does not meet c' , so the latter loop also has winding number 1 around c' . If c' were not contained in region $H(U)$, then continuous function H would map the contractible set U into a non-contractible subset of $\mathbb{R}^2 \setminus c'$, which is impossible. So there indeed exists a pair of angles $(\alpha, \beta) \in U$ with $H(\alpha, \beta) = c'$.


Using these angles α and β , we claim that the configuration C that places a at $a_0 = 0$, b at $\ell_1 e^{i\alpha}$ and c at $H(\alpha, \beta) = c'$ will exhibit all of the desired properties. The angle at b will have changed by at most $1.5h/\ell_1 + 1.5h/\ell_2 \leq 3/16 < \pi/3$, so it will still lie in $(0, \pi)$. Also, b has moved by at most

$$|b - b_0| = \ell_1 \cdot |e^{i\alpha} - 1| = 2\ell_1 \cdot \left| \sin \frac{\alpha}{2} \right| \leq \ell_1 \cdot |\alpha| \leq 1.5h,$$

as required.

For Part 2, apply Part 1 with target position $\hat{c} := c' - a'$ and distance $2h$. Let $\hat{a} = 0, \hat{b}, \hat{c}$ be the resulting positions for vertices a, b, c . This configuration satisfies $|\hat{b} - b_0| \leq 1.5 \cdot 2h = 3h$, the angle of bar ab is in the interval $0 \pm 1.5 \cdot 2h/\ell_1 = 0 \pm 3h/\ell_1$, and likewise the angle of bar bc is in the interval $\theta \pm 1.5 \cdot 2h/\ell_2 = \theta \pm 3h/\ell_2$. Now we translate this configuration by a' to obtain vertices at $a' = \hat{a} + a', b' = \hat{b} + a', c' = \hat{c} + a'$. By the triangle inequality, $|b' - b_0| = |\hat{b} - b_0 + a'| \leq |\hat{b} - b_0| + |a'| \leq 3h + h = 4h$. The bar angles remain the same under translation; we increase the lead constant from 3 to $4 = 2^2$ for future ease of use. $\square$

We rely heavily on parallelograms to force edges to remain parallel. As described in Section 7.3, angle constraints can be used to ensure that parallelograms remain parallelograms, forbidding the contraparallelogram “flip” that Kempe failed to account for. We formalize this as follows:



(a) Parallelogram linkage $\mathcal{P}_1$, initially with opposite corners at $a = (0, 0)$ and $c = (x, y)$ where $x, y \geq 1$.

(b) Parallelogram linkage $\mathcal{P}_2$, with corners initially at $a = (0, 0)$, $b = (4, 0)$, $c = (6, 2)$, and $d = (2, 2)$.

Figure 9: Creating parallelograms with extended linkages. Corners marked with gray sectors are frozen, and vertices enclosed in squares are pinned, as in Convention 7.4.

Lemma 7.16 (Parallelogram Linkage). *Extended linkages $\mathcal{P}_1$ and $\mathcal{P}_2$ of Figures 9a and 9b, whose corners are given ε -tolerance angle constraints, are globally noncrossing with global minimum feature size at least $1/2$, and every configuration of either linkage has $\vec{dc} = \vec{ab}$. Furthermore, the configuration spaces $\text{Conf}(\mathcal{P}_1)$ and $\text{Conf}(\mathcal{P}_2)$ are perfectly described by the map Offset_Λ , which provides homeomorphisms*

$$\text{Conf}(\mathcal{P}_1) \simeq [-\varepsilon, \varepsilon] \quad \text{and} \quad \text{Conf}(\mathcal{P}_2) \simeq [-\varepsilon, \varepsilon].$$

In every configuration of $\mathcal{P}_2$, all vertices other than a and b are configured at least $1/2$ units away from the line through a and b .

Note that the pins in $\mathcal{P}_1$ and $\mathcal{P}_2$ merely serve to fix the reference coordinate frame. Without pins, the angle displacement at corner Λ uniquely determines the locations of c and d relative to a and b .

Proof. Begin with linkage $\mathcal{P}_1$ in Figure 9a, which has initial configuration $a = (0, 0)$ and $c = (x, y)$. In any configuration C ,

$$d = y \cdot (\cos(\pi/2 + \theta), \sin(\pi/2 + \theta))$$

where $\theta = \text{Offset}_\Lambda(C) \in [-\varepsilon, \varepsilon]$. Based on edge lengths, vertex c has two potential positions: $b+d-a$ (forming a parallelogram), or the reflection of this point across diagonal bd (forming a contraparallelogram). The former case indeed satisfies all of its constraints and is noncrossing with feature size at least $\min(x, y)/2$ because $\varepsilon < \pi/4$. In the latter case, corner dcb would have angle $3\pi/2 - \theta$, which is well outside the allowable range of $[\pi/2 - \varepsilon, \pi/2 + \varepsilon]$. So the parallelogram configuration exists uniquely.

In linkage $\mathcal{P}_2$, the vertices have initial positions $a = (0, 0)$, $b = (4, 0)$, $c = (6, 2)$, and $d = (2, 2)$. By similar arguments, each configuration C has $c - d = b - a$ and is uniquely determined by $\theta = \text{Offset}_\Lambda(C)$. If the minimum feature size were to drop below $1/2$, this would happen first when p came that close to edge ab or cd . We may calculate that these events correspond to respective angles $\theta = \pi/3 = 60^\circ$ and

$$\theta = \sin^{-1} \frac{1/2}{\sqrt{5}} - \sin^{-1} \frac{1}{\sqrt{5}} \approx -13.64^\circ.$$

Because $\varepsilon < \pi/16 = 11.25^\circ$, the interval $[-\varepsilon, \varepsilon]$ keeps the minimum feature size above $1/2$, as required. We may also observe that, in this interval, no vertex comes closer than $1/2$ to the x -axis, as required. $\square$

Parallelograms are especially useful in pairs, as with the linkage in Figure 10, which forces bars ab and ef to remain parallel while letting them move freely relative to each other (within some neighborhood). The classical counterpart of this linkage [12] suffers from nonuniqueness: When positions for bars ab and ef are chosen (necessarily parallel), usually there are two possible locations for bar cd . Angle constraints again improve the situation: one of these two positions violates the angle constraints (at all six named vertices!) and is therefore invalid.

Lemma 7.17 (Parallel Gadget).

1. The extended linkage $\mathcal{L}_{\text{parallel}}$ drawn in Figure 10, in which all corners are given an ε angle constraint, is globally noncrossing with minimum feature size at least $\frac{1}{2}$. In every configuration, vectors ab , dc , and ef are equal. Any choice of position for a , b , and e can be completed to at most one configuration of $\mathcal{L}_{\text{parallel}}$.
2. If three points a' , b' , and e' are chosen at distance at most $\varepsilon/2^7$ from their initial positions of $a_0 = (0, 0)$, $b_0 = (4, 0)$, and $e_0 = (0, 4)$ respectively such that $|b' - a'| = 4$, then a unique configuration C exists with $C(a) = a'$, $C(b) = b'$, and $C(e) = e'$. This configuration C moves each vertex no farther than $\varepsilon/2$ from its initial position and rotates each edge by at most $\pm\varepsilon/2$.




Figure 10: Left: The Parallel Gadget allows e to move freely in a neighborhood of its initial position while forcing ef to remain parallel to ab . Corners marked with gray sectors are frozen, as in Convention 7.4. Edge lengths are all in $\{1, 2, 4\}$ and are drawn to scale. Right: a schematic representation of the same gadget.

Note. In this Parallel Gadget, once a and b have been configured, angles $\text{Offset}_{(\Lambda, \Gamma)}(C) = (\alpha, \beta) \in [-\varepsilon, \varepsilon]^2$ uniquely determine the rest of the configuration C , but it is worth mentioning that not all such pairs (α, β) actually give rise to valid configurations. Indeed, the angle to the left of vertex d becomes $\pi + \alpha - \beta$, so the ε -tolerance forces $|\alpha - \beta| \leq \varepsilon$ as an additional constraint (in fact, the only other constraint). This lemma does not need to worry about this fact because it does not attempt to explore the entire configuration space of $\mathcal{L}_{\text{parallel}}$, instead restricting attention to $|\alpha|, |\beta| \leq \varepsilon/2$.

Proof. Let C_0 be the initial configuration of $\mathcal{L}_{\text{parallel}}$ shown in Figure 10, and for each vertex v , we use the shorthand $v_0 = C_0(v)$.

This linkage is made from two instances of $\mathcal{P}_2$ (one reflected) attached along a bar. In any configuration C of $\mathcal{L}_{\text{parallel}}$, by Lemma 7.16, vectors ab and dc must be equal, as must dc and ef . Each vertex other than c and d , so the smallest feature of C comes from one of the two instances of $\mathcal{P}_2$ and is thus at least $1/2$.

For uniqueness, suppose positions a' , b' , and e' have been chosen for a , b , and e . First observe that the position of d is forced: treating the paths from a to d and from d to e as bars ad, de of fixed lengths $\ell(ad) = \ell(de) = 2\sqrt{2}$, and thus ade as a *hook linkage*, Lemma 7.15 shows that there exists a unique point d' at distance $2\sqrt{2}$ from each of a' and e' with $\angle e'd'a' < \pi$, which is required by one of the angle constraints at vertex d . The previous paragraph then shows that c must be placed at $c' = d' + (b' - a')$ and f must be placed at $e' + (b' - a')$. Finally, the rigidified paths (e.g., from a to d) are uniquely determined by their endpoints, so all vertex locations are forced, as claimed.

Define $k := \varepsilon/2^7$. To prove Part 2, first consider the easier case where $a' = a_0$, $b' = b_0$, and $|e' - e_0| \leq 6k$. We will show that a (necessarily unique) configuration C with $C(a) = a' = a_0$, $C(b) = b' = b_0$, and $C(e) = e'$ exists, moves each vertex by at most $28k$, and rotates each edge by at most $9k$.

As described above, Lemma 7.15 (Hook Linkage) allows us to construct the unique positions b' , c' , and f' where vertices b, c, f must lie, which then determines the entire configuration; we must show that this forced configuration C is actually valid, i.e., satisfies

its angle constraints. Because $6k < \min(\ell(ad), \ell(de))/32 = (2\sqrt{2})/32$, Lemma 7.15 proves that $a'd'$ and $d'e'$ are each rotated from a_0d_0 and d_0e_0 by angle at most $\pm 4 \cdot 6k/(2\sqrt{2}) = 6k\sqrt{2} \leq 9k$, so the angles at C 's corners differ from those of C_0 by at most $18k < \varepsilon$ and therefore satisfy their ε -tolerance angle constraints. The same lemma also shows that $|c' - c_0| = |d' - d_0| \leq 24k$. The bound $|f' - f_0| = |e' - e_0| \leq 6k \leq 24k$ follows by assumption.

It remains to show that $|C(v) - v_0| \leq 28k$ for each vertex v of $\mathcal{L}_{\text{parallel}}$ other than a, b, c, d, e, f . Focus on vertex p first, and let $p' = C(p)$. Because the path from e to d through p is frozen in Figure 10, triangles $e_0p_0d_0$ and $e'p'd'$ are congruent. Thus p_0 and p' can be written as the same affine linear combination of e_0, d_0 and e', d' respectively. Because $d_0 = (2, 2) = 2 + 2i$, $e_0 = (0, 4) = 4i$, and $p_0 = (0, 3) = 3i$, we have the (unique) solution

$$p_0 = \left(\frac{1-i}{4}\right) d_0 + \left(\frac{3+i}{4}\right) e_0,$$

and thus the same for p' . Therefore

$$p' - p_0 = \left(\frac{1-i}{4}\right) (d' - d_0) + \left(\frac{3+i}{4}\right) (e' - e_0).$$

Because $|d' - d_0|$ and $|e' - e_0|$ are upper bounded by $24k$ (as argued above), the triangle inequality gives

$$|p' - p_0| \leq \left|\frac{1-i}{4}\right| \cdot 24k + \left|\frac{3+i}{4}\right| \cdot 24k \leq 28k.$$

The remaining vertices other than a, b, c, d, e, f follow the same analysis as p because the four frozen paths are isometric, leading to the same affine weights or their conjugates, and all six points a, b, c, d, e, f are displaced by at most $24k$ (as argued above).

Now we tackle Part 2 in general; as was done for the Hook Linkage in Lemma 7.15, we will reduce this general case to the special case where a and b are fixed. So suppose a', b' , and e' are chosen at distance at most k from a_0, b_0, e_0 respectively and with $|a' - b'| = 4$. Uniqueness is guaranteed as above, so we focus on existence.

Let $\theta = \arg(b' - a')$, i.e., $b' - a' = 4 \exp(i\theta)$, and observe that $|\theta| \leq 1.1 \cdot \frac{2k}{4} \leq k$ by Lemma 7.12, using $A = 4$ and $B = b' - b_0 + a' - a_0$, which satisfy $|B| \leq 2k < \frac{1}{2}|A|$. Define $a'' = a_0$, $b'' = b_0$, and $e'' = (e' - a') \cdot \exp(-i\theta)$, which may be obtained from a', b', e' by translating by $-a'$ and then rotating by angle $-\theta$. We have

$$\begin{aligned} |e'' - e_0| &= |(e' - a') \cdot \exp(-i\theta) - e_0| \\ &\leq |(e' - a') \cdot \exp(-i\theta) - e_0 \cdot \exp(-i\theta)| + |e_0 \cdot \exp(-i\theta) - e_0| \\ &= |(e' - a') - e_0| + |e_0 \cdot \exp(-i\theta) - e_0| \\ &\leq |e' - e_0| + |a'| + |e_0 \cdot \exp(-i\theta) - e_0| \\ &\leq k + k + 4|\theta| \\ &\leq 6k. \end{aligned}$$

By the special case discussed above, there exists a unique configuration C with $C(a) = a_0$, $C(b) = b_0$, and $C(e) = e''$ where each vertex is at most $28k$ away from its initial position, and where each bar is rotated by at most $9k$ from its initial direction. Now define C' as the

configuration formed by rotating configuration C through angle θ (around the origin) and then translating the result by a' . This C' is a valid configuration of $\mathcal{L}_{\text{parallel}}$ because C itself was, and furthermore, $C'(a) = a'$, $C'(b) = b'$, and $C'(e) = e'$. So C' is the promised unique configuration.

Finally, we measure how much C' perturbs the vertices and edge rotations. For each vertex v , we have $|C(v)| \leq |C_0(v)| + |C(v) - C_0(v)| \leq 3\sqrt{5} + 6k \leq 7k$, so rotating by angle θ moves $C(v)$ by at most $7|\theta| \leq 7k$ by Lemma 7.13. Translating by a' moves each vertex by another $|a'| \leq k$, so in total,

$$|C'(v) - v_0| \leq 28k + 7k + k = 35k < \frac{\varepsilon}{2}$$

for each vertex v . Likewise, C has rotated each bar by at most $9k$, so C' rotates each bar by at most $9k + |\theta| \leq 10k < \varepsilon/2$. This finishes the proof. $\square$

7.5 The Cell Gadgets

With the above tools in place, we now present the modular gadgets themselves. Each will inhabit one or a constant number of $Q \times Q$ square cells, where Q was defined in Section 7.3. Each cell has pairs of *transmission edges* attached at *transmission vertices*, which are slice-form vertices at the midpoints of cell edges (usually labelled b_j), and the corners around each transmission vertex will have tolerance δ or 0—never ε —as required by the gadget. All other vertices will have tolerance ε or 0, never δ . If v is a transmission vertex with transmission edge vw , the corner uvw is the corresponding *transmission corner*; the angles at these transmission corners are the gadgets' only means of communicating with each other.

We use a few naming conventions for corners and their angle offsets. Corners are usually labelled by Λ_j or Γ_j , chosen because these letters visually resemble corners, and their corresponding angle offsets are $\alpha_j = \alpha_j(C) := \text{Offset}_{\Lambda_j}(C)$ and $\beta_j = \beta_j(C) := \text{Offset}_{\Gamma_j}(C) = \beta_j$, chosen because ‘ Λ ’ resembles a capital ‘ A ’ and ‘ Γ ’ resembles part of a capital ‘ B ’. These corners should be considered a pair: α_j and β_j together will form the angular coordinates for some meaningful vector. Corners that are *not* paired in this way will be labelled Θ_j , corresponding to angle offsets $\theta_j = \theta_j(C) := \text{Offset}_{\Theta_j}(C)$.

The analyses in this section do not depend on the precise values of ε and δ ; only on the facts that $\varepsilon \leq \pi/16$ and $\delta \leq \varepsilon/2^{18}$. For each gadget $\mathcal{L}$ below, we will prove a statement of the following form: $\varphi : \text{Conf}(\mathcal{L}) \rightarrow U$ is a homeomorphism, where φ is a function defined in terms of angle displacements Offset_Y and possibly vertex positions π_X , and U is some subset of some $\mathbb{R}^k$. This provides an explicit parametrization of $\mathcal{L}$'s configuration space: each configuration is uniquely and continuously determined by just the corners in Y and vertices in X , which are related by the shape of U . For example, the Angle Average Gadget $\mathcal{L}_{\text{avg}}$ is exactly parametrized by

$$U_{\text{avg}} = \left\{ (\theta_1, \theta_2, \theta_3) \in [-\delta, \delta]^3 \mid \theta_2 = \frac{\theta_1 + \theta_3}{2} \right\}.$$

This means the angle offsets θ_j at corners Θ_j must satisfy this average condition, and any triple of angles in $[-\delta, \delta]$ that does satisfy it gives rise to a unique configuration of $\mathcal{L}_{\text{avg}}$.

Furthermore, this configuration depends continuously on the angles θ_j . Because $\text{Conf}(\mathcal{L})$ is compact by Lemma 7.7, the following well-known fact says that this continuity comes for free once we know φ to be a bijection:

Lemma 7.18 ([20, Theorem 4.17]). *If $\varphi : A \rightarrow B$ is a continuous bijection of metric spaces and A is compact, then φ^{-1} is continuous, i.e., φ is a homeomorphism.*

When using some of these gadgets later on, we will not need to work with the exact configuration space U , and it will instead suffice to know that the gadgets have sufficient mobility, i.e., U contains a large enough neighborhood of the initial configuration. The lemma below also specifies these *useful neighborhoods* when appropriate.

For clarity, Figures 11, 13, 15, 16a, and 22 draw Parallel Gadgets using the schematic shorthand from Figure 10 (right). All other vertices in these figures are explicitly marked by a dot or an ‘x’, and we call these vertices *marked*. We will analyze marked vertices separately from the vertices internal to Parallel Gadgets.

We now analyze each gadget in turn. They are organized by similarity of *analysis*, not of *function*. This first set of gadgets all exhibit *scale invariance*: all coordinates are constant multiples of the cell edge length Q , so increasing Q serves only to increase the global minimum feature size of each gadget without otherwise affecting the gadget. This greatly simplifies the analysis.

Lemma 7.19 (Scale Invariant Cell Gadgets). *The abstract extended linkages described below are globally noncrossing with global minimum feature size at least $1/2$, and with their configuration spaces U exactly parameterized as specified in each case. Each gadget is described by an initial configuration where $a_3 = (0, 0)$ and $a_1 = (Q, Q)$, all marked vertices have coordinates that are integer multiples of $Q/10$, and so all vertices have coordinates that are integer multiples of $Q/40$ (an integer). For some gadgets, we prove that the configuration space U contains a conveniently shaped subset, identified below as a *useful neighborhood*.*

Copy Gadget. *The Copy Gadget $\mathcal{L}_{\text{copy}}$ drawn in Figures 11a and 12 constrains its four sliceform angles to remain equal. Specifically, $\text{Offset}_{(\Theta_1, \Theta_2, \Theta_3, \Theta_4)}$ is a homeomorphism from $\text{Conf}(\mathcal{L}_{\text{copy}})$ to*

$$U_{\text{copy}} = \{(\theta, \theta, \theta, \theta) \mid \theta \in [-\delta, \delta]\}.$$

Crossover Gadget. *The Crossover Gadget $\mathcal{L}_{\text{cross}}$ drawn in Figure 11b constrains opposite sliceform angles to remain equal. Specifically, $\text{Offset}_{(\Theta_1, \Theta_2, \Theta_3, \Theta_4)}$ is a homeomorphism from $\text{Conf}(\mathcal{L}_{\text{cross}})$ to*

$$U_{\text{cross}} = \{(\theta_1, \theta_2, \theta_1, \theta_2) \mid \theta_1, \theta_2 \in [-\delta, \delta]\}.$$

Angular Gadget. *The Angular Gadget $\mathcal{L}_{\text{angular}}$ drawn in Figure 11c relates the rectangular coordinates of vertex g with the angles α, β in its angular representation. Specifically, the function $\pi_g \times \text{Offset}_{\Lambda, \Gamma}$ is a homeomorphism from $\text{Conf}(\mathcal{L}_{\text{angular}})$ to*

$$U_{\text{angular}} = \{((x, y), (\alpha, \beta)) \mid \alpha, \beta \in [-\delta, \delta] \text{ and } (x, y) = (Q/2, Q/2) + R \cdot \text{Rect}(\alpha, \beta)\}.$$

Useful Neighborhood: if (x, y) is any point at most $R\delta/2$ away from the center of the cell, a unique configuration exists placing g at (x, y) .




Figure 11: A selection of extended linkage gadgets for manipulating angles and vectors. These scale-invariant gadgets are analyzed in Lemma 7.19. In these figures, $\text{Offset}_\Lambda(C) = \alpha$, $\text{Offset}_\Gamma(C) = \beta$, $\text{Offset}_\Theta(C) = \theta$, and similarly with subscripts. As described in Convention 7.4, vertices surrounded by squares are pinned; those marked with an “x” are sliceform vertices; and corners marked with a solid gray sector are frozen, i.e., have tolerance 0. Unfrozen sliceform vertices b_j are assigned tolerance δ , while all remaining unfrozen corners have tolerance ε . The pins at the vertices a_j are for clarification only; in the overall construction, these nodes a_j are forbidden from moving by other means, so these explicit pins are unnecessary. The thick edges emphasize the rigidified background grid (in gray) and the transmission edges (in black), along with everything rigidly attached to those edges. Each cell has dimensions $Q \times Q$, and all marked vertices have coordinates at integer multiples of $Q/10$, indicated by cyan axis notches.

Vector Rotation Gadget. The Vector Rotation Gadget $\mathcal{L}_{\text{rot}}$ (Figure 11d) enables constructing the 90° rotation of a given vector. Specifically, $\text{Offset}_{(\Lambda_1, \Gamma_1, \Lambda_2, \Gamma_2)}$ is a homeomorphism from $\text{Conf}(\mathcal{L}_{\text{rot}})$ to

$$U_{\text{rot}} = \{(\alpha_1, \beta_1, \alpha_2, \beta_2) \in [-\delta, \delta]^4 \mid \text{Rect}(\alpha_1, \beta_1) = i \cdot \text{Rect}(\alpha_2, \beta_2)\}.$$

Useful Neighborhood: if α_1 and β_1 are such that $|R \cdot \text{Rect}(\alpha_1, \beta_1)| \leq R\delta/2$, then a (necessarily unique) configuration C exists with $\text{Offset}_{(\Lambda_1, \Gamma_1)}(C) = (\alpha_1, \beta_1)$.

Angle Average Gadget. The Angle Average Gadget $\mathcal{L}_{\text{avg}}$ drawn in Figures 13 and 14 constrains one sliceform angle to equal the average of the other two. Specifically, $\text{Offset}_{(\Theta_1, \Theta_2, \Theta_3)}$ is a homeomorphism from $\text{Conf}(\mathcal{L}_{\text{avg}})$ to

$$U_{\text{avg}} = \{(\theta_1, \theta_2, \theta_3) \mid \theta_1, \theta_3 \in [-\delta, \delta] \text{ and } \theta_2 = (\theta_1 + \theta_3)/2\}.$$

Proof. Each of these gadgets is *scale invariant*: scaling down by a factor of $40/Q$ results in the same gadgets with the value of Q replaced by 40. Showing that each of these scaled down gadgets has global minimum feature size at least $1/2$ will show that the original gadgets have global minimum feature size at least $\frac{Q}{40} \cdot \frac{1}{2} \geq \frac{1}{2}$, so we may proceed as if $Q = 40$ instead of being defined as in Section 7.3.

Copy Gadget. Consider first the Copy Gadget $\mathcal{L}_{\text{copy}}$ of Figure 11a, whose movement is illustrated in Figure 12. The idea of the proof is as follows: ε enforces uniqueness and the claimed angle relationships, while δ ensures existence.

Constraint and Uniqueness. Relying only on our choice of ε , we first verify that angles $\theta_1, \theta_2, \theta_3, \theta_4$ must be equal, and that there is *at most* one configuration having $\theta_1 = \theta_2 = \theta_3 = \theta_4 = \theta$ for each value θ . Path $b_1c_1d_1e_1f_1$ must move as a rigid subassembly because of its frozen corners. So angle offset θ_1 determines the locations of vertices c_1, d_1, e_1, f_1 : they are simply rotated around b_1 from their initial positions by angle θ_1 . Positions for the rest of




Figure 12: Snapshots of the configuration space of the Copy Gadget (cf. Figure 11a, Lemma 7.19), not drawn to scale. Angle offsets shown are -3° , 0° and 3° respectively. The gadget has been slightly modified and simplified to emphasize movement. The thick blue edges emphasize the transmission edges.

the marked vertices in Figure 11a are likewise determined by θ_2 , θ_3 , and θ_4 . By Lemma 7.17, the Parallel Gadgets force $d_j c_j \parallel e_{j+1} f_{j+1}$ for $1 \leq j \leq 4$ (with indices taken cyclically) in any configuration of $\mathcal{L}_{\text{copy}}$. It follows that $\theta_1 = \theta_2 = \theta_3 = \theta_4$ in any configuration; let θ be this common value. By the same lemma, the positions of d_j , c_j , and e_{j+1} uniquely determine the positions of the other vertices within their Parallel Gadget, whenever such a configuration exists. So each configuration is indeed uniquely characterized by its single angle θ . Said differently, the map $\varphi = \text{Offset}_{(\Theta_1, \Theta_2, \Theta_3, \Theta_4)}$ injectively maps $\text{Conf}(\mathcal{L}_{\text{copy}})$ to U_{copy} .

Existence. On its own, the above paragraph does not prove that *every* angle $\theta \in [-\delta, \delta]$ gives rise to a configuration of $\mathcal{L}_{\text{copy}}$, or said differently, that φ is a bijection. Indeed, too large a θ would stretch or compress the Parallel Gadgets too far, resulting in broken angle constraints or worse. However, we next show that our choice of δ is small enough to guarantee φ is indeed a bijection. By Lemma 7.13, because the path b_j, c_j, d_j, e_j, f_j has length $\frac{6}{10} \cdot 40 = 24$, the rotations by θ around each b_j move each vertex c_j, d_j, e_j, f_j at most $24\delta < \varepsilon/2^7$ from its starting position, and guarantee that $e_j c_j \parallel e_{j+1} f_{j+1}$ for each $j = 1, 2, 3, 4$. Then Lemma 7.17 (Parallel Gadget) shows that each Parallel Gadget can configure itself at these new endpoints, making sure that each vertex moves by no more than $\varepsilon/2$ from its starting point and each edge rotates by no more than $\varepsilon/2$ from its original direction. This is enough to guarantee that all angle constraints are satisfied, so the configuration indeed exists. By Lemma 7.14, this configuration is noncrossing with minimum feature size at least $1 - 2 \cdot \varepsilon/2 > 1/2$.

Crossover Gadget. The proof proceeds similarly to that of $\mathcal{L}_{\text{copy}}$ above.

Constraint and Uniqueness. In any configuration of $\mathcal{L}_{\text{cross}}$, the Parallel Gadgets and the sliceform vertex g enforce $b_3 c_3 \parallel f_3 g \parallel g f_1 \parallel c_1 b_1$ and likewise $b_4 g \parallel g f_2 \parallel c_2 b_2$. This shows that $\theta_1 = \theta_3$ and $\theta_2 = \theta_4$ throughout $\text{Conf}(\mathcal{L}_{\text{cross}})$, and that these two angles uniquely determine the entire configuration (when such a configuration exists): each bold rigid assembly $b_j c_j d_j$ ($1 \leq j \leq 3$) has rotated by angle θ_j around its (stationary) pivot b_j , edge $b_4 g$ has rotated by angle θ_4 around b_4 , each smaller rigid assembly $g f_j e_j$ has likewise rotated by angle θ_j around the *new* location of g , and by Lemma 7.17 each Parallel Gadget can be drawn to connect its endpoints in at most one way.

Existence. Consider a given pair of angles $\theta_1, \theta_2 \in [-\delta, \delta]$. Because every marked vertex has a path of length at most $\frac{8}{10} \cdot 40 = 32$ to a transmission vertex b_i (avoiding Parallel Gadgets), by Lemma 7.13, the forced location of each marked vertex is at most $32\delta < \varepsilon/2^7$ units away from its initial position. The opposite ends of each Parallel Gadget are indeed parallel, allowing the Parallel Gadgets to configure themselves by Lemma 7.17. As before, vertices have been displaced by at most $\varepsilon/2$ and edges have been rotated by at most $\varepsilon/2$, guaranteeing the configuration's validity and minimum feature size of at least $1 - \varepsilon \geq 1/2$.

Angular Gadget. *Constraint and Uniqueness.* This is straightforward: the Parallel Gadgets force ef and assembly $b_4 c_4 d_4$ to have rotated by α , and likewise fg and the assembly containing b_1, c_1, d_1 to have rotated by β around the new position of f , so $g = (Q/2, Q/2) + R \cdot \text{Rect}(\alpha, \beta)$, so everything is uniquely specified.

Existence. This is also easy: because every marked vertex has a path of length at most $\frac{8}{10} \cdot 40 = 32$ to a vertex of the rigidified background grid (avoiding Parallel Gadgets),

Lemma 7.13 shows that each marked vertex has moved at most $32\delta \leq \varepsilon/2^7$ from its initial position, so the Parallel Gadgets can be configured to connect their opposite edges, which can be checked to be parallel. As above, we may also conclude that each edge has rotated at most $\pm\varepsilon/2$ from its initial direction. We conclude the validity, noncrossing, and minimum feature size of the resulting configuration as before.

Useful Neighborhood. Suppose $(x, y) = g'$ is some point at distance at most $R\delta/2$ from the center of the cell. By Lemma 7.9 with $\theta = \delta$, there is exactly one pair $(\alpha, \beta) \in [-\delta, \delta]^2$ such that $R \cdot \text{Rect}(\alpha, \beta) = g' - (Q/2, Q/2)$, i.e., such that $(g', (\alpha, \beta)) \in U_{\text{angular}}$.

Vector Rotation Gadget. *Constraint and Uniqueness.* This is perfectly analogous to previous analyses—in fact, this gadget is essentially two conjoined Angular Gadgets—so we have omitted this analysis.

Existence. This is also analogous and therefore omitted.

Useful Neighborhood. Vertex g must land in the region $(Q/2, Q/2) + R \cdot \text{Rect}([-\delta, \delta]^2)$ (because of Λ_1, Γ_1) as well as in this same region rotated by $\pi/2$ (because of Λ_2, Γ_2). These two regions are not identical, but by Lemma 7.9 with $\theta = \delta$, they both contain the square $(Q/2, Q/2) + [-R\delta/2, R\delta/2]^2$. Any point within this square therefore gives rise to a valid configuration of $\mathcal{L}_{\text{rot}}$.

Angle Average Gadget. This analysis is more involved than the previous ones, but it follows the same outline. The gadget $\mathcal{L}_{\text{avg}}$ is depicted in Figure 13 (drawn to scale). We choose f_1, f_3 , and e_1 such that $f_3e_3he_2$ and $he_1f_1e_2$ are kites which are initially similar. We will show that these kites remain similar throughout $\text{Conf}(\mathcal{L}_{\text{avg}})$, which is pivotal (so to speak) to how the gadget works. This gadget is based on the *Reflector Gadget* from [2], but as a proof of the Reflector Gadget's movement was not included in that extended abstract, we provide a self-contained proof of our Angle Average Gadget here. The motion of this gadget has been illustrated in Figure 14, which is *not* drawn to scale to emphasize the movements.

Constraint and Uniqueness. We first argue that angles $\theta_1, \theta_2, \theta_3$ uniquely determine their configuration, when one exists. Arguments like those above show that all marked vertices *other* than e_2, k_1 , and k_2 depend straightforwardly on angles $\theta_1, \theta_2, \theta_3$, by rotations and translations. Treating he_2 as a single fixed-length bar, Lemma 7.15 (Hook Linkage) on he_2f_3 (which has angle $\cos^{-1}(-1/\sqrt{5}) < 2\pi/3$) shows that e_2 has at most one valid position, and the same lemma shows that k_1 and k_2 are likewise uniquely determined.

We next argue that, for a configuration to exist, $\theta_2 = (\theta_1 + \theta_3)/2$ must hold. Focus for now on kites $he_1f_1e_2$ and $f_3e_3he_2$. Define $\alpha_1 = \angle e_1he_2$ and $\alpha_2 = \angle f_1e_1h$, and similarly, $\beta_1 = \angle e_3f_3e_2$ and $\beta_2 = \angle he_3f_3$. Considering path $m_1e_1he_2f_3m_3$, segments m_1e_1 and f_3m_3 being parallel (which is guaranteed by the two upper Parallel Gadgets) implies that

$$\alpha_2 + \alpha_1 + (2\pi - \beta_2) + (\pi - \beta_1) = 3\pi,$$

i.e., $\alpha_1 + \alpha_2 = \beta_1 + \beta_2$. We will show that this forces the kites to be similar.

Kite $f_3e_3he_2$ has edge lengths 4 and $4\sqrt{5}$. We may compute that $|e_2 - e_3| = 8 \sin \frac{\beta_1}{2}$,

$\angle e_2 e_3 f_3 = (\pi - \beta_1)/2$, and

$$\angle h e_3 e_2 = \cos^{-1} \frac{|e_2 - e_3|/2}{4\sqrt{5}} = \cos^{-1} \frac{t}{\sqrt{5}},$$

where $t = \sin \frac{\beta_1}{2}$. So, in terms of t ,

$$\beta_1 + \beta_2 = \frac{\pi}{2} + \frac{1}{2}\beta_1 + \angle e_2 e_3 h = \frac{\pi}{2} + \sin^{-1} t + \cos^{-1} \frac{t}{\sqrt{5}}.$$

Analogously, if $s = \sin \frac{\alpha_1}{2}$, we have

$$\alpha_1 + \alpha_2 = \frac{\pi}{2} + \sin^{-1} s + \cos^{-1} \frac{s}{\sqrt{5}}.$$

But the function $f(x) := \sin^{-1} x + \cos^{-1} \frac{x}{\sqrt{5}}$ is strictly monotonically increasing on $x \in [0, 1]$, because its derivative is $(1-x^2)^{-1/2} - (5-x^2)^{-1/2} > 0$. So $f(s) = f(t)$ implies $s = t$, meaning $\alpha_1 = \beta_1$ and $\alpha_2 = \beta_2$.

Revisiting Figure 13, let $\kappa = \cos^{-1} \frac{-1}{\sqrt{5}}$ be the initial value of angles $\alpha_2 = \angle f_1 e_1 h$ and $\beta_2 = \angle h e_3 f_3$. The new values for these angles are $\kappa + \theta_1 - \theta_2$ and $\kappa + \theta_2 - \theta_3$ respectively. But these values must be equal, so indeed $\theta_2 = (\theta_1 + \theta_3)/2$.

Existence. Conversely, we must show that if angles $\theta_j \in [-\delta, \delta]$ are chosen from the set $U_{\angle \text{avg}}$ (i.e., satisfying $\theta_2 = (\theta_1 + \theta_3)/2$), then a configuration indeed exists. Use the θ_j to configure all marked vertices other than e_2 , k_1 , and k_2 by rotations as above. Because every marked vertex has a path of length at most $\frac{11}{10} \cdot 40 = 44$ to a transmission vertex b_i (avoiding Parallel Gadgets), by Lemma 7.13, each marked vertex is displaced by at most $44\delta < 2^6\delta$.

Let p be the reflection of e_1 across $h f_1$, and q the reflection of e_3 across $h f_3$; we will verify that $p = q$, so that e_2 has a well-defined position at the correct distances from h , f_1 , and f_2 simultaneously. Equality $\theta_2 = (\theta_1 + \theta_3)/2$ implies that $\alpha_2 = \beta_2$ and therefore kites $h e_1 f_1 p$ and $f_3 e_3 h q$ are similar. Label their angles as $\alpha'_1 := \angle e_1 h p = \angle e_3 f_3 q$ and $\alpha'_3 := \angle p f_1 e_1 = \angle q h e_3$. We have $2\pi = \angle f_1 e_1 h + \angle e_1 h e_3 + \angle h e_3 f_3 = \angle e_1 h e_3 + 2\alpha_2$ because $f_1 e_1$ and $e_3 f_3$ are parallel, so we may compute

$$\begin{aligned} \angle p h q &= \angle e_1 h e_3 - \angle e_1 h p - \angle q h e_3 \\ &= (2\pi - 2\alpha_2) - \alpha'_1 - \alpha'_3 \\ &= 0, \end{aligned}$$

meaning $p = q$, as claimed.

Each of h and f_3 has been displaced by at most $2^6\delta$ as above, which is less than $\frac{1}{32} \min(|h - e_2|, |e_2 - f_3|) = 1/8$, so the Hook Linkage Lemma shows that e_2 's forced location of $p = q$ is at most $2^8\delta$ from its initial position and that $h e_2$ and $e_2 f_3$ have each rotated by at most $2^6\delta$ from their original directions (since $|h - e_2|$ and $|e_2 - f_3|$ are at least 4). Triangle $h k_2 e_2$ must be rotated by this same amount to match its hypotenuse, so k_2 is displaced by at most $|h - k_2| \cdot 2^6\delta = 2^9\delta$.

Similarly, the Hook Linkage Lemma applied to he_2f_1 shows that edge e_2f_1 has been rotated by at most $2^6\delta$, so the same is true for edges k_1f_1 and k_1e_2 . So k_1 's total displacement may be bounded by e_2 's displacement plus $|e_2 - k_1| \cdot 2^6\delta$, namely $(2^8 + 2^{10})\delta \leq 2^{11}\delta$.

We have thus displaced all marked vertices no farther than $2^{11}\delta \leq \varepsilon/2^7$ and rotated each edge (other than those in Parallel Gadgets) by no more than $2^6\delta < \varepsilon/2$. We may thus configure the Parallel Gadgets and then verify angle constraints and minimum feature size exactly as for prior gadgets, which completes the analysis. $\square$




Figure 13: The Angle Average Gadget forces $\theta_2 = (\theta_1 + \theta_3)/2$. This gadget is analyzed in Lemma 7.19. See this gadget's movement in Figure 14. The drawing follows Convention 7.4. Unfrozen sliceform vertices b_j are assigned tolerance δ , while all remaining unfrozen corners have tolerance ε . The pins at the vertices a_j are for clarification only; in the overall construction, these nodes a_j are forbidden from moving by other means, so these explicit pins are unnecessary. The thick edges emphasize the rigidified background grid (in gray) and the transmission edges (in black), along with everything rigidly attached to those edges. Dotted segments and dashed angles are referred to in the proof. The cell has dimensions $Q \times Q$, and all marked vertices have coordinates at integer multiples of $Q/10$, indicated by cyan axis notches.

The next batch of cell gadgets are not scale invariant: they may have minimum feature size less than $1/2$ when scaled down by $40/Q$, unlike the previous gadgets, so we cannot assume $Q = 40$ as above. We must therefore skip the simplifying scaling step, work instead with the value of Q given in Section 7.3, and be more careful when arguing about feature size.




Figure 14: Snapshots of the Angle Average Gadget's configuration space (cf. Lemma 7.19, Figure 13), shown for each pair of values $\theta_1, \theta_3 \in \{-8^\circ, 0^\circ, 8^\circ\}$ (where $\theta_j = \text{Offset}_{\Theta_j}$). In each configuration, it may be observed that $\theta_2 = (\theta_1 + \theta_3)/2$. Edge lengths have been altered from those in Figure 13 to exaggerate the gadget's movement. The thick colored edges emphasize the transmission edges, and non-black thin edges are colored to indicate which transmission edge controls their angle.

Lemma 7.20 (Cell Gadgets with Unscaled Measurements). *The abstract extended linkages described below are globally noncrossing with global minimum feature size at least $1/2$, and their configuration spaces are exactly parameterized as described. Each gadget is described by an initial configuration having $a_3 = (0, 0)$ and $a_1 = (Q, Q)$, and where all marked vertices have coordinates that are integer multiples of $Q/10$, with a few exceptions listed below.*

Start Gadget. *The Start Gadget $\mathcal{L}_{\text{start}}$ (Figure 15a) relates rectangular and angular coordinates just like the Angular Gadget, but this time the angular coordinates use radius $2r$ instead of R . Recall $r = \lceil d/\delta \rceil$ is half the radius used in the angular change of coordinates $(x, y) = 2r \cdot \text{Rect}(\alpha, \beta)$ applied to each polynomial and that $r \leq Q/20$. The function $\pi_v \times \text{Offset}_{\Lambda, \Gamma}$ is a homeomorphism from $\text{Conf}(\mathcal{L}_{\text{start}})$ to*

$$U_{\text{start}} = \{((x, y), (\alpha, \beta)) \mid \alpha, \beta \in [-\delta, \delta] \text{ and } (x, y) = (Q/5 + 2r, Q/5 + 2r) + 2r \text{Rect}(\alpha, \beta)\}.$$

Exceptional vertices are initially configured at points $u = (Q/5, Q/5) + (2r, 0)$, $v = (Q/5, Q/5) + (2r, 2r)$, and $w = (Q/2, Q/5) + (0, 2r)$, which all have integer coordinates.

Vector Creation Gadget. *For a positive number $1 \leq w \leq R\delta/2$ (not necessarily an integer), the Vector Creation Gadget $\mathcal{L}_{\text{create}}(w)$ (Figure 15b) enables constructing angular coordinates for the vector with complex representation $w \cdot (e^{i\theta} - 1)$ for all $\theta \in [-\delta, \delta]$. Specifically, $\text{Offset}_{(\Theta, \Lambda, \Gamma)}$ provides a homeomorphism from $\text{Conf}(\mathcal{L}_{\text{create}}(c))$ to*

$$U_{\text{create}} = \{(\theta, \alpha, \beta) \in [-\delta, \delta]^3 \mid w \cdot (\cos \theta, \sin \theta) - (w, 0) = R \cdot \text{Rect}(\alpha, \beta)\}.$$

In the initial configuration shown in Figure 15b, the only exceptional vertex has position $c_3 = (Q/2 - w, Q/2)$, which has integer coordinates when w itself is an integer.

End Gadget. *For any chosen real number w with magnitude at most $R\delta/2$, the End Gadget $\mathcal{L}_{\text{end}} = \mathcal{L}_{\text{end}}(w)$ (Figure 15c) constrains a vector to the position $(Q/2 + w, Q/2)$. Specifically, $\text{Conf}(\mathcal{L})$ has exactly one configuration, and this configuration has $R \cdot \text{Rect}(\alpha, \beta) = (w, 0)$. When $w = 0$, there are no exceptional vertices. (When $w \neq 0$, we make no promises about vertices having integer coordinates other than those frozen to the background grid.)*

Proof. We consider each gadget in turn.

Start Gadget. This gadget is simply the Angular Gadget from Lemma 7.19 amended with three vertices u, v, w forming a parallelogram with f .

Constraint and Uniqueness. Vertices u and w must move rigidly with assemblies ef and fg , so their positions are determined by rotations by α and β . By Lemma 7.16, vertex v must complete the parallelogram, so its position is likewise uniquely determined as $e + 2r \exp(i\alpha) + 2r \cdot i \cdot \exp(i\beta) = e + (2r, 2r) + 2r \cdot \text{Rect}(\alpha, \beta)$.

Existence. As in the Angular Gadget, each bar is rotated by at most $\pm\epsilon/2$, so all angle constraints are satisfied. It remains to show that the linkage is noncrossing with minimum feature size at least $1/2$. The only differences from the Angular Gadget are vertices u, v, w and edges uv, vw , so if a crossing or smaller feature size occurs, it must involve one of these




Figure 15: A selection of extended linkage gadgets for manipulating angles and vectors. These gadgets do not scale uniformly with Q , and they are analyzed in Lemma 7.20. In these figures, $\text{Offset}_\Lambda(C) = \alpha$ and $\text{Offset}_\Gamma(C) = \beta$. As described in Convention 7.4, vertices surrounded by squares are pinned; those marked with an “x” are sliceform vertices; and corners marked with a solid gray sector are frozen, i.e., have tolerance 0. However, unfrozen sliceform vertices b_j are assigned tolerance δ , while all remaining corners have tolerance ε . The pins at the vertices a_j are for clarification only; in the overall construction, these nodes a_j are forbidden from moving by other means, so these explicit pins are unnecessary. Each cell has dimensions $Q \times Q$, and all vertices except those specified in Lemma 15 have coordinates at integer multiples of $Q/10$, indicated by cyan axis notches.

five items. Because $2r \leq R/2$, the smallest feature introduced by this parallelogram is at least

$$|u - v| \cdot |\sin \angle vuc| \geq 2r \cdot \sin \left(\frac{\pi}{2} - \varepsilon \right) \geq r > \frac{1}{2},$$

as required.

Vector Creation Gadget. *Constraint and Uniqueness.* Rigid assemblies $b_2c_2d_2e_2$ and $c_3gd_3e_3$ must have been rotated by angle θ , meaning $g = c_3 + w \cdot (\cos \theta, \sin \theta)$. The rest is an instance of the Angular Gadget, which uniquely positions all other marked vertices and locates g at $(Q/2, Q/2) + R \cdot \text{Rect}(\alpha, \beta)$. These two positions of g must agree, proving that

$$R \cdot \text{Rect}(\alpha, \beta) = w \cdot (\cos \theta, \sin \theta) - w \cdot (1, 0),$$


as required.

Existence. Each marked vertex is determined by a rigidified path (from a vertex rigidly attached to the rigidified background grid) of total length at most $\frac{6}{10}Q < Q$ where each edge rotates by at most $\pm\delta$, so each marked vertex is displaced by at most $\frac{6}{10}Q\delta < \varepsilon/2^7 \cdot Q/40$ by Lemma 7.13. Then Lemma 7.17 (Parallel Gadget) shows that each Parallel Gadget (scaled up by $Q/40$) can connect its opposite edges such that each vertex moves by at most $\varepsilon/2 \cdot Q/40$ and each edge rotates by at most $\pm\varepsilon/2$. Any feature not involving c_3 has length at least $\frac{1}{2}Q/40$ by Lemma 7.14, and c_3 's smallest feature is the distance from c_3 to fg , which is at least $|c_3 - g| \cdot \sin(\pi/2 - \varepsilon) \geq w/2 \geq 1/2$.


End Gadget. This is an Angular Gadget with one additional edge keeping g stationary at $g = (Q/2 + w, Q)$. Because this chosen point is at most $R\delta/2$ from the cell's center, the Angular Gadget case of Lemma 7.19 guarantees that the required configuration exists and is unique. When $w = 0$, this unique configuration agrees with the Angular Gadget's initial configuration. $\square$

Finally, we demonstrate a few *compound* gadgets, each inhabiting a constant number of grid cells instead of just one. As described in Section 7.2, the background grid is pinned in place and consists of $Q \times Q$ cells with *transmission edges* incident to *transmission sliceform vertices* at the midpoints of the sides of each cell. Neighboring cells communicate through these sliceform vertices, and similarly, gadgets that are *not* intended to interact must be placed in non-adjacent cells so as to not share a transmission vertex. Each transmission vertex will either be adjacent to one or more gadgets (meaning the angle at its transmission edges is determined by those gadgets) or it will be frozen, so there is no possibility for movement away from the gadgets. In the case of the Vector Average Gadget, a multi-cell cavity is carved into the background grid to make way for a larger assembly: cell edges that are not drawn are indeed missing from the grid, but all remaining cell edges have a transmission vertex and transmission edges at their midpoint, as usual.


Lemma 7.21 (Compound Gadgets). *The abstract extended linkages described below are globally noncrossing with global minimum feature size at least $1/2$, and with their configuration spaces U exactly parameterized as specified in each case. Each gadget is described by an initial configuration where the lower-left corner is configured at $(0, 0)$, all marked vertices*



(a) The Vector Average Gadget forces $g_2 = (g_1 + g_3)/2$, i.e., $\text{Rect}(\alpha_2, \beta_2) = (\text{Rect}(\alpha_1, \beta_1) + \text{Rect}(\alpha_3, \beta_3))/2$.



(b) The Angle Sum Gadget forces $\theta_3 = \theta_1 + \theta_2$.



(c) The Vector Sum Gadget forces $\text{Rect}(\alpha_2, \beta_2) = \text{Rect}(\alpha_1, \beta_1) + \text{Rect}(\alpha_3, \beta_3)$

Figure 16: Compound gadgets for manipulating vectors and angles, built from earlier, more primitive gadgets. In these figures, $\text{Offset}_\Lambda(C) = \alpha$, $\text{Offset}_\Gamma(C) = \beta$, $\text{Offset}_\Theta(C) = \theta$, and similarly with indices. Cells labelled with an angle are Copy Gadgets, those marked with an “x” are Crossover Gadgets, and cells labelled A_j are Angle Average Gadgets. All grid edges (where present) have a transmission sliceform and transmission edges, but only important ones are shown here.

have coordinates that are integer multiples of $Q/10$, and so all vertices have coordinates that are integer multiples of $Q/40$ (an integer). As in Lemma 7.19, for some gadgets, we prove that the configuration space U contains a conveniently shaped subset, identified below as a [useful neighborhood](#).

Vector Average Gadget. The Vector Average Gadget $\mathcal{L}_{\overrightarrow{\text{avg}}}$ (Figure 16a) constrains one vector to be the average of two others. Specifically, the function $\text{Offset}_{(\Lambda_1, \Gamma_1, \Lambda_2, \Gamma_2, \Lambda_3, \Gamma_3)}$ is a homeomorphism from $\text{Conf}(\mathcal{L}_{\overrightarrow{\text{avg}}})$ to

$$U_{\overrightarrow{\text{avg}}} = \{(\alpha_1, \beta_1, \alpha_2, \beta_2, \alpha_3, \beta_3) \in [-\delta, \delta]^6 \mid 2 \text{Rect}(\alpha_2, \beta_2) = \text{Rect}(\alpha_1, \beta_1) + \text{Rect}(\alpha_3, \beta_3)\}.$$

Useful Neighborhood: whenever three vectors satisfying $v_1 + v_3 = 2v_2$ are given with magnitudes at most $|v_j| \leq \delta/2$ for $j = 1, 2, 3$, there exists a (necessarily unique) configuration of $\mathcal{L}_{\overrightarrow{\text{avg}}}$ with $\text{Rect}(\alpha_j, \beta_j) = v_j$ for $j = 1, 2, 3$.

Angle Sum Gadget. The Angle Sum Gadget $\mathcal{L}_{\angle \text{sum}}$ drawn in Figure 16b constrains one sliceform angle to equal the sum of the other two: $\text{Offset}_{(\Lambda_1, \Lambda_2, \Lambda_3)}$ is a homeomorphism from $\text{Conf}(\mathcal{L}_{\angle \text{sum}})$ to

$$U_{\angle \text{sum}} = \{(\theta_1, \theta_2, \theta_3) \mid \theta_1, \theta_2, \theta_3 \in [-\delta, \delta] \text{ and } \theta_3 = \theta_1 + \theta_2\}.$$

Vector Sum Gadget. The Vector Sum Gadget $\mathcal{L}_{\overrightarrow{\text{sum}}}$ (Figure 16c) constrains one vector to be the sum of two others. Specifically, $\text{Offset}_{(\Lambda_1, \Gamma_1, \Lambda_2, \Gamma_2, \Lambda_3, \Gamma_3)}$ is a homeomorphism from $\text{Conf}(\mathcal{L}_{\overrightarrow{\text{sum}}})$ to

$$U_{\overrightarrow{\text{sum}}} = \{(\alpha_1, \beta_1, \alpha_2, \beta_2, \alpha_3, \beta_3) \in [-\delta, \delta]^6 \mid \text{Rect}(\alpha_3, \beta_3) = \text{Rect}(\alpha_1, \beta_1) + \text{Rect}(\alpha_2, \beta_2)\}.$$

Useful Neighborhood: whenever three vectors satisfying $v_1 + v_2 = v_3$ are given with magnitudes at most $|v_j| \leq \delta/2$ for $j = 1, 2, 3$, there exists a (necessarily unique) configuration of $\mathcal{L}_{\overrightarrow{\text{sum}}}$ with $\text{Rect}(\alpha_j, \beta_j) = v_j$ for $j = 1, 2, 3$.

Note. The Vector Average Gadget uses the well-known *pantograph linkage* to keep g_2 at the midpoint of g_1 and g_3 .

Proof. We consider each gadget in turn.

Vector Average Gadget. The Vector Average Gadget, $\mathcal{L}_{\overrightarrow{\text{avg}}}$, illustrated in Figure 16a, is formed within a 3×11 rectangle of grid cells, where some grid edges have been removed as shown. Three Angular Gadgets (missing their top and left cell edges) are placed with lower-left corners at $(0, 0)$, $(4Q, Q)$, and $(8Q, 2Q)$, and their vertices g_1, g_2, g_3 are connected to parallelogram $defg_2$ as shown. The angles α_j and β_j ($j = 1, 2, 3$) at the three Angular Gadgets are duplicated by Copy Gadgets and output along the bottom edge. Observe that this gadget is scale invariant, like the gadgets from Lemma 7.19, so we may assume $Q = 40$ in our analysis.

Constraint and Uniqueness. Temporarily ignoring d, e, f , and their edges, the Angular and Copy Gadgets ensure that the six angles α_j and β_j are each free to move within $[-\delta, \delta]$,

and together they uniquely determine the positions of all vertices (other than d , e , and f). Lemma 7.15 (Hook Linkage) shows that the positions of d , e , and f are determined by g_1, g_2, g_3 . Furthermore, parallelogram $defg_2$ (which remains a parallelogram by Lemma 7.16) ensures that triangle g_1dg_2 remains similar to g_1eg_3 , so g_2 must lie at the midpoint of g_1g_3 . This is equivalent to $\text{Rect}(\alpha_2, \beta_2)$ being the average of $\text{Rect}(\alpha_1, \beta_1)$ and $\text{Rect}(\alpha_3, \beta_3)$, as required.

Existence. Suppose angles α_j and β_j are chosen for $j = 1, 2, 3$ satisfying this average constraint, i.e., lying in U_{avg} . As in Lemma 7.19 (the Copy Gadget and Angular Gadget cases), all vertices other than d, e, f may be configured at most $\varepsilon/2$ from their initial positions such that all edges other than those incident to d, e, f have rotated by at most $\pm\varepsilon/2$ from their initial directions. Membership in U_{avg} guarantees that g_2 is the midpoint of g_1g_3 . Lemma 7.15 (Hook Linkage) applied to g_1dg_2 , which is applicable because $\max((\varepsilon/2)/|g_1 - d|, (\varepsilon/2)/|d - g_2|) = \varepsilon/80 < 1/32$, shows that d may be configured at most $4 \cdot \varepsilon/2 = 2\varepsilon$ from its initial position such that edges g_1d and dg_2 have each rotated by at most $\max((\varepsilon/2)/|g_1 - d|, (\varepsilon/2)/|d - g_2|) \leq \varepsilon/20 < \varepsilon/2$ from their initial directions. Similar computations apply to hooks h_2fg_3 and g_1eg_3 , and because g_2 is the midpoint of g_1g_3 , similar triangles ensure that d and f land at the midpoints of g_1e and g_3e , respectively. The result is thus a valid configuration of $\mathcal{L}_{\text{avg}}$ with minimum feature size at least $1 - 4\varepsilon \geq 1/2$.

Useful Neighborhood. This final claim follows from the Angular Gadget case of Lemma 7.19.

Angle Sum Gadget. The Angle Sum Gadget in Figure 16b is constructed in a 6×7 grid of $Q \times Q$ cells, with transmission sliceforms and transmission edges as described immediately preceding this Lemma (most not depicted). The cells labeled A_1 and A_2 are Angle Average Gadgets ($\mathcal{L}_{\text{avg}}$), where one of A_2 's transmission sliceforms (drawn in gray) has been frozen, fixing that input to 0. The other gray cells are Copy Gadgets, $\mathcal{L}_{\text{copy}}$. This analysis follows straightforwardly from the characterizations of the Angle Average and Copy Gadgets from Lemma 7.19, as we now show. Define $\theta'_j = \theta'_j(C) = \text{Offset}_{\Theta'_j}(C)$ for a given configuration C , analogous to the definitions of θ_j .

Constraint and Uniqueness. Suppose $\theta_1, \theta_2, \theta_3 \in [-\delta, \delta]$ are given. We must show that they give rise to at most one configuration of $\mathcal{L}_{\text{sum}}$, and that $\theta_3 = \theta_1 + \theta_2$ must hold for a configuration to exist. The Copy Gadgets labeled θ_1 and θ_2 show that $\theta'_1 = \theta_1$ and $\theta'_2 = \theta_2$. Then gadget A_1 forces $\theta_4 = (\theta_1 + \theta_2)/2$. Copy Gadgets transfer this same value to $\theta'_4 = \theta_4 = (\theta_1 + \theta_2)/2$, and then gadget A_2 and the remaining Copy Gadgets ensure that $\theta'_4 = (\theta_1 + \theta_2)/2$ is the average of 0 and θ_3 , or in other words, $\theta_3 = \theta_1 + \theta_2$. All copy and Angle Average Gadget configurations are uniquely determined by these values. Because transmission edges not part of a gadget on either side are frozen, angles $\theta_1, \theta_2, \theta_3$ uniquely determine the entire configuration of $\mathcal{L}_{\text{sum}}$ (when it exists), as required.

Existence. The conditions $\theta_j \in [-\delta, \delta]$ for $j = 1, 2, 3$ show that $\theta_4 = (\theta_1 + \theta_2)/2$ also lies in $[-\delta, \delta]$, and the condition $\theta_3 = \theta_1 + \theta_2$ ensures that the inputs at A_1 and A_2 satisfy the average constraint required by each Angle Average Gadget, so the desired configuration indeed exists by Lemma 7.19.

Vector Sum Gadget. In this gadget illustrated in Figure 16c, each $Q \times Q$ cell labeled with an angle α_j or β_j is a Copy Gadget, and each cell with an “x” is a Crossover Gadget. As with

the previous compound gadgets, all unused transmission edges are frozen, as are the two transmission sliceforms illustrated in gray instead of black at the right of the bottom-most Vector Average Gadget. The analysis is similar to that of the Angle Sum Gadget, so we omit the details. $\square$

7.6 Combining the Gadgets

We are now prepared to use these gadgets to construct an extended linkage that draws a piece of $Z(F)$, up to a translation. Recall that $F = \{f_1, \dots, f_s\}$ is a family of polynomials in $\mathbb{R}[x_1, y_1, \dots, x_m, y_m]$, each with total degree at most d . We apply the change of coordinates $(x_k, y_k) = 2r \cdot \text{Rect}(\alpha_k, \beta_k)$ (for $1 \leq k \leq m$) to write each polynomial f_j (for $1 \leq j \leq s$) in angular form as in Lemma 7.10:

$$f_j(\vec{xy}(\vec{\alpha\beta})) - f_j(\vec{0}) = \sum_{u=0}^3 \sum_{I \in \text{Coeffs}(2m, d)} i^u \cdot d_{j,u,I} \cdot \left(\exp(i \cdot (I \cdot \vec{\alpha\beta})) - 1 \right), \quad (8)$$

where the numbers $d_{j,u,I}$ are nonnegative. If the polynomials in F do not have integer coefficients, we may need to modify the input slightly: scale up F uniformly (which scales the $d_{j,u,I}$ coefficients by the same amount) until each nonzero term $d_{j,u,I}$ is at least 1, and let M be an upper bound on the scaled up coefficients in F (as standard polynomials in $\vec{xy}$). If F originally had integer coefficients, the $d_{j,u,I}$ were already integers, so this scaling is unnecessary.

For each $1 \leq j \leq s$, the numbers $d_{j,u,I}$ are nonnegative and add to at most $6^d \cdot r^d \cdot M \cdot \binom{2m+d}{d} \leq R\delta/2$, by Lemma 7.10 and by the choice of R in Section 7.3. This shows that every partial sum of the right-hand side of (8) has magnitude at most $R\delta/2$. In particular, the Vector Creation, Vector Rotation, Vector Average, and Vector Sum Gadgets can all safely construct these terms and partial sums, according to the gadget specifications in Lemmas 7.19 through 7.21. The same holds for the value $f_j(\vec{0})$ and the End Gadget, because $|f_j(\vec{0})| \leq M \leq R\delta/2$.

The extended linkage $\mathcal{E} = \mathcal{E}(F)$ that will be built in this section will draw a translation of $Z(F) \cap (2r \cdot \text{Rect}([- \delta/d, \delta/d]^2))^m$, which by Lemma 7.9 contains $Z(F) \cap [-1, 1]^2$ because $2r \cdot \delta/(2d) \geq 1$ by choice of r .

Step 1: Grid. To begin, define extended linkage $\mathcal{E}$ as a rectangular grid of $Q \times Q$ cells, with $O(\text{poly}(m^d, d^d, s))$ cells on each side. Each cell is given tolerance-0 corners and *transmission edges* attached at cell edge midpoints. The transmission vertices (which are sliceforms) initially have frozen corners. Add pins to three noncollinear vertices of this grid, so the entire grid $\mathcal{E}$ is globally rigid. There will be no other pins anywhere in $\mathcal{E}$; pins illustrated in the gadgets above are removed before use here, though they are *effectively* still present because the grid $\mathcal{E}$ is globally rigid.

In the rest of the construction, gadgets will be added to this grid, with their transmission vertices' corners upgraded to tolerance- δ as required by the gadgets.

Step 2: Start Gadgets to set up $\alpha_1, \beta_1, \dots, \alpha_m, \beta_m$. Add m Start Gadgets $\mathcal{L}_{\text{start}}(k)$, $1 \leq k \leq m$ to $\mathcal{E}$ as illustrated at the top of Figure 17, and let v_k be the vertex in $\mathcal{L}_{\text{start}}(k)$ corresponding to vertex v in Figure 15a; these vertices $X = \{v_1, \dots, v_m\}$ will be the drawing vertices of $\mathcal{E}$. Likewise define corners Λ_k and Γ_k as the corners from $\mathcal{L}_{\text{start}}(k)$ that are labelled Λ and Γ in Figure 15a, and set $(\alpha_k, \beta_k) = (\alpha_k(C), \beta_k(C)) = \text{Offset}_{(\Lambda_k, \Gamma_k)}(C)$ for a configuration C , as usual.

With $\mathcal{E}$ in this state, by Lemma 7.20 (Start Gadget), $\text{Conf}(\mathcal{E})$ is perfectly described by corners $Y = (\Lambda_1, \Gamma_1, \dots, \Lambda_m, \Gamma_m)$: the map Offset_Y is a homeomorphism of $\text{Conf}(\mathcal{E})$ with $[-\delta, \delta]^{2m}$.




Figure 17: Constructing the sum $\gamma = \beta_1 + \alpha_2$. Cells with “S” are Start Gadgets; those with “X” are Crossover Gadgets; and those with α_k, β_k , or γ are Copy Gadgets.

Step 3: Construct all linear combinations θ_I . In the next step, we modify $\mathcal{E}$ further to construct all the angles $\theta_I := I \cdot \vec{\alpha\beta}$, where $\vec{\alpha\beta} = (\alpha_1, \beta_1, \dots, \alpha_m, \beta_m)$ and I ranges over all vectors of $\text{Coeffs}(2m, d)$.

We start with base cases of the angle 0 (corresponding to $I = (0, \dots, 0)$) and the angles α_k and β_k (corresponding to $I = (0, \dots, 0, 1, 0, \dots, 0)$). The angle 0 can be constructed by placing a 0 angle constraint on any transmission edge. The angles α_k and β_k are already constructed by the Start Gadgets. Use Copy Gadgets to build *wires* that copy all of these base cases vertically along columns, as shown in Figure 17. We call this collection of wires the *wire column*, and it will accumulate other wires constructing other angles as the construction continues.

By ordering all the remaining vectors $I \in \text{Coeffs}(2m, d)$ according to the sum of the absolute values of entries in I , each successive θ_I may then be computed as $\theta_I = \theta_{I'} \pm \alpha_k$ or $\theta_I = \theta_{I'} \pm \beta_k$, which can be constructed using a single Addition Gadget. To do this, as illustrated in Figure 17, we use Crossover Gadgets to transmit the desired input angles to the Addition Gadget, and then the constructed θ_I joins the wire column. Repeat this for each remaining $I \in \text{Coeffs}(2m, d)$.

For each I , let Θ_I be any transmission corner in the wire coming from the construction of θ_I above. We claim that $\text{Conf}(\mathcal{E})$ is still perfectly described by the corners $Y = (\Lambda_1, \Gamma_1, \dots, \Lambda_m, \Gamma_m)$, but this time on a smaller range: we will show that Offset_Y is a homeomorphism of $\text{Conf}(\mathcal{E})$ with $[-\delta/d, \delta/d]^{2m}$, and furthermore that the corners Θ_I indeed have angle $\text{Offset}_{\Theta_I}(C) = \theta_I(C)$. The latter claim follows from induction on I . To see that the α_k and β_k are each precisely (and independently) constrained to the interval $[-\delta/d, \delta/d]$, first consider the corner Θ_I where $I = (d, 0, 0, \dots, 0)$. We have $\theta_I = d \cdot \alpha_1$, and the δ -tolerance at corner Θ_I requires this value to remain in $[-\delta, \delta]$, so α_1 must remain in $[-\delta/d, \delta/d]$. The other entries in $\vec{\alpha\beta}$ are likewise constrained by analogous arguments. In the other direction, for any vector $\vec{\alpha\beta} \in [-\delta/d, \delta/d]^{2m}$ and any $I \in \text{Coeffs}(2m, d)$, the value $I \cdot \vec{\alpha\beta}$ lies in the interval $[-\delta, \delta]$ by triangle inequality. Thus, by our gadget analyses (Lemmas 7.19, 7.20, and 7.21), any such vector $\vec{\alpha\beta}$ may be realized by a unique configuration of $\mathcal{E}$.

Step 4: Construct the terms in the angular expansions of $f_1, \dots, f_s$. For each $1 \leq j \leq s$ and each nonzero coefficient $d_{j,u,I}$ in the angular representation of $f_j(\vec{xy}(\vec{\alpha\beta}))$, insert a Vector Creation Gadget using angle θ_I and length $w = d_{j,u,I}$ (which indeed satisfies $1 \leq |w| \leq R\delta/2$), and if $u \neq 0$, chain this with u Vector Rotation Gadgets. The result constructs the desired vector $i^u \cdot d_{j,u,I} \cdot (\exp(i \cdot \theta_I) - 1)$: specifically, the result is a pair of wires with transmission corners $\Lambda'_{j,u,I}$ and $\Gamma'_{j,u,I}$ whose angles $(\alpha'_{j,u,I}(C), \beta'_{j,u,I}(C)) = \text{Offset}_{\Lambda'_{j,u,I}, \Gamma'_{j,u,I}}(C)$ satisfy

$$R \cdot \text{Rect}(\alpha'_{j,u,I}, \beta'_{j,u,I}) = i^u \cdot d_{j,u,I} \cdot (\exp(i \cdot \theta_I) - 1).$$

These additional gadgets do not constrain $\mathcal{E}$'s movement in any way, so the map Offset_Y is still a homeomorphism of $\text{Conf}(\mathcal{E})$ with $[-\delta/d, \delta/d]^{2m}$.

Step 5: Add the vectors to construct $f_j(\vec{\alpha\beta}) - f_j(\vec{0})$. For each $1 \leq j \leq s$, use Vector Addition Gadgets to successively construct the sum of the nonzero vectors among $R \cdot \text{Rect}(\alpha'_{j,u,I}, \beta'_{j,u,I})$, resulting in a single pair of wires with transmission corners Λ'_j, Γ'_j whose angle offsets α'_j, β'_j satisfy

$$R \cdot \text{Rect}(\alpha'_j, \beta'_j) = \left(f_j(\vec{\alpha\beta}) - f_j(\vec{0}), 0 \right).$$

(Join these wires to the wire column, as usual.) As above, these additional gadgets do not constrain $\mathcal{E}$'s movement, so $\text{Offset}_Y : \text{Conf}(\mathcal{E}) \rightarrow [-\delta/d, \delta/d]^{2m}$ is still a homeomorphism.

Step 6: Conclude with End Gadgets Finally, for each $1 \leq j \leq s$, feed the α'_j and β'_j wires into an End Gadget that enforces the constraints

$$R \cdot \text{Rect}(\alpha'_j, \beta'_j) = \left(-f_j(\vec{0}), 0 \right).$$

The configuration space of the resulting linkage $\mathcal{E}$ is then homeomorphic to the subset of $\vec{\alpha\beta} \in [-\delta/d, \delta/d]^{2m}$ that satisfies the constraints of these End Gadgets, i.e.,

$$\text{Conf}(\mathcal{E}) = \left\{ \vec{\alpha\beta} \in [-\delta/d, \delta/d]^{2m} \mid f_j(\vec{\alpha\beta}) = 0, 1 \leq j \leq s \right\}.$$




Figure 18: The Sliceform Gadget keeps w_1, v, w_3 and w_2, v, w_4 collinear. All marked vertices have coordinates at integer multiples of $|v - w_j|/4$, indicated by cyan axis notches.

And because the offset of each vertex v_k from the lower-left corner of its cell is precisely $(Q/5, Q/5) + (2r, 2r) + 2r \cdot \text{Rect}(\alpha_k, \beta_k) = (Q/5, Q/5) + (2r, 2r) + (x_k, y_k)$, these vertices $\{v_1, \dots, v_m\}$ indeed draw a translation of

$$Z(F) \cap (2r \cdot \text{Rect}([- \delta/d, \delta/d], [- \delta/d, \delta/d]))^m,$$

as required.

7.7 Sliceform and Angle Restrictor Gadgets

In Section 7.8, we will convert the extended linkage $\mathcal{E}(F)$ constructed above into a partially rigidified linkage (with no other constraints). In this section, we describe two gadgets to help with this process: the *Sliceform Gadget* obviates the need for sliceform vertices, and the *Angle Restrictor Gadget* enforces $\mathcal{E}$'s angle constraint, after replacing each edge with a rigidified tree.

Lemma 7.22 (Sliceform Gadget). *The extended linkage that has four edges meeting at a sliceform vertex with tolerance θ (where θ is in $\{0, \delta, \varepsilon\}$) is perfectly simulated by the extended linkage $\mathcal{L}_{\angle\text{slice}}(\theta)$, where all non-frozen corners are given tolerance θ (see Figure 18). Furthermore, the global minimum feature size is at least $|v - w_j|/8$ (which is the same for all $j = 1, 2, 3, 4$ and in all configurations), and all vertices lie within distance $\frac{3}{2}|v - w_j|$ of v , in all configurations.*




Figure 19: Angle Restrictor Gadget, $\mathcal{L}_{\angle\text{restrict}}(n)$, shown in full (left) and closeup (right). Tree A is bold and black, with leaves at o and a_j , $j = 1, 2, 3$; B is bold and gray, with leaves at o and b_j , $j = 1, 2, 3, 4, 5$. Coordinates depend on input integer n such that the range of motion of the gadget approaches zero as n increases. Dashed segments show alternative configurations for edges, with vertices constrained to dashed circles.

Proof. For convenience, scale so that $|v - w_1|$ is 4 in the initial configuration, so that the cyan notches in Figure 18 have unit distance between them, and all edge lengths are integers. By Lemma 7.16 applied to the many rectangles of Figure 18, all parallelograms in the figure remain parallelograms in all configurations. This, paired with the frozen angles keeping collinear edges collinear, ensures the desired collinearity and that the distances $|v - w_j|$ remain constant, namely 4. The gadget admits the desired motion by the simultaneous motion of all the parallelograms. All edge lengths are at least 1, so by Lemma 7.16, the global minimum feature size is at least $1/2 = |v - w_j|/8$. By triangle inequality, all vertices remain at distance at most $6 = \frac{3}{2}|v - w_j|$ from v . $\square$

The last gadget is built from *partially rigidified* linkages instead of *extended* linkages. Recall that a *partially rigidified* linkage is a constrained linkage where all constraints are *rigid constraints*, $\text{RigidCon}(H, C_H)$, each forcing a certain subgraph H to maintain a chosen, rigid shape given by a specified configuration C_H .

Lemma 7.23 (Angle Restrictor Gadget). *For an integer $n \geq n_\epsilon = 5000$ (from Equation 6), let $\mathcal{L}_{\angle\text{restrict}} = \mathcal{L}_{\angle\text{restrict}}(n)$ be the partially rigidified linkage shown in Figure 19, with two rigid constraints fixing orthogonal trees A and B in the configurations shown (but not governing their orientations or relative position). Tree A is drawn with bold black lines, while B has bold gray lines. Tree B is also pinned to the plane for convenience, with o at $(0, 0)$.*

Then $\mathcal{L}_{\angle\text{restrict}}$ is globally noncrossing, and vertex a_1 draws the locus

$$\{\ell(oa_1) \cdot (\cos \theta, \sin \theta) \mid \pi/2 - \gamma \leq \theta \leq \pi/2 + \gamma\}$$

liftably and rigidly, where $\gamma = \gamma(n)$ is given by

$$\gamma(n) = \cos^{-1} \left(1 - \frac{3}{10} \cdot \frac{2n}{n^2 + 1} \right).$$

Furthermore,

1. *The initial configuration shown in the figure has rational coordinates.*
2. *All configurations are noncrossing and agree with the same combinatorial embedding. In particular, tree A is always configured with the same orientation as in Figure 19.*
3. *At every vertex that is not interior to one of the rigidified trees (in other words, every vertex given a name in Figure 19), every pair of consecutive edges forms an angle strictly between 60° and 240° in every configuration, except for the external angle $\angle a_1 o b_1 \in [3\pi/2 - \gamma, 3\pi/2 + \gamma]$.*
4. *There is a unique configuration of $\mathcal{L}_{\angle\text{restrict}}$ having $\angle b_1 o a_1 = \pi/2$.*
5. *Finally, the limiting case $\mathcal{L}_{\angle\text{restrict}}(\infty)$ is well-defined, satisfies all of the above, and is globally rigid.*

Note. The assumption of pinning B to the plane is equivalent to choosing a coordinate frame relative to B 's configuration. Thus, if we do not pin B , then this lemma still describes tree A 's position relative to that of B . In particular, if either tree is known to have orientation matching that of Figure 19, then they both must—one cannot “flip” relative to the other—and $\angle b_1 o a_1$ must lie in $[\pi/2 - \gamma, \pi/2 + \gamma]$.




Figure 20: A simpler Angle Restrictor Gadget is possible if uniqueness at $\theta = \pi/2$ is not required. This is shown only for contrast with Figure 19; this simpler version is not formally analyzed or used in this paper.

Note. The uniqueness property of Part 4 is the crucial step in proving hardness of *global* rigidity, but it also complicates the gadget. Without this requirement, a much simpler construction is possible, such as that shown in Figure 20.

Proof. First we specify the initial configuration C_0 , drawn solid in Figure 19, more precisely. Vertex o rests at the origin, $(0, 0)$. Let $\kappa = \sin^{-1}(5/13)$, which is a *Pythagorean angle*, meaning $\sin(\kappa) = 5/13$ and $\cos(\kappa) = 12/13$ are both rational, and define another Pythagorean angle $\varphi = \varphi(n) = \sin^{-1}(2n/(n^2 + 1))$ (its cosine is $(n^2 - 1)/(n^2 + 1) \in \mathbb{Q}$), where $\varphi(n) \leq \varphi(n_\varepsilon) < 0.0004$. In the initial configuration, $a_2 = (\cos \kappa, \sin \kappa)$, $c_1 = 2a_2$, $\angle b_3 c_1 o = \angle o c_1 b_2 = \pi/2 - \varphi$, and d_1 is the midpoint of $c_1 b_2$. Edges $c_1 d_1$, $d_1 b_2$, and $c_1 b_3$ have length $t, t, 2t$ respectively, where $t = 3/20$. (This arbitrary choice of t allowed for a prettier figure than others we tried.) Vertices a_3, c_2, b_4, b_5, d_2 are initially configured at the reflections through the line $y = x$ of the initial positions of vertices a_2, c_1, b_3, b_2, d_1 respectively. The upper right vertex of B has coordinates $(2 + t, 2 + t)$. The initial configuration has minimum feature size realized by point d_1 to segment $a_2 c_1$ (and d_2 to $a_3 c_2$), namely $t \sin \angle a_2 c_1 d_1 = t \cos \varphi \geq t \cos \varphi(n_\varepsilon) > 0.14999$. Because we chose Pythagorean angles, all coordinates are initially rational. We will proceed as if $|o - a_1| = |o - b_1| = 2.75$ (as drawn), but they may be longer without affecting the global minimum feature size.

We now investigate a general configuration C of $\mathcal{L}_{\angle \text{restrict}}$. Bar $b_3 c_1$ restricts c_1 to the dashed circle with radius $2t$, and bars $b_2 d_1$ and $d_1 c_1$ further restrict vertex c_1 to the disk with radius $2t$, as shown. As a result, c_1 must lie on the small circular arc $g_1 g_2$ (see Figure 19), centered at b_3 with angle 2φ . (Note that g_1 and g_2 are fixed points in the plane, not linkage vertices.) In turn, because $|o - a_2| = |a_2 - c_1| = 1$, vertex a_2 is confined to the circular arc $h_1 h_2$ centered at o with radius 1, where the endpoints of confinement, h_1 and h_2 , are determined by rhombus $oh_1 g_1 h_2$ of side length 1. Because o , g_1 , and g_2 are collinear, we may compute that $|o - g_1| = 2 - 4t \sin \varphi$, and so $\gamma := \frac{1}{2} \angle h_1 o h_2$ is given by $\gamma = \cos^{-1}(1 - 2t \sin \varphi)$. We have $\gamma(n) \leq \gamma(n_\varepsilon) < 0.0155$, which is less than the fixed value $|\angle a_2 o a_3| = \pi/2 - 2\kappa > 0.78$, so arcs $h_1 h_2$ and $h_3 h_4$ are disjoint. This forces $\angle a_2 o a_3$ to have counterclockwise orientation, and so tree A must maintain its orientation as claimed, and $\theta = \angle b_1 o a_1$ must remain in the closed interval between $\pi/2 \pm \gamma$.

We must show that $\mathcal{L}_{\angle \text{restrict}}$ cannot move far enough to intersect itself or to violate the angle conditions in Part 3. Toward this goal, let us first argue that the angle of each edge in configuration C differs from that in C_0 by at most $\pm 1/60$ radians. The edges in tree A can rotate by at most $\pm \gamma$, and likewise edges $a_2 c_1$ and $a_3 c_2$ can change their direction by at most $\pm \gamma$, where $\gamma < 1/60$ as above. Edges $b_3 c_1$ and $b_4 c_2$ can rotate by at most $2\varphi \leq 2\varphi(n_\varepsilon) < 0.0008 < 1/60$. Only the edges incident to d_1 and d_2 remain, and by symmetry, we may consider only those at d_1 . We will measure these edges relative to the vector $b_2 c_1$, which is not an edge of the linkage but is nevertheless a useful reference. This vector $b_2 c_1$ can change its direction by at most $2\varphi \leq 2\varphi(n_\varepsilon) < 1/120$. Its length is at least $|b_2 - b_3| - 2t = 2t(2 \cos(\varphi) - 1) \geq 2t(2 \cos(\varphi(n_\varepsilon)) - 1)$, which may be checked to be at least $2t \cdot \cos \frac{1}{120}$. Isosceles triangle $b_2 d_1 c_1$ then shows that $|\angle d_1 b_2 c_1| = |\angle d_1 c_1 b_2| \leq 1/120$ radians. So each of $d_1 b_2$ and $d_1 c_1$ can rotate by at most $1/120 + 1/120 = 1/60$, as claimed.

Each vertex is connected to a vertex in the stationary tree B by a path of length at most 3, so by Lemma 7.13, each vertex has been offset from its initial position by at most

$3 \cdot 1/60 = 1/20$. Because C_0 has minimum feature size > 0.14999 , we conclude by Lemma 7.14 that C is noncrossing (with minimum feature size at least $0.14999 - 2/20 > 0.04999 > 0$). Each corner at each named vertex initially had an angle between $\pi/2 - \varphi$ and $\pi + \kappa + \varphi$, and these angles are at most $2 \cdot 1/60 = 1/30$ different in configuration C , which confirms that the corresponding angles in C lie safely between $\pi/3$ and $4\pi/3$: indeed,

$$\frac{\pi}{2} - \varphi - \frac{1}{30} > 88.06^\circ > 60^\circ \quad \text{and} \quad \pi + \kappa + \varphi + \frac{1}{30} < 204.56^\circ < 240^\circ,$$

confirming Part 3. It also proves that C and C_0 agree with the same combinatorial embedding (Part 2).

Conversely, given any angle $\theta \in [\pi/2 - \gamma, \pi/2 + \gamma]$, we will argue there are either 1 or 4 configurations C having $\angle b_1 o a_1 = \theta$. This angle determines the position of a_2 along arc $h_1 h_2$, and then there is a unique point c_1 on arc $g_1 g_2$ with $|a_2 - c_1| = 1$. If c_1 lies at g_2 (i.e., $\theta = \pi/2$), then because $|b_2 - c_1| = 2t$, vertex d_1 must be configured at the midpoint of $b_2 c_1$. The case is similar if c_1 is at g_1 , i.e., $\theta = \pi/2 \pm \gamma$. Otherwise, d_1 lies in the interior of the solid gray disk, so there are two possible choices for the location of d_1 . The same is true for the other assembly anchored at a_3 , b_4 , and b_5 , and furthermore, vertices d_1 and d_2 may vary continuously with θ . This proves that vertex a_1 draws liftably and rigidly.

In the limiting case $n = \infty$, the angles φ and γ are 0, so the resulting linkage $\mathcal{L}_{\angle\text{restrict}}(\infty)$ is globally rigid. $\square$

7.8 Implementing Extended Linkages with Partially Rigidified Linkages

We now use the gadgets from Section 7.7 to transform the extended linkage $\mathcal{E}$, constructed in Section 7.6, into a partially rigidified linkage $\mathcal{L} = \mathcal{L}(F)$ that simulates $\mathcal{E}$ —not perfectly, but at least liftably and rigidly.

Extended linkage $\mathcal{E}$ has all edge lengths at least 1 and has global minimal feature size at least $1/2$, because each gadget individually has these properties. Construct a new linkage $\mathcal{E}' = \mathcal{E}'(F)$ by replacing each sliceform vertex v in $\mathcal{E}$ with a Sliceform Gadget $\mathcal{L}_{\angle\text{slice}}$ of Lemma 7.22 with the same angle tolerance, scaled so that $|v - w_j| = 1/8$. By Lemma 7.22, each sliceform gadget individually has minimum feature size at least $|v - w_j|/8 = 1/64$ and lives within a disk of radius $\frac{3}{2}|v - w_j| = 3/16$ centered at the replaced vertex. Thus the feature size between different sliceform gadgets is at least $\frac{1}{2} - 2 \cdot \frac{3}{16} = \frac{1}{8}$. Therefore $\mathcal{E}'$ perfectly simulates $\mathcal{E}$, is globally noncrossing, and has global minimum feature size at least $1/64$.

We transform $\mathcal{E}'$ into a partially rigidified linkage $\mathcal{L} = \mathcal{L}(F)$ as follows: first, consider each edge uv of $\mathcal{E}'$ as a rigidified tree T_{uv} initially containing just one edge. For each vertex v with neighbors w_j (for $1 \leq j \leq \deg v$), we will modify the trees T_{vw_j} in a small neighborhood of v with (scaled down) Angle Restrictor Gadgets (from Lemma 7.23) adapted as necessary to v 's structure, to faithfully implement v 's angle constraints. Specifically, look at the angles $A(\Lambda)$ for corners Λ around v : the 5 possibilities, up to cyclic reordering, are $(90^\circ, 90^\circ, 90^\circ, 90^\circ)$, $(180^\circ, 90^\circ, 90^\circ)$, $(180^\circ, 180^\circ)$, $(270^\circ, 90^\circ)$, or (360°) . The last case corresponds to $\deg v = 1$, in which case its angle constraint does nothing and may be ignored. In each of the first four cases, we modify the trees T_{vw_j} in a small region around




Figure 21: To turn extended linkage $\mathcal{E}'$ into $\mathcal{L}$, a neighborhood of each vertex is replaced with an *Angle Restrictor Gadget neighborhood* tailored to the base angles of its corners, as shown. The possible vertex types are $(90^\circ, 90^\circ, 90^\circ, 90^\circ)$ (top left), $(180^\circ, 90^\circ, 90^\circ)$ (top right), $(180^\circ, 180^\circ)$ (bottom left), $(270^\circ, 90^\circ)$ (bottom right), and (360°) (not shown). In the bottom two cases of degree 2, a single Angle Restrictor Gadget would suffice, but we draw both for consistency. Each rigidified tree is a different color of thick edges.

v with the appropriate *Angle Restrictor Gadget neighborhood* from Figure 21. Each Angle Restrictor Gadget within this neighborhood is given the shape of $\mathcal{L}_{\angle\text{restrict}}(n_\varepsilon)$, $\mathcal{L}_{\angle\text{restrict}}(n_\delta)$, or $\mathcal{L}_{\angle\text{restrict}}(\infty)$ depending on whether the corner Λ has $\Delta(\Lambda) = \varepsilon$, δ , or 0 respectively. (Recall n_ε and n_δ are given by Equations 6–7.) Finally, remove the three pins of $\mathcal{E}'$ (that serve only to fix the background grid in place), and replace them in $\mathcal{L}$ with three noncollinear pins in one of the rigidified trees $T_{u_1u_2}$ built from an edge u_1u_2 in the background grid.

We claim that $\mathcal{L}$ simulates $\mathcal{E}'$ (and hence $\mathcal{E}$) liftably and rigidly. For each rigidified tree T_{uv} , at least one of u or v has degree greater than 1 in $\mathcal{E}'$ because $\mathcal{E}'$ is connected, so T_{uv} participates in at least one Angle Restrictor Gadget neighborhood. As shown in Lemma 7.23 (Angle Restrictor Gadget), rigidified trees connected to each other by an Angle Restrictor Gadget cannot change orientation relative to each other. Because tree $T_{u_1u_2}$ is pinned rigidly in place and all rigidified trees in $\mathcal{L}$ are connected to each other by a sequence of Angle Restrictor Gadgets, none of the rigidified trees in $\mathcal{L}$ can change orientation at all. This same lemma guarantees that all angle constraints are enforced by the Angle Restrictor Gadgets. In particular, the entire background grid in $\mathcal{L}$ is fixed rigidly in place by the pins in $T_{u_1u_2}$. The Angle Restrictor Gadgets maintain their constraints in a liftable and rigid fashion, so $\mathcal{L}$ indeed simulates $\mathcal{E}'$ liftably and rigidly.

Below we will show that this linkage $\mathcal{L} = \mathcal{L}(F)$ satisfies all requirements of the Main Theorem.

Theorem 7.24. *This linkage $\mathcal{L} = \mathcal{L}(F)$, built from $F = \{f_1, \dots, f_s\}$, satisfies all of the requirements of Theorem 2.14, with*

$$D = 2^{12} \cdot 13 \cdot 20 \cdot (n_\varepsilon^2 + 1) \cdot (n_\delta^2 + 1) \approx 4.26 \cdot 10^{42}.$$

- Proof. Part I.* 1. Define $X \subset V(\mathcal{L})$ as the set of vertices of $\mathcal{L}$ that simulate vertices $v_1, \dots, v_m$ of $\mathcal{E}$. Define translation $T(\vec{xy}) = \vec{xy} + (a_1, b_1, \dots, a_m, b_m)$ where for each $1 \leq k \leq m$, $(a_k, b_k) - (Q/5 + 2r, Q/5 + 2r)$ are the coordinates of the lower left corner of the cell containing Start Gadget $\mathcal{L}_{\text{start}}(k)$; in other words, (a_k, b_k) is the position of vertex v_k that would correspond to $(\alpha_k, \beta_k) = (0, 0)$. We showed in Section 7.6 that $\mathcal{E}$ perfectly draws the region $T(Z(F) \cap U)$, where $U = (2r \cdot \text{Rect}([-d/\delta, d/\delta]))^m$. Because $\mathcal{L}$ liftably and rigidly simulates $\mathcal{E}$, $\mathcal{L}$ indeed liftably and rigidly draws the same trace.
2. Following the construction above, there are m Start Gadgets, $O(|\text{Coeffs}(2m, d)|) = O(\text{poly}(m^d, d^d))$ Angle Addition, Vector Creation, Vector Rotation, and Vector Addition Gadgets, and s End Gadgets. Stacked one on top of the other to the side of the wire column (other than the Start Gadgets above the wire column), these gadgets require a height of $O(\text{poly}(m^d, d^d, s))$. The wire column likewise has no more than $O(\text{poly}(m^d, d^d, s))$ wires, so the entire construction fits in an $O(\text{poly}(m^d, d^d, s)) \times O(\text{poly}(m^d, d^d, s))$ grid of cells. Each cell contains $O(1)$ vertices and edges, proving the claim.
3. We already concluded that $\mathcal{E}'$ is globally noncrossing with global minimum feature size at least 2^{-8} . In the insertion of Angle Restrictor Gadget neighborhoods to transform $\mathcal{E}'$ to $\mathcal{L}$, scale each gadget by a factor of 2^{-12} , so that it fits within a 2^{-10} neighborhood of its vertex and (on its own) has a global minimum feature size at least $2^{-12} \cdot 1/20$. $\mathcal{E}'$ ensures that each Angle Restrictor Gadget neighborhood stays well separated from

other edges or Angle Restrictor Gadgets, so $2^{-12} \cdot 1/20 > 1/D$ is a lower bound on the global minimum feature size of $\mathcal{L}$, as desired.

4. Each rigidified tree T_{uv} in $\mathcal{L}$ is constructed by attaching a portion of an Angle Restrictor Gadget neighborhood (as illustrated in Figure 21) to one or both endpoints of edge uv . These rigidified trees are indeed orthogonal and edge disjoint, possess at least three noncollinear vertices, and only connect to other trees or edges via leaves.
5. The combinatorial embedding σ of $\mathcal{L}$ follows from the combinatorial embedding of extended linkage $\mathcal{E}'$ (which is unique by Lemma 7.22), augmented with the unique combinatorial embedding of Angle Restrictor Gadgets (Lemma 7.23, Part 2). Any vertex u of $\mathcal{L}$ that is *not* internal to a rigidified tree must either belong to an Angle Restrictor Gadget neighborhood, or come from a vertex of $\mathcal{E}'$ of degree 1. In the latter case, u still has degree 1 in $\mathcal{L}$, so we need not consider it. If instead u belongs to the Angle Restrictor Gadget neighborhood of some vertex $v \in \mathcal{E}'$, it may be seen in Figure 21 that u 's angles remain within $(60^\circ, 240^\circ)$: indeed, by Lemma 7.23 (Angle Restrictor Gadget), the only place this condition might be violated is at the center of the neighborhood, $u = v$, but Figure 21 shows that the angles at v 's corners remain in $\pi/2 \pm \gamma$ or $\pi \pm \gamma$.
6. This is true by construction.

Part II. 7. When the coefficients of $f_1, \dots, f_s$ are integers, all coefficients $d_{j,u,I}$ are also integers bounded by $O(\text{poly}(m^d, d^d, s, M))$ (by Lemma 7.10). Parameters r, Q, R are also integers bounded by $O(\text{poly}(m^d, d^d, s, M))$, so it may be seen that all edge lengths of $\mathcal{E}$ are *integers* no greater than Q (including the Vector Creation Gadgets with integers $w = d_{j,u,I}$ and End Gadgets with integers $w = -f_j(0)$). The edges of $\mathcal{E}'$ have lengths in $\frac{1}{32}\mathbb{Z}$, because the scaled Sliceform Gadgets $\mathcal{L}_{\angle\text{slice}}$ that were inserted into $\mathcal{E}$ to form $\mathcal{E}'$ have these edge lengths. Finally, the scaled Angle Restrictor Gadget neighborhoods (with each Angle Restrictor Gadget having shape $\mathcal{L}_{\angle\text{restrict}}(n_\epsilon)$, $\mathcal{L}_{\angle\text{restrict}}(n_\delta)$, or $\mathcal{L}_{\angle\text{restrict}}(\infty)$) may be checked to have edge lengths in $\frac{1}{D}\mathbb{Z}$, so the same can be said for $\mathcal{L}$.

The only edges of $\mathcal{L}$ not contained in a rigidified tree correspond to those edges in Figure 19 (Angle Restrictor Gadget) that are drawn with thin lines. In $\mathcal{L}$, these edges all have lengths 2^{-12} , $2^{-12} \cdot 3/20$, and $2^{-12} \cdot 3/10$, which are less than D .

8. Angular forms $g_j(\overrightarrow{\alpha\beta})$ may be computed from polynomials $f_j(\overrightarrow{xy})$ in deterministic time $O(\text{poly}(m^d, d^d, s, M))$ by Lemma 7.10, and the magnitudes of the coefficients never exceed $O(\text{poly}(m^d, d^d, s, M))$. Our transformations from $f_1, \dots, f_m$ to $\mathcal{L}(F)$ are explicit and deterministic.

Part III. 9. When the given polynomials f_j satisfy $f_j(\vec{0}) = 0$, our construction of $\mathcal{E}(F)$ uses End Gadgets only with input $w = 0$, which come with initial configurations with coordinates that are integer multiples of $Q/40$ and are therefore integers, as in Lemma 7.20. The Vector Creation Gadgets, with integer input $w = d_{j,u,I}$, likewise have integer coordinates initially. All other gadgets start at integer coordinates unconditionally. Together, these specify an initial configuration of $\mathcal{E}(F)$ corresponding to $\overrightarrow{\alpha\beta} = \vec{0}$, i.e., $\overrightarrow{xy} = \vec{0}$, which indeed has integer coordinates.

The initial configurations of the Sliceform Gadgets and Angle Restrictor Gadgets were also illustrated in Figures 18 and 19 and may be checked to have rational coordinates in $\frac{1}{D}\mathbb{Z}$. These induce the desired initial configuration C_0 of $\mathcal{L}(F)$.

10. Linkage $\mathcal{E}'$ perfectly draws $T(Z(F) \cap U)$, meaning there is only one configuration of $\mathcal{E}'$ mapping to $T(\vec{0})$. In this unique configuration, each Angle Restrictor Gadget must be in its initial state (corresponding to $\theta = \pi/2$ in Lemma 7.23), and there is only one such configuration for each gadget. So C_0 is indeed unique.
11. All edges of $\mathcal{E}'$ are initially axis aligned, so the rigidified trees built from these edges are likewise axis aligned.
12. As described above, our construction of C_0 is explicit and deterministic.

All properties have been verified, so this concludes the proof. $\square$

7.9 Modifications for Strong Matchstick Universality

We may subtly modify the above proof of Theorem 2.14 (Main Theorem) to prove that the proper subsets of $\mathbb{R}^2$ drawn by matchstick linkages are *exactly* the bounded semialgebraic sets. We will use one extra cell gadget when constructing extended linkage $\mathcal{E}(F)$, the Crossing End Gadget (Figure 22, Lemma 7.26), which creates a crossing precisely when $g(\vec{xy}) = 0$ for a given polynomial g . When linkage $\mathcal{E}(F)$ is simulated by a matchstick linkage $\mathcal{M}(F)$ as described in Section 6, all of $\mathcal{E}(F)$'s noncrossing configurations transfer to $\mathcal{M}(F)$, i.e., thickening does not introduce unintended crossings. This allows us to draw semialgebraic sets of the form

$$\{\vec{x} \in \mathbb{R}^k \in \mathbb{R}^2 \mid f_1(\vec{x}) = \dots = f_s(\vec{x}) = 0, g_1(\vec{x}) \neq 0, \dots, g_r(\vec{x}) \neq 0\},$$

as well as coordinate projections thereof. This is sufficient to draw any bounded semialgebraic set in the plane. Details follow.


Theorem 7.25 (Universality of Matchstick Linkages). *The proper subsets $R \subsetneq \mathbb{R}^2$ that are drawable by a matchstick linkage are exactly the bounded semialgebraic sets.*

Proof. Let $\mathcal{M} = (\mathcal{L}, \text{NX}_{\mathcal{L}})$ be a matchstick linkage that is connected and has at least one pin. For the underlying, unconstrained linkage $\mathcal{L}$, the set $\text{Conf}(\mathcal{L})$ is described by closed polynomial conditions and is therefore compact and algebraic. Noncrossing can be described by the nonvanishing of polynomials, so $\text{Conf}(\mathcal{M}) = \text{NXConf}(\mathcal{L}) \subset \text{Conf}(\mathcal{L})$ is bounded and semialgebraic (not necessarily closed). Then, for any vertex v , the trace $\pi_v(\text{Conf}(\mathcal{M}))$ is also bounded and semialgebraic.

Conversely, to show that every bounded, semialgebraic set can be drawn with a matchstick linkage, it suffices (by Lemma 3.2) to show that every bounded, basic semialgebraic set of the form

$$U = \{\vec{x} \in \mathbb{R}^{2m} \mid f_1(\vec{x}) = \dots = f_s(\vec{x}) = 0, g_1(\vec{x}) \neq 0, \dots, g_r(\vec{x}) \neq 0\}$$

can be drawn by a matchstick linkage, up to translation. We may assume U does not contain the origin (by translating if necessary), and by replacing each $g_j(\vec{x})$ by $g_j(\vec{x}) \cdot |\vec{x}|^2$ (which



does not modify U), we may further assume that each g_j satisfies $g_j(\vec{0}) = 0$. We first prove the special case where $U \subset [-1, 1]^{2m}$; the general case is discussed at the end.

$$\frac{1}{D} \cdot U = \{\overrightarrow{xy} \mid F_j(\overrightarrow{xy}) = 0 \text{ for } 1 \leq j \leq s \text{ and } G_j(\overrightarrow{xy}) \neq 0 \text{ for } 1 \leq j \leq r\}. \quad (9)$$

Now we will use the “coefficients as variables” trick from the proof of Theorem 5.9. Namely, for each polynomial $P \in \{F_1, \dots, F_s, G_1, \dots, G_r\}$, and for each monomial $c_P j \vec{x}^j \vec{y}^j$ in P , we create new variables $a_{P,J}$ and $b_{P,J}$, gather all of these new variables into a vector $\vec{ab}$, and define the new polynomials

$$P'(\overrightarrow{xy}, \overrightarrow{ab}) := \sum_{J \text{ such that } c_{P,J} \neq 0} a_{P,J} \overrightarrow{xy}^J.$$

Polynomials $F'_1, \dots, F'_s, G'_1, \dots, G'_r$ now have integer coefficients: in fact, all coefficients equal 1. It remains to implement the equations $F'_j(\vec{xy}, \vec{ab}) = 0$, $G'_j(\vec{xy}, \vec{ab}) \neq 0$, $(a_{F_j, J}, b_{F_j, J}) = (c_{F_j, J}, 0)$, and $(a_{G_j, J}, b_{G_j, J}) = (c_{G_j, J}, 0)$, for all j and J , which exactly recover the solutions $\vec{xy}$ from (9).

To implement each non-equality $G'_j(\overrightarrow{xy}, \overrightarrow{ab}) \neq 0$, we create a custom extended linkage cell gadget, the *Crossing End Gadget* $\mathcal{L}_{\text{X-end}}$, that is *not* globally noncrossing, but instead takes as input a real number in some range (namely, the horizontal offset of vertex g) and induces a crossing precisely when that input is zero:

Lemma 7.26 (Crossing End Gadget). *The Crossing End Gadget, $\mathcal{L}_{\text{X-end}}$, is a copy of $\mathcal{L}_{\text{angular}}$ with one additional edge b_2h of length $Q/2$. (It is drawn short to emphasize that*

g and h are distinct and do not share an edge, but in the initial configuration depicted, g and h actually overlap.) Let $H \subset \text{Conf}(\mathcal{L}_{\text{X-end}})$ consist of those configurations in which g has y coordinate equal to $Q/2$. Then the only configuration in H with a crossing has $g = (Q/2, Q/2)$. All configurations in H have minimum feature size at least $1/2$ if the distance between edges b_2h and kg is ignored.

Proof. This follows from Lemma 7.19 (Angular Gadget). $\square$

We may now proceed as in the proof of the Main Theorem (Theorem 2.14) using polynomials $F'_1, \dots, F'_s$ and $G'_1, \dots, G'_r$, and concluding each G'_j with a Crossing End Gadget instead of an End Gadget. The rest of the linkage is already designed to ensure that the angles α, β fed into each Crossing End Gadget satisfy $R \cdot \text{Rect}(\alpha, \beta) = (G'_j(\vec{xy}) - G'_j(\vec{0}), 0)$ (see Step 5 of Section 7.6), so only configurations in the set H from Lemma 7.26 are possible. Then, as in Theorem 5.9, we pin all the vertices $v_{P,J}$ in the plane to force variables $(a_{P,J}, b_{P,J})$ to take the values $(c_{P,J}, 0)$. For the resulting extended linkage $\mathcal{E}$ with drawing vertices X , we conclude that the map π_X is a homeomorphism between $\text{NXConf}(\mathcal{E})$ and a translation of $\frac{1}{D} \cdot U$, and furthermore, the only crossings in any configurations of $\mathcal{E}$ come from the Crossing End Gadgets.

As in Section 7.8, eliminate sliceforms in $\mathcal{E}$ to form $\mathcal{E}'$ and process this further (as before) into partially rigidified linkage $\mathcal{L}$. Finally, modify $\mathcal{L}$ into a matchstick linkage $\mathcal{M}$ as in Construction 6.4.

We claim that $\mathcal{M}$ simulates $D \cdot \mathcal{E}$. This means that each noncrossing configuration C of $\mathcal{E}$ gives rise to a noncrossing configuration of $\mathcal{M}$, i.e., no *unintended* crossings arise in $\mathcal{M}$'s underlying linkage. Supposing the contrary, if C is a crossing configuration, the crossing must happen in the vicinity of a Crossing End Gadget by Lemma 7.26: the edge polyiamonds ultimately built from edges kg and b_2h of $\mathcal{L}_{\text{X-end}}$ must intersect. But the former edge polyiamond lies entirely in the lower half of the gadget, except for its endpoint on the line $y = Q/2$: indeed, in the Crossing End Gadget, vertex g remains on the line $y = Q/2$ and the direction of edge kg changes by at most $\pm\delta$, far less than the 30° required for any other part of the edge polyiamond of kg to reach the line $y = Q/2$. For similar reasons, the edge polyiamond for b_2h remains above $y = Q/2$ except for its stationary endpoint at $h = (Q/2, Q/2)$, so they could intersect only at h . But then C is a crossing configuration of $\mathcal{E}$, contrary to assumption. So $\mathcal{M}$ indeed simulates $D \cdot \mathcal{E}$ and thus draws the set $D \cdot \frac{1}{D} \cdot U = U$.

Finally, we consider the general case where U does not necessarily reside in $[-1, 1]^{2m}$. First choose an integer n large enough so that $\frac{1}{n} \cdot U$ lies in $[-1, 1]^{2m}$, and construct a matchstick linkage $\mathcal{L}$ that draws $\frac{1}{n} \cdot U$ as detailed above. Build a new matchstick linkage $\mathcal{L}'$ that copies $\mathcal{L}$ except that each edge polyiamond is n -times longer. This linkage $\mathcal{L}'$ will indeed draw U , as required. $\square$

8 Open Problems

Table 1 settles most problems in this area, but a few interesting open problems remain. The one combination in the table that remains unsolved is the complexity of deciding global

rigidity in graphs with unit edge lengths, allowing crossings. In particular, are there *any* such graphs larger than a triangle? (If not, the decision problem has an easy algorithm!)

We could also consider additional graph types (rows) in Table 1. For example, we considered unit edge lengths and edge lengths in $\{1, 2\}$, both when allowing crossings and when forbidding crossings, but we did not consider globally noncrossing graphs with edge lengths restricted to $\{1\}$ or $\{1, 2\}$. Do these linkages remain universal for compact semialgebraic sets? Is realizing them $\exists\mathbb{R}$ -complete? Is testing their rigidity and global rigidity $\forall\mathbb{R}$ -complete? Our results for globally noncrossing graphs use only integer edge lengths bounded by a universal constant D , so we are “only” a constant factor away from a bound of 1 or 2.

On the practical side, $D \approx 10^{42}$ is rather large. We have been loose with our constants to keep the analysis as simple as possible, but it would be interesting to tune the constants and see how practical the whole construction can be made. After all, a primary motivation for avoiding crossings was to make it possible to physically construct universal linkages.

We also introduced the class of globally noncrossing graphs. Is it $\forall\mathbb{R}$ -complete to determine whether a graph with edge-length constraints is globally noncrossing, that is, whether all its realizations are noncrossing? Our $\exists\mathbb{R}$ -completeness proof for realizing globally noncrossing graphs shows that distinguishing between unrealizable graphs and realizable globally noncrossing graphs is $\exists\mathbb{R}$ -complete, but both cases are technically “globally noncrossing”.

What if we relax our edge-length constraints to allow approximate solutions, modelling a small amount of pliancy in the bars or tolerance at the hinges? Approximate realizations where each edge can be stretched by at most some α factor are considered extensively in the field of metric embedding, with many interesting approximation algorithms. Saxe [21] proved that it is strongly NP-complete to distinguish between graphs realizable with stretch $\alpha = 1 + \frac{1}{18}$ from graphs realizable with stretch $\alpha = 1 + \frac{1}{9}$, when embedding into one dimension. What about embedding into two dimensions? What about when restricted to globally noncrossing graphs, matchstick graphs, graphs with unit edge lengths, or graphs with edge lengths in $\{1, 2\}$? (Saxe’s proof uses edge lengths in $\{1, 2, 3, 4\}$.)

Acknowledgments

We are grateful to the anonymous referees whose thorough suggestions greatly improved the presentation of this paper.

References

- [1] Timothy Good Abbott. Generalizations of Kempe’s Universality Theorem. Master’s thesis, Massachusetts Institute of Technology, 2008. Joint work with Reid W. Barton and Erik D. Demaine. URL: <http://hdl.handle.net/1721.1/44375>.

- [2] Zachary Abel, Erik D. Demaine, Martin L. Demaine, Sarah Eisenstat, Anna Lubiw, André Schulz, Diane Souvaine, Giovanni Viglietta, and Andrew Winslow. Algorithms for designing pop-up cards. In *Proceedings of the 30th International Symposium on Theoretical Aspects of Computer Science*, pages 269–280, Kiel, Germany, 2013. doi:[10.4230/LIPIcs.STACS.2013.269](https://doi.org/10.4230/LIPIcs.STACS.2013.269).
- [3] Kenneth Appel and Wolfgang Haken. Every planar map is four colorable, Part I: Discharging. *Illinois Journal of Mathematics*, 21:429–490, 1977. doi:[10.1215/ijm/1256049011](https://doi.org/10.1215/ijm/1256049011).
- [4] Jacek Bochnak, Michel Coste, and Marie-Françoise Roy. *Real Algebraic Geometry*. Springer-Verlag, 1998. doi:[10.1007/978-3-662-03718-8](https://doi.org/10.1007/978-3-662-03718-8).
- [5] Sergio Cabello, Erik D. Demaine, and Günter Rote. Planar embeddings of graphs with specified edge lengths. *Journal of Graph Algorithms and Applications*, 11(1):259–276, 2007. doi:[10.1007/978-3-540-24595-7_26](https://doi.org/10.1007/978-3-540-24595-7_26).
- [6] John Canny. Some algebraic and geometric computations in PSPACE. In *Proceedings of the 20th Annual ACM Symposium on Theory of Computing*, pages 460–469, Chicago, Illinois, 1988. doi:[10.1145/62212.62257](https://doi.org/10.1145/62212.62257).
- [7] Erik D. Demaine and Joseph O’Rourke. *Geometric Folding Algorithms: Linkages, Origami, Polyhedra*. Cambridge University Press, 2007. doi:[10.1017/CBO9780511735172](https://doi.org/10.1017/CBO9780511735172).
- [8] Peter Eades and Nicholas C. Wormald. Fixed edge-length graph drawing is NP-hard. *Discrete Applied Mathematics*, 28(2):111–134, 1990. doi:[10.1016/0166-218X\(90\)90110-X](https://doi.org/10.1016/0166-218X(90)90110-X).
- [9] Denis Jordan and Marcel Steiner. Configuration spaces of mechanical linkages. *Discrete & Computational Geometry*, 22:297–315, 1999. doi:[10.1007/PL00009462](https://doi.org/10.1007/PL00009462).
- [10] Michael Kapovich and John J. Millson. Universality theorems for configuration spaces of planar linkages. *Topology*, 41(6):1051–1107, 2002. doi:[10.1016/S0040-9383\(01\)00034-9](https://doi.org/10.1016/S0040-9383(01)00034-9).
- [11] Toshikazu Kawasaki. On the relation between mountain-creases and valley-creases of a flat origami. In *Proceedings of the 1st International Meeting of Origami Science and Technology*, pages 229–237, 1989.
- [12] Alfred Bray Kempe. On a general method of describing plane curves of the n^{th} degree by linkwork. *Proceedings of the London Mathematical Society*, 7:213–216, 1876. doi:[10.1112/plms/s1-7.1.213](https://doi.org/10.1112/plms/s1-7.1.213).
- [13] Alfred Bray Kempe. *How to Draw a Straight Line: A Lecture on Linkages*. Macmillan and co., London, 1877. URL: <https://www.gutenberg.org/ebooks/25155>.
- [14] Alfred Bray Kempe. On the geographical problem of the four colours. *American Journal of Mathematics*, 2:183–200, 1879. doi:[10.2307/2369235](https://doi.org/10.2307/2369235).

- [15] Henry C. King. Planar linkages and algebraic sets. *Turkish Journal of Mathematics*, 23(1):33–56, 1999. URL: <https://journals.tubitak.gov.tr/math/vol23/iss1/3/>.
- [16] Pascal Koiran. The complexity of local dimensions for constructible sets. *Journal of Complexity*, 16(1):311–323, 2000. doi:10.1006/jcom.1999.0536.
- [17] William S. Massey. Sufficient conditions for a local homeomorphism to be injective. *Topology and its Applications*, 47(2):133–148, 1992. doi:10.1016/0166-8641(92)90068-B.
- [18] Nikolai E. Mnëv. The universality theorems on the classification problem of configuration varieties and convex polytopes varieties. In Oleg Viro and Anatoly Vershik, editors, *Topology and Geometry — Rohlin Seminar*, volume 1346 of *Lecture Notes in Mathematics*, pages 527–543. Springer, 1988. doi:10.1007/BFb0082792.
- [19] Neil Robertson, Daniel Sanders, and Paul Seymour. The four-colour theorem. *Journal of Combinatorial Theory, Series B*, 70:2–44, 1997. doi:10.1006/jctb.1997.1750.
- [20] Walter Rudin. *Principles of Mathematical Analysis*. McGraw-Hill, Inc., 3rd edition, 1976.
- [21] James B. Saxe. Embeddability of weighted graphs in k-space is strongly NP-hard. In *Proceedings of the 17th Annual Allerton Conference on Communication, Control, and Computing*, Monticello, IL, 1979. URL: <https://www.cs.duke.edu/brd/Teaching/Bio/asmb/current/Readings3/saxe-embeddability.pdf>.
- [22] Marcus Schaefer. Realizability of graphs and linkages. In János Pach, editor, *Thirty Essays on Geometric Graph Theory*, chapter 23. Springer, 2013. doi:10.1007/978-1-4614-0110-0_24.
- [23] Peter W. Shor. Stretchability of pseudoline arrangements is NP-hard. In Peter Gritzmann and Bernd Sturmfels, editors, *Applied Geometry and Discrete Mathematics: The Victor Klee Festschrift*, volume 4 of *DIMACS Series in Discrete Mathematics and Theoretical Computer Science*, pages 531–554. American Mathematical Society, 1991. URL: https://math.mit.edu/~shor/papers/Stretchability_NP-hard.pdf.