

Parallel-Redrawing Mechanisms, Pseudo-Triangulations and Kinetic Planar Graphs

Ileana Streinu

Computer Science Department, Smith College, Northampton, MA 01063
streinu@cs.smith.edu

Abstract. We study *parallel redrawing graphs*: graphs embedded on moving point sets in such a way that edges maintain their slopes all throughout the motion.

The configuration space of such a graph is of an oriented-projective nature, and its combinatorial structure relates to rigidity theoretic parameters of the graph. For an appropriate parametrization the points move with constant speeds on linear trajectories. A special type of kinetic structure emerges, whose events can be analyzed combinatorially. They correspond to collisions of subsets of points, and are in one-to-one correspondence with contractions of the underlying graph on rigid components. We show how to process them algorithmically via a *parallel redrawing sweep*.

Of particular interest are those *planar* graphs which maintain non-crossing edges throughout the motion. Our main result is that they are (essentially) pseudo-triangulation mechanisms: pointed pseudo-triangulations with a convex hull edge removed. These kinetic graph structures have potential applications in morphing of more complex shapes than just simple polygons.

1 Introduction

Consider a straight-line drawing of a graph $G = (V, E)$ in the plane. A *parallel redrawing* of G is another drawing so that for every edge $ij \in E$, the corresponding line segments in the two drawings are parallel (Fig. 1). A parallel redrawing is *trivial* if it is similar to the original drawing (via a rescaling or translation of). See Fig. 1(A-B). Classical results in Rigidity Theory, see [16], establish combinatorial criteria for a graph to admit non-trivial parallel redrawings, generically. In particular, such a graph has a certain number of *degrees of freedom* (for the example in Fig. 1(C-D) this number is one).

A *kinetic point set* is (for the purpose of this paper) a set of points in the plane moving with constant velocities, and a *kinetic graph* is a graph drawn on a kinetic point set. An illustration is found in the rightmost example of Fig. 2. As the points move, the shape of the embedding changes: edges may shrink to zero-length, points may collide and edges may cross for certain time intervals, and be non-crossing for others. A natural graph drawing problem is whether these phenomena may be




Fig. 1. A direction network (A) with only similar (trivial) parallel redrawings (B), and another one (C) with a non-similar parallel redrawing (D)




Fig. 2. The kinetic point set underlying a Laman mechanism: a parallel redrawing, the point velocities and the linear trajectories of the vertices

predicted from the law of motion of the points (given by their initial positions and velocities). In this paper, we give a complete answer to this problem under the additional constraint that throughout the motion, the graphs are parallel redrawings of each other. To keep the presentation short, we’ll discuss here only the case when the graphs have exactly *one degree of freedom*. Our **main algorithmic result** is an efficient technique for predicting the relevant events via a new process called the *parallel redrawing sweep*, of which Figure 3 is a preview.

Of particular interest is the *planar* case. Given a plane embedding of a planar graph, start moving the points (kinetically): *are there graphs which remain non-crossing throughout the whole duration of the motion?*

The answer to this question is intrinsically related to a special class of planar graphs, the *pointed pseudo-triangulation mechanisms*. A pseudo-triangulation is a plane graph with a convex outer face and with all inner faces embedded as pseudo-triangles: simple polygons with exactly three inner convex vertices (called the *corners* of the pseudo-triangle). In a *pointed pseudo-triangulation*, each vertex is incident to an angle larger than π . Removing a convex hull edge from such a graph produces a *pointed pseudo-triangulation mechanism* (ppt-mechanism), see Fig. 3. Our **main theorem** is that these are (essentially) the only one-degree-of-freedom graphs which *maintain planarity throughout a parallel redrawing motion*. Figure 3 gives a preview.

Historical background. Points moving with constant velocities appear in a popular morphing technique for planar polygonal shapes. It is known that the sim-

plicity of the shape may be violated during the morph. While there are morphing techniques which achieve simplicity, e.g. via compatible planar triangulations as in [5], or by maintaining the slopes of the polygon edges, as in [7], nobody seems to have analyzed theoretically what this simple paradigm for motion has to say about edge crossings.




Fig. 3. A parallel redrawing pointed pseudo-triangulation mechanism: snapshots from its motion. The second and third snapshot have the same combinatorial structure, but all the others differ in the convex/reflex angles per face. Note the flips in rigid bodies between consecutive snapshots with different combinatorial structure.

Pseudo-triangulations are relatively new objects, applied in Computational Geometry for problems such as visibility [10] and kinetic data structures [1]. The *pointed* or *minimal* pseudo-triangulations [14, 15], have interesting rigidity-theoretical properties and applications in designing efficient motion planning algorithms for planar robot arms. Recently many papers have investigated their rich combinatorial, rigidity theoretic and polyhedral properties, e.g. [8, 11, 2, 3] and applications [12].

The main result of this paper, as stated above, is in fact a new rigidity-theoretic property of pointed pseudo-triangulation *mechanisms*. Whereas in the fixed-edge-length model of rigidity, these graphs are 1-dof *expansive* mechanisms, the *new result* of this paper is that in the *fixed-edge-direction* model, they capture the essence of kinetic parallel-redrawing *planar graphs*. The property is surprisingly simple, has an elementary proof and is entirely new: I am not aware of anybody even asking such a question before. I am presenting it to the Graph Drawing community with the expectation that it may find applications beyond those that originally motivated my investigation.

How to read this extended abstract. To formally state the results (theorems and algorithms), I have no choice but to plunge into a fair amount of definitions: it will take a few pages, since there is no standard reference where I can send the reader to gather them all (but please skip directly to page 426 if this is familiar material). To help the reader less fluent with all these concepts get faster to the ideas, and especially since many of them have a kinetic nature which can hardly be conveyed only with static printed images, I have assembled a web site with animations and interactive applets illustrating them: <http://cs.smith.edu/~streinu/Research/KineticPT/>

2 Preliminaries

A (planar) *kinetic point set* $p(t) = \{p_1(t), \dots, p_n(t)\}$ is a finite collection of time-dependent points $p_i(t) \in R^2$ in the Euclidean plane, each one moving on a *linear trajectory* at constant *speed*. The continuous *time* parameter t runs over the real line R , but it will become apparent during our proofs and analysis that it is very useful to think of it sometimes as running over the *projective line*: R extended with a *point at infinity*.

A *graph drawing* or *embedding* $G(p)$ is a mapping of the vertices V of a graph $G = (V, E)$ to a set $p = \{p_1, \dots, p_n\}$ of points in the plane, $i \mapsto p_i$. A *planar graph embedding* (or a *plane graph*) has no crossing edges. A *topological plane graph* is the (planar) graph with the additional information regarding its faces and their incidences, given for instance by the *rotations* of edges around vertices (but not necessarily by a concrete embedding).

A *kinetic graph* $G(p(t))_{t \in I}$ is a graph G embedded on a kinetic point set $p(t)$. Of particular interest are the kinetic graphs which remain *planar* throughout the entire motion $t \in R$. To formally state our results on planar kinetic graphs, we give now the necessary definitions regarding Parallel Drawings, Rigidity Theory and Pseudo-triangulations.

A *direction network* (G, D) is a graph together with a set of directions (slopes) d_{ij} associated to its edges $ij \in E$. A drawing¹ or *realization* $G(p)$ of (G, D) is an embedding $G(p)$ of G on a set of points p such that for each edge $ij \in E$, the direction of the line through p_i and p_j is d_{ij} .

If we denote the coordinates of the unknown points as $p_i = \{x_i, y_i\}$ and the known directions as $d_{ij} = [a_{ij} : b_{ij}]$ (as projective points, expressed with homogeneous coordinates), we obtain each realization as a solution of the homogeneous linear system:

$$\langle p_i - p_j, d_{ij}^\perp \rangle = 0, \quad \forall ij \in E \tag{1}$$

where $d_{ij}^\perp = [b_{ij} : -a_{ij}]$ is the vector orthogonal to d_{ij} . Since every translation of a realization yields another realization, we can factor them out by pinning down a vertex at the origin, e.g. as $x_1 = y_1 = 0$. This shows that the set of all the realizations (modulo translations) is a linear space (of some dimension k), and contains the *trivial realization*, where all the points coincide (with the origin).

A *parallel redrawing* of a graph embedding $G(p)$ is another embedding $G(p')$ of the same graph, such that the corresponding edges have parallel directions in the two drawings. A parallel redrawing is *trivial* if it is a translation $G(p + p_0)$ or a rescaling $G(\alpha p)$ of $G(p)$. In this case the two figures are *similar*. See Fig. 1. The *configuration space* of a direction network is the space of all possible realizations, modulo the trivial parallel redrawings.

A direction network is *consistent* if it has a realization with not all the points coinciding (i.e. its configuration space is not trivial). It is *generic* if small per-

¹ It is worth alerting the reader at this point that although our graph drawings are taking place in the Euclidean plane, the resulting space of *all possible drawings* will turn out to be a *real projective space*.




Fig. 4. (A) An inconsistent direction network: only the trivial realization matches for the edge ij the direction of the thick segment. (B) Changing the direction of the inconsistent edge to a well defined value turns it into a consistent network.

turbations of the directions do not change the dimension of the configuration space. It is (*generically*) *tight* if it is generic and the configuration space contains exactly one non-trivial embedding (modulo parallel redrawings). Otherwise, it is called a (*generically*) *loose* graph. A graph is *tight* if for some choice of directions, it becomes a generic tight network, and it is *minimally tight* if it is tight, but the removal of any edge makes it loose. See Fig. 4 for some examples.

A point set is *trivial* if all the points coincide, $p_i = p_j, \forall i, j \in [n]$, otherwise it is non-trivial. A point set is *degenerate* if some of its points coincide. Otherwise, if *all the points* are distinct, it is called *non-degenerate* or *sharp*.

What is the point of all these definitions? The realizations of direction networks are solutions to linear systems, which may or may not have solutions, or have too many (if they are *under-determined*). The interesting thing is that the nature of the linear solution space is controlled by the combinatorial structure of the graph, generically. For almost all possible choices of slopes, we can read off whether a network is consistent, tight, etc. using a simple counting criterion (this is well known from rigidity theory). And where I am trying to get to, eventually (and this is one of the results of the paper), is that even questions such as whether there are realizations where some points *coincide*, or whether *all* the realizations of a direction network are non-crossing (planar) can be answered combinatorially. To get there, we need two more concepts: Laman graphs and pseudo-triangulations.




Fig. 5. A generically minimally tight graph (Laman graph) in a generic (A) and non-generic (B) embedding. A parallel redrawing is shown for (B), which is loose in this non-generic embedding. (C) A generically tight graph which is not minimally tight (Laman-plus-one), and (D) a generically loose graph (a 1dof Laman mechanism).

In a graph G , a subset of some k vertices *satisfies the Laman count* if it spans at most $2k - 3$ edges. G is a *Laman graph* if $m = 2n - 3$ and every subset satisfies the Laman count. See Figure 5(A).

Laman graphs are known in the Rigidity Theory literature as *generic minimally infinitesimally rigid graphs*, due to their rigidity theoretic properties (on which we won't elaborate here, since they will not be used in this paper). It is a well known fact in Rigidity Theory (see [16]) that *minimally tight graphs* coincide with the class of Laman graphs. Their configuration space is zero-dimensional, for generic choices of directions. Removing any subset of k edges creates a (*redrawing*) *mechanism*: a direction network with a k -dimensional configuration space.

In a Laman graph, a subset for which the Laman count is satisfied with equality is called a *block*. An *independent set of edges*² is a Laman graph with some edges (possibly none) removed. If at least one edge is missing, it is called a *Laman mechanism*. Of special interest is the *one-degree-of-freedom (1dof) Laman mechanism*, with only one missing edge. It has $2n - 4$ edges and every subset satisfies the Laman count. See Fig. 5 (D). A *maximal block* of a Laman mechanism is called a *rigid component* (shortly, an *r-component*).

A *pointed planar graph* is a non-crossing (planar) graph embedding where each vertex is incident to an angle larger than π . A *pointed pseudo-triangulation* is a pointed planar graph where each face is a pseudo-triangle. In [14] it was shown that pointed pseudo-triangulations are *maximal* pointed plane graph embeddings (maximal with respect to edge set inclusion), and that they are planar Laman graphs (and hence generically minimally rigid). When a convex hull edge is removed, the resulting *pseudo-triangulation mechanism* has several rigid components, which are themselves pointed pseudo-triangulations if they contain more than two vertices.

3 Overview of the Results

We show that *generic* parallel redrawings of 1-dof Laman mechanisms induce *kinetic point sets* and *kinetic graphs* with a rich inner structure. Although the definition of configuration space is static and seems to be purely algebraic, it inherently - and naturally - has *motion* in it! Our results include a complete characterization of all the collision situations in such kinetic point sets by relating them to the *rigid components* in the associated kinetic graph. In particular, we prove that, generically, at most $2n - 4$ collision events may occur and that this bound is attained. We also describe how to efficiently predict combinatorially which clusters of points will collide, and when.

For this purpose, we introduce an algorithmic tool, the *parallel redrawing sweep*, which follows the time parameter of the kinetic point set and predicts combinatorially all the collision events in its configuration space.

² This terminology relates to the matroidal point of view of rigidity in dimension two, see [6].

Finally, we give a characterization of the kinetic parallel redrawing planar 1-dof Laman mechanism graphs which distort linearly³ for *any* interval of time without the occurrence of crossings, except at collision events, which take place in clusters corresponding to rigid components. Surprisingly, they turn out to be exactly⁴ the **pointed pseudo-triangulation mechanisms**.

Let us emphasize that the results described in this paper hold *generically*. One can always find non-generic situations, resulting in more collisions and other special types of behavior not addressed in this paper. The set of non-generic direction networks is described as an algebraic condition (technically, the rank of the parallel redrawing matrix is not maximum). The complement of this algebraic set, consisting in all the generic cases, is an open set characterizing *most of* the situations one would encounter in a random parallel redrawing kinetic point set.

The following theorems relate the configuration spaces of arbitrary *generic* and pseudo-triangulation direction networks of Laman mechanisms to certain types of kinetic point sets. The inner structure of the configuration space, viewed as an oriented-projective (defined in the spirit of [13]), rather than as an affine or projective real space) is captured in Theorem 1. A (*collision*) *cluster* is a maximal set of colliding points.

Theorem 1. Parallel-Redrawing Laman Mechanisms. *The oriented-projective configuration space of a 1-dof Laman mechanism is a (topological) circle. It contains exactly r antipodal pairs of isolated points (configurations), each pair corresponding to a colliding cluster of points induced by one of the r rigid components of the mechanism. The regions obtained by removing the isolated points are one-dimensional segments, corresponding to classes of collision free configurations.*

Figure 6 illustrates this theorem (and also Theorems 3 and 4) projectively (rather than oriented-projectively, which would just wrap around the sequence of events once more and show them rotated around by π). The reader is advised to use the interactive applet from my web page to get a better sense of what it actually means for *moving points*.

The theorem can be extended to k -dof parallel redrawing Laman mechanisms. We focus now on the 1-dof case. An affine part of the configuration space, obtained by removing one *point at infinity*, can be described as follows. Pin down to the origin an arbitrary vertex (called the *grounded vertex*) of the direction network. This eliminates translations. Then pin down a whole incident edge. This eliminates *rescalings*. By varying the length of another edge incident to the grounded vertex (i.e. choosing one of its coordinates to act as *time-parameter*), we *sweep* an affine part of the configuration space. The point at infinity corresponds to the *scale edge* shrinking to zero-length. We show that the underlying

³ A technical detail, not addressed in this extended abstract, is that the kinetic point set *may* obey other absolute laws of motion (not necessarily linear). It is their *relative* motion that must be *linear*.

⁴ Modulo a technical detail. The precise concept is that of a collapsed pseudo-triangulation mechanism, as in [11].




Fig. 6. The (projective) configuration space of a 4-bar parallel redrawing Laman mechanism. The sweep events and sweep segments are illustrated by a representative embedding. The first image represents the *event at infinity*: the collapse of the *scale edge* 14.

point set of a *grounded 1-dof parallel redrawing Laman mechanism*, parametrized by time, behaves like a kinetic point set, and call it the *parallel redrawing kinetic point set*. Alternate linear parametrizations, or circular versions sweeping the entire oriented projective trajectory may also be considered.

Theorem 2. Kinetic Parallel Redrawing Point Sets. *The points of a grounded parallel redrawing Laman mechanism with 1-dof move with constant velocities. As the time parameter changes, all the configurations (modulo translations and scalings) are encountered, except the one corresponding to the scale edge being reduced to zero length.*

See Fig. 2. The next Theorem 3 refines the previous result for pseudo-triangulation mechanisms.

Theorem 3. Parallel Redrawing Pseudo-Triangulation Mechanisms. *All the realizations with fixed directions of a pseudo-triangulation mechanism maintain the pseudo-triangulation mechanism property.*

All these pseudo-triangulations have the same plane graph structure, but the combinatorial pseudo-triangulation structure (defined in [8]) varies.

From the algorithmic point of view, we investigated a continuous process, called the *parallel redrawing sweep*, which generates (sweeps) the configuration space of such a one-degree-of-freedom mechanism. This can be seen as an animation of the kinetic graph in time, with particular attention being paid to the *events*, when *the combinatorial structure of the graph embedding changes*. See Fig. 3 for an illustration.

Theorem 4. Parallel Redrawing Sweep. *As the time parameter sweeps the real line, the events encountered correspond each to the collapse of a rigid component to a single point (and thus capture a collision of an entire cluster in the kinetic point set).*

Each event is characterized combinatorially by a Laman graph, obtained as the contraction of the Laman mechanism on a rigid component.

The events and their sequence can be predicted algorithmically as follows:

1. **Prediction of all the events.** The events can be computed a priori in $O(n^2)$ time using an algorithm for computing rigid components in a Laman mechanism (see [9] and the references given there).
2. **Prediction of the next event.** Given a time t and the direction of the sweep, the time of the next collision event can be computed in linear time.
3. **Computing the combinatorial event.** Given a collision time t , the colliding cluster can be computed combinatorially in $O(n^2)$ time, or algebraically in the time needed to solve a linear system.

Next section defines the concepts needed to understand the concept of *shape similarity* used during the parallel redrawing sweep. The rest of the abstract is a high level sketch of the proofs.

4 Shape Similarity of Embedded Graphs

Let $G(p)$ be a graph embedding. We define and collect several types of combinatorial information from the embedding, depending on both p and G . See the example in Fig. 7.

Signed circular hyperlines. for each vertex index i , the i th hyperline ($1 \leq i \leq n$) is the signed sequence of indices $j \neq i$, in the circular order in which a line rotating through the point p_i encounters all the other points p_j . See [4].

Unsigned circular hyperlines. same as before, but we ignore the signs.

Signed linear hyperlines. same as the circular ones, but recorded starting from the vertical direction.

Unsigned linear hyperlines. the unsigned version of the previous ones.




Fig. 7. (a) The signed linear hyperline is $3 : 24\bar{1}5$. (b) A graph embedding, with the set of signed linear hyperlines $1 : 23$, $2 : \overline{31}4$, $3 : 24\bar{1}5$, $4 : \overline{53}2$, $5 : 4\bar{3}$.

Let G be a graph and let $G(p)$, $G(p')$ be two non-degenerate embeddings. We say that the two embeddings are *combinatorially equivalent* if they have the same set of signed circular hyperlines. For a fixed graph G , a *combinatorial class* of embeddings $[G(p)]$ contains all the embeddings which are combinatorially equivalent to $G(p)$, and is identified by their common set of signed circular hyperlines.

This concept suggests the use of the partial signed hyperlines as a combinatorial criterion for discerning *shape similarity* of two embeddings of the same

graph. This measure is invariant under *translations, rotations and scalings*. For pseudo-triangulations, roughly the same information is contained in the *combinatorial pseudo-triangulations* defined in [8].

5 The Parallel Redrawing Sweep

Let G be a Laman graph with some edges (possibly none) added or removed. Let $m = 2n - 3 - k$ be the number of edges, $k < 0$, $= 0$ or $k > 0$. Let (G, D) be a generic direction network on G .

Lemma 1. (The Dimension Lemma). *The projective configuration space of (G, D) is trivial for extra added edges ($k \leq 0$). Otherwise, it is a projective subspace of dimension equal to the number $0 < k \leq 2n - 3$ of removed edges, embedded in the projective $(2n - 3)$ -dimensional space P^{2n-3} .*

We can visualize, via a continuous linear motion, an affine part of the projective configuration space of a 1dof Laman direction network. This is the **Parallel redrawing sweep**.

A *parallel redrawing mechanism* $G(p(t))_{t \in R}$ is a continuously deforming family of embeddings of a 1dof Laman mechanism direction network (G, D) . Out of the many possible parametrizations, here's one. An edge $gs \in E$ is pinned down, to eliminate translations and rescalings. This eliminates the point at infinity of the projective configuration space. The affine part can now be swept through by the continuous motion of a *free vertex* f on a linear trajectory by varying one of its coordinates, the *time parameter*. All the vertices of the mechanism (except those that were pinned down) move with constant speed along linear trajectories. If we lift the arrangement to the third dimension as *time*, it becomes a line arrangement in space, which is swept by a plane orthogonal to the time dimension. Combinatorially, we prove that the sweep events (the vertices of the arrangement) correspond to a *reorientation* of the local hyperline sequences at certain vertices and are in one-to-one correspondence with the rigid components of the mechanism.

The following lemmas clarify the occurrence of degenerate embeddings in Laman networks and contribute to the proof of Theorem 1. The proofs are elementary and make use of basic combinatorial properties of Laman graphs.

Lemma 2. *Let (G, D) be a generic Laman direction network. Then all its non-trivial realizations are non-degenerate.*

Lemma 3. *Let (G, D) be a generic Laman direction network.*

1. *For every non-zero scaling factor $s \neq 0$ there exists a unique embedding $G(p)$. For opposite scale factors s and $-s$, the two embeddings are one the rotation by π of the other.*
2. *All the non-trivial realizations of a Laman direction network have the same signed circular hyperlines. The realizations fall into two classes of signed linear hyperlines: one for positive and one for negative scale factors, which are reversal equivalent.*

We introduce now the operation of *contracting a rigid component* and relate it the different signed linear hyperlines of realizations $G(p)$ and to degenerate embeddings of (G, D) .

Let G be an arbitrary graph, $e = ij \in G$ an edge. The *contraction* G/e of G on e is a new graph (possibly not simple) obtained by *collapsing* the endpoints i and j into a single vertex i' . In the contraction, every vertex incident to i or j becomes instead incident to i' and the edge ij is discarded. The contraction operation is extended naturally to subsets E' of edges, which are contracted one by one in an arbitrary order. The contracted graph is denoted by G/E' .

Lemma 4. *Let G be a Laman mechanism and $R \subset V$ a subset of vertices spanning a rigid component (R, E_R) . Then the contraction G/E_R of a 1dof Laman mechanism on a rigid component is a Laman graph. The contraction G/E_R of a connected k -dof Laman mechanism on a rigid component is a $(k - 1)$ -dof Laman graph.*

Let (G, D) be a direction network, $E' \in E$ a subset of edges, V' the set of endpoints of the edges in E' and G/E' the contraction of G on E' . As a result of the contraction on E' , all the vertices in V' have been collapsed into one. We denote this new vertex by $[V']$, when we want to remind the set of vertices that have been collapsed into one, or shortly by i' .

We define the *contracted direction network* $(G/E', D')$ obtained from (G, D) by keeping the directions of the edges not touched by the contraction operation, discarding the deleted edges and giving the direction $d_{i'j} = d_{ij}$ to a new edge $i'j$ resulting from an old edge ij , $i \in E'$, $j \notin E'$. We are interested in characterizing the configuration space of the contracted network in terms of the configuration space of the original network, for Laman graphs and mechanisms. We show first that the contraction can be realized algebraically by the addition of a single linear equation.

Lemma 5. *Let (G, D) be a direction network. Let $e = ij \in E$. If d_{ij} is not vertical, add the constraint $x_i = x_j$ to its direction network system, otherwise add the constraint $y_i = y_j$. This is equivalent to asking for an embedding where this edge has length zero, hence its endpoints are collapsed. Then, the resulting system has the same realization space as the direction network on the contraction G/e .*

The proof follows from elementary linear algebra. We apply now this fact to Laman graphs and mechanisms, under the same assumptions as in Lemma 5.

Lemma 6. *Adding the constraint $x_i = x_j$ to the direction network system of a Laman graph produces only trivial realizations.*

Lemma 7. *Adding the constraint $x_i = x_j$ to the direction network system of a Laman mechanism produces a unique realization (up to translations and scalings). This realization is non-trivial, and contains a unique trivial part corresponding to the collapsed vertices of the rigid component to which the edge ij belongs.*

Lemma 8. *The only degenerate embeddings of a 1dof Laman mechanism are those obtained by the collapsing of rigid components.*

The Parallel Redrawing Sweep. The inner structure of the configuration space of a generic 1dof Laman mechanism consists of:

Events. If G has r rigid components, the (oriented projective) configuration space contains exactly r pairs of *special* degenerate embeddings. Each pair corresponds to the two *antipodal* realizations of a contraction of (G, D) on one of the r rigid components.

Sweep segments. The removal of the points corresponding to the collapsed embeddings from the circular configuration space leaves r pairs of connected components (antipodal circle, or *sweep segments*, in the oriented projective space), each one corresponding to a pair of reversal-equivalent combinatorial classes of embeddings. By Lemma 8, these realizations are non-degenerate.

Theorem 5. *All the embeddings within a sweep segment have the same signed linear hyperlines. Embeddings in antipodal segments differ by a complete reversal of signs in their signed linear hyperlines.*

6 Parallel Redrawing Pseudo-Triangulation Mechanisms

A *parallel redrawing pseudo-triangulation mechanism* inherits all the properties proved in Section 5. Theorem 3 is a consequence of the following lemmas.

Lemma 9. *Let $G(p)$ be a generic embedding of a Laman mechanism and let $G(p')$ be a parallel redrawing of $G(p)$. Viewing now $G(p)$ as a mechanism with fixed edge lengths, let \mathcal{V} be the linear space of all the infinitesimal motions of $G(p)$. Let $v = \{v_1, \dots, v_n\} \in \mathcal{V}$ be a non-trivial infinitesimal motion and define $\sigma_{ij}(p, v) := \langle p_i - p_j, v_i - v_j \rangle, \forall i, j \in [n]$. Then $\sigma_{ij}(p, v) = \sigma_{ij}(p', v), \forall i, j, \forall v \in \mathcal{V}$.*

Under the same assumptions as in Lemma 9, the following two properties hold:

Corollary 6. *When viewed as fixed-edge lengths mechanisms, $G(p)$ and $G(p')$ have the same linear space of infinitesimal motions, in particular they have the same pattern of infinitesimal expansion.*

Indeed, the pattern of infinitesimal expansion is given by the signs of $\sigma_{ij}(p, v)$.

Corollary 7. *If $G(p)$ is a pseudo-triangulation, then $G(p')$ is a pseudo-triangulation. We pin down an edge and parameterize by the position of a third vertex k . Then all the resulting parallel redrawings are pseudo-triangulations, and the plane graph structure doesn't change.*

This completes the high-level sketch of the results and proofs.

References

1. Pankaj Agarwal, Julien Basch, Leonidas Guibas, John Hershberger, and Li Zhang. Deformable free space tilings for kinetic collision detection. *International Journal of Robotics Research*, 21:179–197, March 2003. Preliminary version appeared in Proc. 4th International Workshop on Algorithmic Foundations of Robotics (WAFR), 2000.

2. Oswin Aichholzer, Günter Rote, Bettina Speckmann, and Ileana Streinu. The zig-zag path of a pseudo-triangulation. In *Proc. 8th International Workshop on Algorithms and Data Structures (WADS)*, Lecture Notes in Computer Science 2748, pages 377–388, Ottawa, Canada, 2003. Springer Verlag.
3. Sergei Bespamyatnikh. Enumerating pseudo-triangulations in the plane. *Comput. Geom. Theory Appl.*, 30(3):207–222, 2005.
4. Jürgen Bokowski, Susanne Mock, and Ileana Streinu. The folkman-lawrence topological representation theorem for oriented matroids - an elementary proof in rank 3. *European Journal of Combinatorics*, 22:601–615, July 2001.
5. Craig Gotsman and Vitaly Surazhsky. Guaranteed intersection-free polygon morphing. *Computers and Graphics*, 25(1):67–75, 2001.
6. Jack Graver, Brigitte Servatius, and Herman Servatius. *Combinatorial Rigidity*. Graduate Studies in Mathematics vol. 2. American Mathematical Society, 1993.
7. Leonidas Guibas, John Hershberger, and Subash Suri. Morphing simple polygons. *Discrete and Computational Geometry*, 24:1–34, 2000.
8. Ruth Haas, David Orden, Günter Rote, Francisco Santos, Brigitte Servatius, Herman Servatius, Diane Souvaine, Ileana Streinu, and Walter Whiteley. Planar minimally rigid graphs and pseudo-triangulations. *Computational Geometry: Theory and Applications*, pages 31–61, May 2005.
9. Audrey Lee, Ileana Streinu, and Louis Theran. Finding and maintaining rigid components. In *Proc. Canad. Conf. Comp. Geom.*, Windsor, Canada, August 2005.
10. Michel Pocchiola and Gert Vegter. Topologically sweeping visibility complexes via pseudo-triangulations. *Discrete & Computational Geometry*, 16(4):419–453, Dec. 1996.
11. Günter Rote, Francisco Santos, and Ileana Streinu. Expansive motions and the polytope of pointed pseudo-triangulations. In Janos Pach Boris Aronov, Saugata Basu and Micha Sharir, editors, *Discrete and Computational Geometry - The Goodman-Pollack Festschrift*, Algorithms and Combinatorics, pages 699–736. Springer Verlag, Berlin, 2003.
12. Bettina Speckmann and Csaba Tóth. Allocating vertex π -guards in simple polygons via pseudo-triangulations. In *Proc. ACM-SIAM Symp. Discrete Algorithms (SODA)*, pages 109–118, 2003. To appear in *Discrete and Computational Geometry*, 2004.
13. Jorge Stolfi. *Oriented Projective Geometry: A Framework for Geometric Computations*. Academic Press, New York, NY, 1991.
14. Ileana Streinu. A combinatorial approach to planar non-colliding robot arm motion planning. In *IEEE Symposium on Foundations of Computer Science*, pages 443–453, 2000.
15. Ileana Streinu. Pseudo-triangulations, rigidity and motion planning. *Discrete and Computational Geometry*, to appear, 2005. A preliminary version appeared in [14].
16. Walter Whiteley. Some matroids from discrete applied geometry. In J. Oxley J. Bonin and B. Servatius, editors, *Matroid Theory*, volume 197 of *Contemporary Mathematics*, pages 171–311. American Mathematical Society, 1996.