A Generalized Carpenter's Rule Theorem for Self-Touching Linkages

Timothy G. Abbott*† Erik D. Demaine*‡ Blaise Gassend§

Abstract

The Carpenter's Rule Theorem states that any chain linkage in the plane can be folded continuously between any two configurations while preserving the bar lengths and without the bars crossing. However, this theorem applies only to strictly simple configurations, where bars intersect only at their common endpoints. We generalize the theorem to self-touching configurations, where bars can touch but not properly cross. At the heart of our proof is a new definition of self-touching configurations of planar linkages, based on an annotated configuration space and limits of nontouching configurations. We show that this definition is equivalent to the previously proposed definition of self-touching configurations, which is based on a combinatorial description of overlapping features. Using our new definition, we prove the generalized Carpenter's Rule Theorem using a topological argument. We believe that our topological methodology provides a powerful tool for manipulating all kinds of self-touching objects, such as 3D hinged assemblies of polygons and rigid origami. In particular, we show how to apply our methodology to extend to self-touching configurations universal reconfigurability results for open chains with slender polygonal adornments, and single-vertex rigid origami with convex cones.

*MIT Computer Science and Artificial Intelligence Laboratory, 32 Vassar St., Cambridge, MA 02139, USA, {tabbott,edemaine}@mit.edu

†Partially supported by an NSF Graduate Research Fellowship and an MIT-Akamai Presidential Fellowship.

‡Partially supported by NSF CAREER award CCF-0347776, DOE grant DE-FG02-04ER25647, and AFOSR grant FA9550-07-1-0538.

§Exponent Failure Analysis Associates, 149 Commonwealth Drive, Menlo Park, CA 94025, USA, blaise.gassend@m4x.org. Work done while at MIT.

1 Introduction

In the mathematics of geometric folding [O'R98, Dem00, DD01, DO05, DO07], a common idealization is to model the underlying real-world object—a mechanical linkage, robotic arm, protein, piece of paper, or another object or surface—as having zero thickness. The rods or bars that make up a linkage become perfect mathematical line segments of fixed length; the joints or hinges that connect them become mathematical points; a piece of paper can be folded repeatedly ad infinitum. While these idealizations are not entirely realistic (see [Gal02]), the zero-thickness model has led to a wealth of powerful theorems that rarely abuse the lack of thickness and are therefore practical.

Almost all forms of folding forbid folding objects from crossing, matching a natural physical constraint, but at the same time allow folding objects to touch. Figure 1 illustrates the distinction between touching and crossing. For example, overlapping multiple layers of paper enables origamists to form arbitrarily complicated shapes, both in practice and in theory [DDM00].

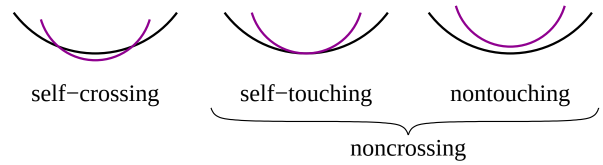

Figure 1: *The different types of configurations.*

Touching is easy to model for objects with positive thickness: allow the boundaries, but not the interiors, to intersect. But in the zero-thickness model, formally distinguishing between touching and crossing is difficult. In particular, when two portions of the object overlap, the geometry alone is insufficient to distinguish which portion is on top of which. The approach taken so far to resolving the ambiguity is to express the information missed by the geometry with additional combinatorial information. A simple example is map folding of an $m \times n$ grid of squares [ABD$^+$04]. In this context, the geometry of the squares is completely determined, independent of the folding: in any successful folding that uses all the creases, all of the squares will end up on top of each other, with orientations specified by a checkerboard pattern in the grid. The folding itself can be specified by a purely combinatorial object: the permutation of the panels that describes their total order in the folding. The challenge is to determine what constraints on this combinatorial object correspond to the paper not self-crossing. A generalization of this approach is essentially the one taken by [DDMO04] for defining general origami.

There are two concerns with this type of approach.

First, how do we know that the combinatorial definition corresponds to the intended meaning of self-touching configurations? The combinatorial definitions inherently lack geometric intuition, so it is hard to "feel" that they are correct, even though we believe they are.

Second, how do we manipulate these definitions to prove interesting theorems? The complexity of the definitions makes them hard to use. While some problems were successfully attacked in [CDR02, DDMO04], many other problems about self-touching configurations remain open. An alternate, equivalent definition would give a new way to examine and attack these problems.

Recently, touching has been studied for both linkages and origami. In the context of linkages, Connelly et al. [CDR02] show that self-touching configurations of linkages could be used to prove theorems about nearby non-self-touching perturbations. This result essentially reduces proving a planar linkage to be locked to an automatic, algorithmic procedure, whereas previous arguments that dealt solely with non-self-touching configurations were tedious and ad-hoc. To do this, [CDR02] introduced a combinatorial definition of touching linkages that we will describe later.

Many results in computational origami construct folded states with the desired properties, but do not show that the state can be reached by a continuous folding motion. Demaine et al. [DDMO04] showed that this was always possible. To do this, they defined origami using a combinatorial definition to handle self-touching folded states and their folding motions. This combinatorial definition turned out to be tedious to work with.

Our results. In this paper, we study the self-touching analog of the Carpenter's Rule Theorem, posed at FOCS 2000 [CDR03]. Consider a polygon or open polygonal chain in the plane, where the edges represent rigid bars of fixed length and the vertices represent hinges that can take arbitrary angles. Connelly et al. [CDR03] proved the Carpenter's Rule Theorem: every such linkage can be unfolded to a convex configuration while preserving connectivity, edge lengths, and without self-crossing (see also [Str00] for a more algorithmic approach). But it remained open whether this result held when the original configuration was self-touching (but not self-crossing). We solve this open problem, proving that the Carpenter's Rule Theorem extends to self-touching polygons and polygonal chains: every such linkage can be convexified starting from any (possibly self-touching) configuration.

The basis for this result is a new technique for defining self-touching configurations. Our approach is based on the intuition that self-touching configurations are limits of non-self-touching configurations. This intuition may seem obvious, but in its literal form, it is false. In the example of map folding, taking the limit to zero separation still, in the end, discards all of the information about the folding. Nonetheless, when working with self-touching configurations, people draw configurations with overlapping layers separated slightly for visibility, and imagine the limit as those separations go to zero.

We show how to turn this intuitive idea into a definition. Our main idea is to *annotate* the geometry of the configuration with additional continuous information. Previous combinatorial definitions could added annotations only at places where self-touching occurs, as needed to resolve ambiguity. In contrast, the topology of our configuration space places self-touching configurations near nontouching configurations. We therefore annotate all configurations, whether they are self-touching or not. Generally, the annotations are made up of the output of an order function applied to each ordered pair of independently mobile parts of the object to be modeled. For a linkage, bars are the independently mobile parts; for paper, each point of the paper is an independently mobile part.

Taking limits of annotated configurations is unfortunately insufficient to get all self-touching configurations. Figure 2(a) shows a linkage that has no nontouching configurations. In order to get these configurations, we allow flexibility in the limit-taking. That is, we allow the object to "stretch" while the limit is being taken. In the case of linkages, stretching means varying the length of the bars 2(b).

We believe that uniform annotation and stretchiness can be used to define self-touching configurations for a wide range of foldable objects. In this paper, we use our limit-based definition to prove the self-touching Carpenter's Rule Theorem using a primarily topological argument. At the end of the paper, we discuss how our topological approach might be applied to origami, rigid origami, and 3D hinged assemblies of planar panels.

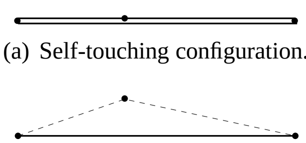

(a) Self-touching configuration.

(b) Stretched configuration.

Figure 2: *This self-touching configuration of a degenerate triangle linkage needs some added flexibility to be a limit of nontouching configurations.*

Outline. In Section 2, we review 2D linkages and define ε-related configurations of linkages. Then, in Section 3 we present an order function designed for linkages and apply it to self-touching linkage configurations. To increase the credibility of this definition, Section 4 compares it with the previously existing combinatorial definition and a variant on our definition that allows vertices to be split into a pair of vertices connected by a zero-length edge, and concludes that all three definitions are equivalent. The new definition is then used in Section 5 to prove the self-touching carpenter's rule theorem. We generalize this universal reconfigurability result to strictly slender polygonal adornments in Section 6. Finally, in Section 7 we discuss extensions of our definition methodology to objects that cannot be modeled as 2D linkages.

2 Linkage Preliminaries

This section will introduce the definitions that are important throughout this paper.

Definition 1 *A* linkage *is a pair $\mathcal{L} = (G, \ell)$ consisting of a graph G and a function $\ell : E(G) \to \mathbb{R}_{\geq 0}$ assigning nonnegative lengths to the edges. We will refer to the edges of G as* bars.

Definition 2 *A* configuration C *of a linkage* $\mathcal{L} = (G, \ell)$ *in the plane is a map* $C : V(\mathcal{L}) \to \mathbb{R}^2$ *obeying the length constraints, so if* $(v, w) \in E(G)$ *then* $|C(v) - C(w)| = \ell(v, w)$*. The set of all such configurations is called the* configuration space $\mathrm{Conf}(\mathcal{L})$ *of* $\mathcal{L}$*.*

Definition 3 *A* nontouching configuration *of* $\mathcal{L} = (G, \ell)$ *is a configuration in which no two edges intersect except at endpoints, and two endpoints coincide if and only if they are connected by a path in* G *of zero-length bars. Let* $\mathrm{NConf}(\mathcal{L}) \subset \mathrm{Conf}(\mathcal{L})$ *be the subspace of nontouching configurations.*

Simple linkages include *open chains* and *closed chains* for which the underlying graph is a single path, or a single loop, respectively.

Our definition of linkage is unusual in that it allows zero-length bars. Allowing zero-length bars will be necessary for some of our topological arguments, because without them, certain configuration spaces are not closed.

The definition we have taken for nontouching may yield surprising results with linkages having zero-length bars, as in Figure 3. Our definition considers pairs of vertices that are connected by a zero-length edge to be nontouching, effectively merging the two endpoints of such edges into a single vertex.

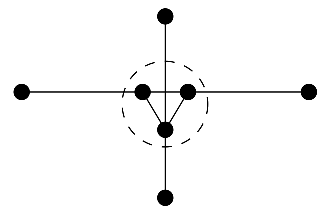

Figure 3: *This configuration, where the bars within the dotted circle have zero length, is nontouching, despite the apparent intersection between the upward-pointing edge and the topmost zero-length edge.*

2.1 ε-related Configurations

Definition 4 *For* $\varepsilon \geq 0$*, we say two linkages* (G_1, ℓ_1) *and* (G_2, ℓ_2) *are* ε-related *if* $G_1 = G_2$ *and*

$$|\ell_1(e) - \ell_2(e)| \leq \varepsilon \text{ for all } e \in E(G_1).$$

Definition 5 *For* $\varepsilon \geq 0$*, an* ε-related configuration *of a linkage* $\mathcal{L}$ *is a configuration of a linkage* $\mathcal{L}'$ *that is* ε*-related to* $\mathcal{L}$*.*

The set of ε-related configurations of $\mathcal{L}$ is denoted by $\mathrm{Conf}_\varepsilon(\mathcal{L})$. Notice $\mathrm{Conf}_0(\mathcal{L}) = \mathrm{Conf}(\mathcal{L})$. Similarly, $\mathrm{NConf}_\varepsilon(\mathcal{L})$ is the set of nontouching configurations of linkages ε-related to $\mathcal{L}$.

2.2 Real Algebraic Geometry

We will need to use some results from real algebraic geometry in this work. In particular, all of the objects discussed in this work will be semi-algebraic.

Definition 6 *A (real)* semi-algebraic set *is subset of* $\mathbb{R}^n$ *that is a finite Boolean combination of sets of the form* $\{x : f(x) * 0\}$ *where* f *is a polynomial of* n *variables, and* $*$ *is* $=$ *or* $<$*. A function* $f : \mathbb{R}^m \to \mathbb{R}^k$ *is* semi-algebraic *if its graph* $\{(x, f(x)) : x \in \mathbb{R}^m\} \subseteq \mathbb{R}^{m+k}$ *is a semi-algebraic set.*

In this work, we use a number of important topological properties of semi-algebraic sets. For a comprehensive reference on the theory of semi-algebraic sets, see [BCR98].

3 Noncrossing Configurations

3.1 The Order Function

In defining noncrossing configurations, we must allow two bars to "overlap". When this happens, the geometry of the configuration space does not have the information necessary to determine when they cross.

We define an order function, Ord, that determines the relative positions of two edges in noncrossing configurations. Its key property is that Ord is continuous wherever the edges do not touch, but Ord has a discontinuity where two edges share a common segment. The two different limits as the edges approach each other will encode the relative position information.

Definition 7 *Let e_1 and e_2 be two oriented edges in the plane. Using coordinates where the x-axis is directed along e_1, define*

$$d_+(e_1, e_2) = \textit{len}\{x \in e_1 \mid \exists y : y \geq 0, (x, y) \in e_2\};$$

$$d_-(e_1, e_2) = \textit{len}\{x \in e_1 \mid \exists y : y \leq 0, (x, y) \in e_2\}.$$

d_+ can be thought of as the length of the projection of the part of e_2 above e_1 onto the ***segment*** *e_1 (and similarly for d_-). See Figure 4. Define the* order function $\mathrm{Ord}(e_1, e_2) = d_+(e_1, e_2) - d_-(e_1, e_2)$.

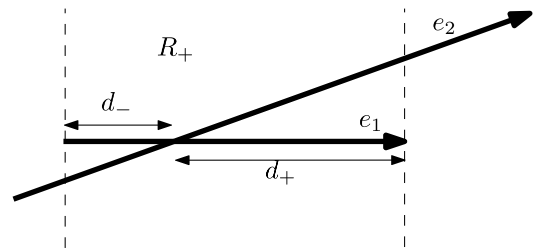

Figure 4: *Defining the order function.*

Notice that indeed Ord is not continuous when the two edges are tangent; but Ord is defined everywhere, so that $\mathrm{Ord}(e_1, e_1) = 0$.

Lemma 1 *The order function has the following properties:*

1. Ord *is a semi-algebraic function.*
2. Ord *is continuous for edges that do not intersect in their interior.*
3. *Consider two sequences of oriented edges e_1^n and e_2^n converging to edges e_1 and e_2 respectively, such that e_1^n and e_2^n do not intersect in their interior for any n (note that e_1 and e_2 might intersect in their interior). Then the sequence $\mathrm{Ord}(e_1^n, e_2^n)$ either converges to $\mathrm{Ord}(e_1, e_2)$, or has at most two accumulation points: d and $-d$, where d is the length of the overlap between e_1 and e_2.*

Proof: We will prove that d_+ is a semi-algebraic function that is continuous over edges that do not intersect in their interior; d_- will have these properties by a similar argument. We use the reference frame centered at the first vertex of e_1, with the x axis directed along e_1. In this reference frame, edge e_1 extends from $(0, 0)$ to $(l, 0)$.

Now, replace e_2 with the its intersection with the region $R_+ = \{(x, y) \mid y \geq 0, l \geq x \geq 0\}$. If the intersection is empty (a semi-algebraic condition), then $d_+ = 0$. Otherwise, the coordinates (x_1, y_1) and (x_2, y_2) are a semi-algebraic function of the endpoints of $e_2 \cap R_+$, since whether e_2 intersects each boundary of R_+ can be determined using line-segment intersections and thus can be tested with a boolean combination of polynomial inequalities. For each possible set of intersections, the points of intersection can be semi-algebraically computed from the original coordinates.

Clearly, $d_+ = 0$ if $l = 0$. The reader can check that $d_+ = |f(x_2/l) - f(x_1/l)|\, l$, where $f(z) = \max(\min(z, 1), 0)$. f is a boolean combination of polynomial inequalities, and thus d_+ is a semi-algebraic function everywhere.

Because d_+ is uniformly zero when e_2 does not intersect R_+, d_+ is continuous when $e_2 \cap R_+$ is empty. Notice that our formula for d_+ in terms of f would still be true if we had replaced e_2 with its intersection with the closed upper half plane H instead of R_+. Since f is continuous, d_+ is continuous everywhere where $e_2 \cap H$ is a continuous function of the coordinates of e_2. Similarly, d_+ is continuous everywhere where $e_2 \cap R_+$ is a continuous function of the coordinates of e_2. Thus, d_+ is continuous except when both endpoints of $e_2 \cap R_+$ lie along the boundary of H intersect the boundary of R_+. This exceptional case occurs only if e_2 and e_1 intersect in the interior. Thus d_+ is continuous everywhere except when e_2 intersects e_1 in the interior.

Now, consider sequences e_1^n and e_2^n defined in the statement of Part 3 of this lemma. By continuity, $\mathrm{Ord}(e_1^n, e_2^n)$ converges to $\mathrm{Ord}(e_1, e_2)$, unless the limits e_1 and e_2 share a common interval. If they do, notice that e_1^n and e_2^n are nontouching, and so for sufficiently large n, $\mathrm{Ord}(e_1^n, e_2^n)$ must be within ε of either l or $-l$, where l is the length of that common interval. The result follows. $\square$

3.2 Annotations

An annotated configuration is a configuration augmented with the values of Ord on each pair of edges in the configuration.

Definition 8 *Let C be a configuration of a linkage $\mathcal{L}$ with graph $G = (V, E)$. For $e_i = (u, v) \in E$, write $C(e_i) = (C(u), C(v))$.*

The annotation $M_C \in \mathbb{R}^{E \times E}$ *is the matrix whose i-j entry is* $\mathrm{Ord}(C(e_i), C(e_j))$.[1] *An* annotated configuration *is the pair (C, M_C). Let $Annot_{\mathcal{L}} : \mathbb{R}^{|V|} \to \mathbb{R}^{|V|} \times \mathbb{R}^{|E| \times |E|}$ be the function that maps a configuration C of $\mathcal{L}$ to the corresponding annotated configuration of $\mathcal{L}$.*

Fixing a configuration C of $\mathcal{L} = (G, \ell)$, $\mathrm{Annot}_{\mathcal{L}}(C)$ depends on $\mathcal{L}$ only on the graph G (since C defines the distances). Since ε-related linkages have the same graph, the notion of annotated configurations extends naturally to ε-related configurations.

Definition 9 *The set of* annotated nontouching configurations *is* $\mathrm{Annot}_{\mathcal{L}}(\mathrm{NConf}(\mathcal{L}))$. *For $\varepsilon \geq 0$, the set of* ε-related annotated nontouching configurations *is* $\mathrm{Annot}_{\mathcal{L}}(\mathrm{NConf}_{\varepsilon}(\mathcal{L}))$.

Lemma 2 *The annotation function $Annot_{\mathcal{L}}$ is injective, continuous over nontouching configurations, and semi-algebraic.*

Proof: The annotation function is injective, since the first component of $\mathrm{Annot}_{\mathcal{L}}(C)$ is C. Since the annotation function simply applies Ord to all pairs of edges in G, and by definition, in nontouching configurations no two edges intersect in their interior, the remaining properties follow directly from Lemma 1. □

3.3 Noncrossing Configurations

We are now ready to define noncrossing configurations in terms of limits of nontouching configurations.

Definition 10 *A* noncrossing configuration *of $\mathcal{L}$ is an element of* $\mathrm{Conf}_0(\mathcal{L}) \times \mathbb{R}^{|E(\mathcal{L})| \times |E(\mathcal{L})|}$ *that is the limit of a sequence of annotated nontouching configurations of linkages 1-related to $\mathcal{L}$. The space of noncrossing configurations of $\mathcal{L}$ is denoted $A(\mathcal{L})$.*

Equivalently, $A(\mathcal{L}) = (\mathrm{Conf}_0(\mathcal{L}) \times \mathbb{R}^{|E(\mathcal{L}) \times E(\mathcal{L})|}) \cap \overline{\mathrm{Annot}_{\mathcal{L}}(\mathrm{NConf}_1(\mathcal{L}))}$, *where $\overline{X}$ denotes the topological closure of X.*

The following characterization implies that replacing 1 with any $\varepsilon > 0$ in Definition 10 would define the same set.

Lemma 3 $A(\mathcal{L}) = \cap_{n=1}^{\infty} \overline{\mathrm{Annot}_{\mathcal{L}}(\mathrm{NConf}_{1/n}(\mathcal{L}))}$

Proof: Clearly, $\cap_{n=1}^{\infty} \overline{\mathrm{Annot}_{\mathcal{L}}(\mathrm{NConf}_{1/n}(\mathcal{L}))} \subset \overline{\mathrm{Annot}_{\mathcal{L}}(\mathrm{NConf}_1(\mathcal{L}))}$. To see that $\cap_{n=1}^{\infty} \overline{\mathrm{Annot}_{\mathcal{L}}(\mathrm{NConf}_{1/n}(\mathcal{L}))} \subset \mathrm{Conf}_0(\mathcal{L}) \times \mathbb{R}^{|E(\mathcal{L})| \times |E(\mathcal{L})|}$, notice that the length of any bar e in the left-hand side must differ from $\ell(e)$, the length of that bar in $\mathcal{L}$, by at most $1/n$ for all n, and thus must equal $\ell(e)$, so that we in fact have an annotated configuration of $\mathcal{L}$.

Conversely, if $x \in \mathrm{Conf}_0(\mathcal{L})$ is the limit of the sequence $\alpha_1, \alpha_2, \ldots, \alpha_i \in \mathrm{Annot}_{\mathcal{L}}(\mathrm{NConf}_1(\mathcal{L}))$, then for any n, there exists $N_0(n)$ such that for any $k > N_0(n)$, α_k is an annotated $1/n$-related configuration of $\mathcal{L}$, i.e. $\alpha_k \in \overline{\mathrm{Annot}_{\mathcal{L}}(\mathrm{NConf}_{1/n}(\mathcal{L}))}$. Since that set is closed, the limit x must also be in $\overline{\mathrm{Annot}_{\mathcal{L}}(\mathrm{NConf}_{1/n}(\mathcal{L}))}$. Thus $A(\mathcal{L}) \subset \cap_{n=1}^{\infty} \overline{\mathrm{Annot}_{\mathcal{L}}(\mathrm{NConf}_{1/n}(\mathcal{L}))}$, as desired. □

[1]We assume that the edges have been assigned some canonical orientation.

3.4 Semi-Algebraic

Theorem 1 *For $\varepsilon \geq 0$ the following sets are semi-algebraic:* $\mathrm{Conf}_\varepsilon(\mathcal{L})$, $\mathrm{NConf}_\varepsilon(\mathcal{L})$, $\mathrm{Annot}_\mathcal{L}(\mathrm{NConf}_\varepsilon(\mathcal{L}))$, $A(\mathcal{L})$.

Proof: We start with $\mathrm{Conf}_\varepsilon(\mathcal{L})$, which is defined by requiring each bar length to be within ε of its length in $\mathcal{L}$. For a bar between points (x_i, y_i) and (x_j, y_j), with a length l_k in $\mathcal{L}$, we have the following constraints:

$$\begin{aligned}(x_i - x_j)^2 + (y_i - y_j)^2 \leq &\quad (l_k + \varepsilon)^2, \\ (x_i - x_j)^2 + (y_i - y_j)^2 \geq &\quad (l_k - \varepsilon)^2 \quad \text{if } l_k \geq \varepsilon.\end{aligned}$$

Because these conditions are semi-algebraic, $\mathrm{Conf}_\varepsilon(\mathcal{L})$ is semi-algebraic.

$\mathrm{NConf}_\varepsilon(\mathcal{L})$ is defined like $\mathrm{Conf}_\varepsilon(\mathcal{L})$, but with additional nontouching constraints. We reuse a strategy presented in Equation (3.6) of [CDR02], based on the following idea: if two bars do not intersect, then one of the bars lies completely on one side of the other bar, i.e., both ends of the first bar are on the same side of the other bar. The condition in [CDR02] has to be slightly changed by making the inequalities in it strict, to prevent self-touching in addition to self-intersection.

Unfortunately, this condition is too strong, as it prevents bars from touching at their endpoints when the distance in the graph between the endpoints is zero. If there is a path $v_{i_0}, \ldots, v_{i_n}$ in the graph between two vertices v_{i_0} and v_{i_n}, then we can test that the distance between them in the graph is zero using the equation

$$\sum_{j=0}^{n-1} (x_{i_j} - x_{i_{j+1}})^2 + (y_{i_j} - y_{i_{j+1}})^2 = 0.$$

When this condition holds, we augment the strict nontouching condition by explicitly allowing v_{i_0} and v_{i_n} to touch. We allow this as long as the other vertex of each bar does not touch the other bar, or one of the bars has zero length. All these conditions can be expressed using polynomials.

Combining all these polynomial conditions, we find that $\mathrm{NConf}_\varepsilon(\mathcal{L})$ is semi-algebraic. It follows that $\mathrm{Annot}_\mathcal{L}(\mathrm{NConf}_\varepsilon(\mathcal{L}))$ is semi-algebraic, since it is the image of the semi-algebraic set $\mathrm{NConf}_\varepsilon(\mathcal{L})$ under the semi-algebraic annotation map (by Lemma 2). Similarly, $\mathrm{Annot}_\mathcal{L}(\mathrm{Conf}_\varepsilon(\mathcal{L}))$ is semi-algebraic.

Finally, $A(\mathcal{L})$ is the intersection of the topological closure of the semi-algebraic set $\mathrm{Annot}_\mathcal{L}(\mathrm{NConf}_1(\mathcal{L}))$ with the semi-algebraic set $\mathrm{Annot}_\mathcal{L}(\mathrm{Conf}_0(\mathcal{L}))$, and is thus semi-algebraic. □

s

4 Equivalent Definitions

In this section, we describe two definitions of noncrossing configurations that are equivalent to our definition. Section 4.1 describes a variation on our definition which is useful for analyzing the linkages with very short edges.

Section 4.2 describes a translation of the combinatorial definition of [CDR02] into our terminology, and proves that all three definitions are equivalent.

4.1 Extended Linkages

In this section, we introduce *extended linkages* and *extended noncrossing configurations*, and use them to define a natural and seemingly larger class of noncrossing configurations. We use extended linkages as a tool in our proof that our noncrossing configurations are the same as the noncrossing configurations of [CDR02], which we will refer to as combinatorial noncrossing configurations.

When defining noncrossing configurations as a limit of nontouching configurations, we could allow vertices to be split into two vertices with an edge of length at most ε between them. In the limit as $\varepsilon \to 0$, the extra edges have length 0, and we remerge their endpoints, so that the resulting configuration is naturally an element of $\mathrm{Annot}_\mathcal{L}(\mathrm{Conf}_\varepsilon(\mathcal{L}))$. Linkage extensions and reductions formalize this notion of splitting vertices into two vertices separated by a zero-length edge. A simple example is shown in Figure 4.1.

This formulation is useful for analyzing linkages that have very short edges, by understanding their self-touching limit. This technique is can be used to simplify a proof that there exists a locked orthogonal tree [CDD+07]. In particular, the orthogonal tree has horizontal edges of length at most ε, but it is easiest to argue it is rigid in the limiting case $\varepsilon = 0$, and then conclude for some small ε it is locked using the equivalence between extended configurations and combinatorial configurations along with results of [CDR02].

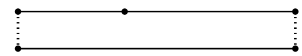

Figure 5: *Extended Linkage version of Figure 2*

Definition 11 *Given a linkage $\mathcal{L}$ with an edge $e = (u, v)$ of length 0, we can construct a* single-step reduced linkage *in which e has been removed and u and v have been merged into a single vertex. A* reduction *of a linkage $\mathcal{L}$ is the result of zero or more single-step reductions starting from $\mathcal{L}$. An* extension *of $\mathcal{L}$ is a linkage $\mathcal{L}'$ of which $\mathcal{L}$ is a reduction.*

One useful way to extend a linkage is to replace each vertex with one vertex per incident edge, all connected together by 0-length bars.

Given an annotated configuration C of a linkage $\mathcal{L}$, a configuration of the extended linkage $\mathcal{L}'$ can be generated by placing the fragments of a newly split vertex v where v was before, and setting the annotations of the new zero-length edge with other edges to 0. In every configuration of $\mathcal{L}'$, the two fragments must be in the same location, so this defines a correspondence between $\mathrm{Annot}_{\mathcal{L}}(\mathrm{Conf}(\mathcal{L}))$ and $\mathrm{Annot}_{\mathcal{L}'}(\mathrm{Conf}(\mathcal{L}'))$. However, the ε-related configurations of $\mathcal{L}'$ are a larger class than those of $\mathcal{L}$.

Definition 12 *An* extended noncrossing configuration *of $\mathcal{L}$ is the reduction to $\mathcal{L}$ of a noncrossing configuration of some extension $\mathcal{L}'$ of $\mathcal{L}$. The set of extended noncrossing configurations is denoted $E(\mathcal{L})$.*

4.2 Combinatorial Characterization of a Noncrossing Configuration

So far we have defined a noncrossing configuration of a linkage as an annotated configuration which is a limit of annotated nontouching configurations. This topological definition is markedly different from combinatorial definition given in [CDR02]. We shall now see that the annotated configurations which are nontouching can be characterized combinatorially as well as topologically. The combinatorial characterization will also allow us to see that in fact, $E(\mathcal{L}) = A(\mathcal{L})$, so extensions are not necessary to generate all nontouching linkage configurations.

We now present a combinatorial definition of noncrossing. This definition is expressed as constraints on configurations with an annotation matrix $(C, A) \in \mathrm{Conf}_0(\mathcal{L}) \times \mathbb{R}^{E(\mathcal{L}) \times E(\mathcal{L})}$. These constraints are equivalent to the constraints placed on noncrossing configurations in [CDR02], however, we have translated combinatorial configurations into our configuration space structure to help clarify the equivalence. We will not detail the (fairly straightforward) correspondence between our combinatorial formulation and theirs, though we will introduce “corridor segments”, “vertex locations”, and “magnified views” that directly relate to the edges, vertices, and magnified views in [CDR02].

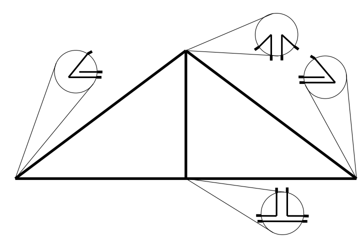

(a) Combinatorial noncrossing configuration.

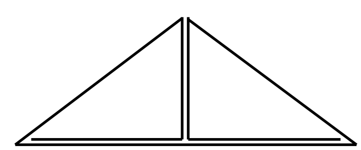

(b) Nearby ε-related nontouching configuration. A noncrossing configuration is a limit of these as $\varepsilon \to 0$.

Figure 6: *Combinatorial noncrossing configuration example.*

Definition 13 *For any configuration C of a linkage $\mathcal{L}$, we consider a magnified view around each* vertex location *(i.e., each point at which at least one vertex is located). For each vertex location, define the* inbounds *at that location as follows. There is one inbound per nonzero-length bar that has an endpoint co-located with the vertex location, and two inbounds per bar that goes through the vertex location. Inbounds to a vertex location are grouped into* entrances, *the directions from which they are incident. We write an inbound as a pair (θ, e) where θ is the entrance and e is the edge. Two inbounds to a vertex location are* directly connected *when there is a zero length path in* $\mathcal{L}$ *between them, including when they are part of the same bar that passes through the vertex location.*

Figure 6 gives a simple example that will be helpful for visualizing the situation. The disks are the magnified views, the directions from which lines approach the disks from outside are the entrances, and the intersections between the lines and boundaries of the disks are the inbounds. Two inbounds are directly connected if they are in the same connected component of the graphs inside the disks.

We define a combinatorial noncrossing configuration *of a linkage* $\mathcal{L}$ *to be a pair* $(C, A) \in \mathrm{Conf}_0(\mathcal{L}) \times \mathbb{R}^{E(\mathcal{L})\times E(\mathcal{L})}$ *which satisfies the following constraints:*

1. Macroscopically Noncrossing: *Bars* e_i *and* e_j *cannot have a strict crossing (they can touch at their endpoints or overlap over a finite length).*
2. Well-Annotated: *If* $i \neq j$ *and bars* e_i*,* e_j *overlap over a nonzero length* l*, then* $A_{i,j} = \pm l$*. Otherwise, set* $A_{i,j} = \mathrm{Ord}(e_i, e_j)$*.*[2]
3. Well-Ordered: *At each vertex location* v*, there is a total ordering* $\succeq$ *on inbounds* (θ, e_i)*, defined by the angle of* e_i *out from* v *with ties broken by the annotations. Let* $\mathbf{dir}(e)$ *be* $+1$ *if edge* e *is directed towards* v*, and* -1 *otherwise. Then* $\succeq$ *is defined as follows:*

$$(\theta_i, e_i) \succeq (\theta_j, e_j) \iff \begin{array}{ll} \theta_i > \theta_j & \textit{when } \theta_i \neq \theta_j \\ A_{i,j}\mathbf{dir}(e_i) \geq 0 & \textit{when } \theta_i = \theta_j \end{array} \tag{1}$$

   Notice that $A_{ij} = -A_{ji}$*.*
4. Microscopically Noncrossing: *The ordering of inbounds around a vertex location is compatible with the direct connections between those inbounds. More precisely, for inbounds* $t_1 \succeq t_2 \succeq t_3 \succeq t_4$*, if there are direct connections both between* t_1 *and* t_3*, and between* t_2 *and* t_4*, then all four inbounds are directly connected.*

We denote the space of combinatorial noncrossing configurations by $C(\mathcal{L})$*.*

Theorem 2 *The sets of noncrossing configurations, extended noncrossing configurations, and combinatorial noncrossing configurations are identical, i.e.* $E(\mathcal{L}) = A(\mathcal{L}) = C(\mathcal{L})$*.*

There are three inclusions to prove. $A(\mathcal{L}) \subset E(\mathcal{L})$ follows directly from the definitions. $E(\mathcal{L}) \subset C(\mathcal{L})$ will be proven by showing that all the conditions in the definition are indeed met, and $C(\mathcal{L}) \subset A(\mathcal{L})$ will be shown by constructing a converging sequence of nontouching configurations.

Definition 14 *A* δ*-perturbation of a combinatorial self-touching configuration* C *is a nontouching configuration in which each vertex is within* δ *of its location in* C *and the relative positions of the bars are preserved.*

Lemma 4 $C(\mathcal{L}) \subset A(\mathcal{L})$

Proof: Let $(C, A) \in C(\mathcal{L})$. By definition, $C \in \mathrm{Conf}_0(\mathcal{L})$. Consequently, it suffices to show $(C, A) \in \overline{\mathrm{Annot}_{\mathcal{L}}(\mathrm{NConf}_1(\mathcal{L}))}$. By Theorem 3.1 of chapter 1 of Ares Ribó Mor's Thesis [Mor], for any $C \in C(\mathcal{L})$, for any $\delta > 0$, there is a nontouching δ-perturbation C_δ of C. Since a δ-perturbation changes bar lengths by at most 2δ, $C_\delta \in \mathrm{Annot}_{\mathcal{L}}(\mathrm{NConf}_{2\delta}(\mathcal{L}))$. Because the relative positions of the bars are preserved in a δ-perturbation, and the annotation function is continuous for nontouching configurations, the C_δ converge to C as $\delta \to 0$. Thus $C \in \overline{\mathrm{Annot}_{\mathcal{L}}(\mathrm{NConf}_1(\mathcal{L}))} \Rightarrow C \in A(\mathcal{L})$, as desired. □

Since the argument of [Mor] is quite involved, we provide a simpler proof that $C(\mathcal{L}) \subset E(\mathcal{L})$ to give some intuition for this result. The basic strategy is to perturb the bars within each geometric location containing bars (which we call corridor segments) so that the bars within the corridor segment are parallel to each other and are ordered in a consistent fashion. We then use the information from the magnified views to implement the direct connections at the vertex locations. The details follow.

[2]Note that the annotations in combinatorial configurations are only meaningful for overlapping bars; we chose $A_{i,j} = \mathrm{Ord}(e_i, e_j)$ for nonoverlapping bars to simplify the statement of Theorem 2.

Definition 15 *A* corridor *of $\mathcal{L}$ is a line containing at least one bar of $\mathcal{L}$. A* corridor segment *is an interval in a corridor which has a vertex location at each end and no other vertex locations intersecting it.*

In Figure 6(a), the segments between vertex locations are the corridor segments, and the two corridor segments along the bottom combine to form a single corridor.

Since our combinatorial noncrossing configuration is well-annotated, the annotations define a total order on the bars within each corridor segment.

Lemma 5 *In a combinatorial noncrossing configuration C, each corridor S has a total ordering on its bars, that when restricted to any corridor segment is the order determined by the annotations on that corridor segment.*

Proof: We piece together the ordering for the corridor by proceeding down the corridor, successively merging the ordering so far with that of each corridor segment. At vertex location v, we can merge the ordering so far with the ordering for the next corridor segment if the these two orderings are consistent. Since each bar exists for a contiguous interval along the corridor, it suffices to check that the two corridor segments of S incident to v have consistent orderings. The claim that these orderings are consistent is a special case of the microscopic noncrossing condition at v. □

Lemma 6 $C(\mathcal{L}) \subset E(\mathcal{L})$

Proof: Suppose $(C, A) \in C(\mathcal{L})$. We will construct a sequence of nontouching configurations of an extension $\mathcal{L}'$ of $\mathcal{L}$, converging to an extension (C', A') whose reduction to $\mathcal{L}$ is (C, A).

$\mathcal{L}'$ is constructed by splitting from each vertex of $\mathcal{L}$ into one vertex for each incident edge (with each new vertex incident with its edge and a zero-length bar to the lexically first new vertex). Observe that pairs of vertices directly connected in $\mathcal{L}'$ are precisely those that are directly connected in $\mathcal{L}$. We will call the zero-length bars extension bars, and the others original bars. Since both endpoints of each extension bar are endpoints of original bars, specifying the locations of the original bars defines the configuration of $\mathcal{L}'$. In C', each vertex will (necessarily) lie in the same place as the vertex of C that it was split from.

Let $0 < \delta < \min(1/n, l_{\min}, (\sin\theta_{\min})/(2n))$ be a real number, where $l_{\min}$ is the minimum bar length in $\mathcal{L}$, $\theta_{\min}$ is the minimum angle between nonparallel bars in (C, A), and n is the number of bars in $\mathcal{L}$.

Let S be a corridor with m original bars in it (we treat extension bars as not belonging to any corridor). By Lemma 5, there is a total ordering on the bars in S compatible with the annotation orderings. Thus we can assign distinct offsets $\psi(e) \in \{0, 1, \ldots, m-1\}$ to the bars in S in a way compatible with the annotation orderings.

Arbitrarily select a unit vector $\overrightarrow{u}$ normal to S. Imagine shifting each original bar (v, w) contained in S from its location in C' by $\delta^2\psi((v, w))\overrightarrow{u}$. Consider also the circle of radius δ centered at $C'(v)$. $\delta^2\psi((v, w)) \leq \delta^2 n < \delta$, so v's shifted location is inside this circle. Because $\delta < l_{\min}$, w's shifted location is outside, so the shifted bar intersects v's circle exactly once. We set $C_\delta(v)$ to be this unique intersection of v's circle and shifted bar (and similarly for all the other vertices in S).

We now show that C_δ is nontouching for sufficiently small δ. Original bars never intersect extension bars except at common vertices because the former lie entirely outside the circles of radius δ, and the latter entirely inside. Intersections between original bars in a common corridor are impossible by construction. Since C_δ converges to C' as $\delta \to 0$, original bars that have nonzero separation in C' do not intersect in C_δ for small enough δ. It remains to handle pairs of original bars that touched in C' precisely at a vertex location v. Take two such bars, with offsets i and j, and with a relative angle of θ in C'. If they intersect in C_δ, it is at a distance at most $(i + j)\delta^2/\sin|\theta| \leq 2n\delta^2/\sin|\theta_{\min}| < \delta$ from $C'(v)$. Since neither bar intersects the circle of radius δ about v, no two original bars cross in C_δ.

We have constructed C_δ so that the orderings of vertices around the circles of radius δ are compatible with the ordering of inbounds at each vertex location. The microscopic noncrossing condition therefore forbids extension bars from crossing. Thus C_δ is noncrossing.

Having shown that C_δ is noncrossing and converges to C' as δ goes to 0, all that remains is to show that A_δ, the corresponding annotations, converge to A'. Because (C, A) is well annotated, Lemma 1 implies each annotation in A_δ for pairs of bars not sharing a corridor segment converges to the corresponding annotation in A'. By Lemma 1(3), the bars A_δ for pairs of bars sharing a corridor segment have accumulation points at $\pm$ the corresponding annotations in A'. But the offsets for bars in the corridors were chosen precisely so the signs of the annotations in A_δ matched the signs of annotations in A. Thus, the annotations converge to A'.

Taking any sequence of δs that converges to zero, we conclude that $C(\mathcal{L}) \subset E(\mathcal{L})$. □

Lemma 7 $E(\mathcal{L}) \subset C(\mathcal{L})$.

Proof: Take any extended noncrossing configuration $(c, a) \in E(\mathcal{L})$; we need to prove that (c, a) is macroscopically noncrossing, well-annotated, well-ordered, and microscopically noncrossing. Let (c_k, a_k) be a sequence of nontouching configurations of some extension $\mathcal{L}'$ of $\mathcal{L}$ that converges to an extension (c', a') of (c, a). The macroscopic noncrossing condition is easily met because the configurations in which bars have a strict crossing form an open set, so that a limit of nontouching configurations cannot have a strict crossing. The well-annotated condition follows immediately from Lemma 1(3) and the continuity of Ord over non-interior-intersecting edges.

To prove the well-ordered and microscopically noncrossing conditions, we draw small circles around each vertex location. Take δ small enough that, in c', the circle of radius 4δ drawn around a vertex location does not contain any other vertex locations, and does not intersect any edges that are not inbounds to the vertex location. For some k_0 and all $k \geq k_0$, each vertex is less than δ away from its final location, so each bar with nonzero length in c' crosses the circles corresponding to its endpoints, and each bar with length 0 in c' is contained within the circle that is common to both its endpoints. Furthermore, for some k_1 and all $k \geq k_1$, annotations in a_k that have nonzero limits have strictly the same sign as in a'. Henceforth, we assume that $k \geq k_0, k_1$.

Suppose e_i is an edge connecting vertex location v' to vertex location v. Let C_v be the circle centered at v with radius 2δ. Then e_i intersects C_v somewhere between vertex v and v'. We define the angle $\alpha_{i,k}$ to be the angle from a reference direction to this intersection between e_i and C_v. Without loss of generality, we may assume the reference direction is not $\lim_{k\to\infty} \alpha_{i,k}$ for any bar entering any vertex location in $\mathcal{L}$. Then because for k sufficiently large, e_i and e_j , there exists k_2 such that if $\lim_{k\to\infty} \alpha_{i,k} > \lim_{k\to\infty} \alpha_{j,k}$, then for all $k \geq k_2$, $\alpha_{i,k} > \alpha_{j,k}$. Henceforth, we assume that $k \geq k_2$.

We now define the necessary well-ordering. We say $Ent(E, e_i) \succeq Ent(E, e_j)$ if for all sufficiently large k, $\alpha_{i,k} > \alpha_{j,k}$. This is a well-ordering on inbounds at v for $k \geq k_2$. Set $\theta_i = \lim_{k\to\infty} \alpha_{i,k}$. Inbound edges e_i and e_j share the same entrance E at v if and only if $\theta_i = \theta_j$.

If $\theta_i > \theta_j$, then since $\lim_{k\to\infty} \alpha_{i,k} = \theta_i > \theta_j = \lim_{k\to\infty} \alpha_{j,k}$, for all sufficiently large k, $\alpha_{i,k} > \alpha_{j,k}$, and thus the well-ordering condition is satisfied in this case.

If inbound edges e_i and e_j share a common entrance, then in c' they overlap. The annotations now give the relationship between e_i and e_j in c_k. Assume e_i is oriented from v towards v', and $\mathrm{Ord}(e_i, e_j) > 0$ in c' (the other cases are symmetric). Then for sufficiently large k, in c_k, $\mathrm{Ord}(e_i, e_j) > 0$. Since in c_k, e_i and e_j are nontouching, it follows from the fact that $\alpha_{i,k} - \alpha_{j,k}$ goes to zero as $k \to \infty$ that $\alpha_{i,k} > \alpha_{j,k}$ for sufficiently large k. This completes the proof of the well-ordering condition.

Consider now the portion of the linkage in configuration c_k which is contained inside C_v. This portion of the linkage must be a planar graph, since c_k is noncrossing. Two intersection points of bars with C_v are connected by this graph if and only if the corresponding inbounds are directly connected. Given that the order of the intersection points around C_v matches the order of the inbounds to v, the fact that the graph is planar is precisely the microscopic noncrossing condition.

Thus $E(\mathcal{L}) \subset C(\mathcal{L})$. □

Theorem 2 now follows immediately from Lemmas 4 and 7 and the fact that $\mathcal{L}$ is a (trivial) extension of $\mathcal{L}$.

5 The Generalized Carpenter's Rule Theorem

The Carpenter's Rule Theorem says that any nontouching configuration of an open or closed chain linkage can be convexified through a continuous motion [CDR03] [Str00]. In this section, we use our definition of a noncrossing linkage to extend the Carpenter's Rule Theorem to all noncrossing linkages. That is, we shall show that when $\mathcal{L}$ is an open chain linkage, $A(\mathcal{L})$ is connected.

One might hope to show $A(\mathcal{L})$ is connected for closed chains. Such a result is not true, since the configuration space for a closed chain may have two connected components, one that turns clockwise, and one that turns counter-clockwise. We instead generalize by showing that any connected component of the noncrossing configuration space contains a connected component of the corresponding nontouching configuration space. To express this concept, we use the following definition:

Definition 16 *We say that a semi-algebraic set A* path-expands *a semi-algebraic subset B if every connected component of A contains a connected component of B.*

We will implicity use in the following discussion that connected components of semi-algebraic sets are path-connected and semi-algebraic [BCR98, Proposition 2.5.13].

Lemma 8 *If $\mathcal{L}$ is a chain linkage and $\mathrm{NConf}_0(\mathcal{L}) \neq \emptyset$, and $\varepsilon \geq 0$, then $\mathrm{NConf}_\varepsilon(\mathcal{L})$ path-expands $\mathrm{NConf}_0(\mathcal{L})$*

Proof: Each configuration C of a chain linkage has a corresponding canonical configuration. For an open chain, it is the straight configuration; for a closed chain, it is the configuration where the vertices are concyclic, turning in the same direction as C.

For any $C \in \mathrm{NConf}_\varepsilon(\mathcal{L})$, C is connected to its canonical configuration by the nontouching Carpenter's Rule Theorem. Thus, if we can show C's canonical configuration is connected to an element of $\mathrm{NConf}_0(\mathcal{L})$, it will follow that $\mathrm{NConf}_\varepsilon(\mathcal{L})$ path-expands $\mathrm{NConf}_0(\mathcal{L})$.

Fix a vertex location and an edge direction from that vertex (to factor out translations and rotations). Then there is a unique map sending each configuration $C \in \mathrm{NConf}_\varepsilon(\mathcal{L})$ to a corresponding canonical configuration $C' \in \mathrm{NConf}_0(\mathcal{L})$ turning in the same direction. By linearly interpolating between the canonical configuration for C and C', we obtain a path between the two in $\mathrm{Conf}_\varepsilon(\mathcal{L})$. Since at the endpoints of the path, each vertex is in convex position, the vertices remain in convex position along the path, so that the path is contained entirely in $\mathrm{NConf}_\varepsilon(\mathcal{L})$. Thus the canonical configuration for C is in the same connected component as $C' \in \mathrm{NConf}_0(\mathcal{L})$, as desired. □

Lemma 9 *Suppose a linkage $\mathcal{L}$ is connected, and there exists a $\delta > 0$ such that for all $\varepsilon \leq \delta$, $\mathrm{NConf}_\varepsilon(\mathcal{L})$ path-expands $\mathrm{NConf}_0(\mathcal{L})$. Then $A(\mathcal{L})$ path-expands $\mathrm{Annot}_{\mathcal{L}}(\mathrm{NConf}_0(\mathcal{L}))$.*

Proof: This lemma is trivially true if for some $\varepsilon > 0$, $\mathrm{NConf}_\varepsilon(\mathcal{L})$ is empty, in which case $A(\mathcal{L})$ is also empty. In the rest of this proof, we assume that this is not the case.

Suppose $n > 1/\delta$ (so that $0 < 1/n < \delta$). Since the annotation function is continuous on nontouching configurations, and $\mathrm{NConf}_{1/n}(\mathcal{L})$ path-expands $\mathrm{NConf}_0(\mathcal{L})$, $\mathrm{Annot}_{\mathcal{L}}(\mathrm{NConf}_{1/n}(\mathcal{L}))$ path-expands $\mathrm{Annot}_{\mathcal{L}}(\mathrm{NConf}_0(\mathcal{L}))$. It follows that the topological closure $\overline{\mathrm{Annot}_{\mathcal{L}}(\mathrm{NConf}_{1/n}(\mathcal{L}))}$ of $\mathrm{Annot}_{\mathcal{L}}(\mathrm{NConf}_{1/n}(\mathcal{L}))$ path-expands $\mathrm{Annot}_{\mathcal{L}}(\mathrm{NConf}_0(\mathcal{L}))$, since the closure of a connected set is connected.

Let $Y_1, \ldots, Y_r$ be the connected components of $\mathrm{Annot}_{\mathcal{L}}(\mathrm{NConf}_0(\mathcal{L}))$ (semi-algebraic sets always have finitely many connected components). Since $\overline{\mathrm{Annot}_{\mathcal{L}}(\mathrm{NConf}_{1/n}(\mathcal{L}))}$ path-expands $\mathrm{Annot}_{\mathcal{L}}(\mathrm{NConf}_0(\mathcal{L}))$, each connected component of $\overline{\mathrm{Annot}_{\mathcal{L}}(\mathrm{NConf}_{1/n}(\mathcal{L}))}$ contains one of the Y_j. Thus, we can write

$$\overline{\mathrm{Annot}_{\mathcal{L}}(\mathrm{NConf}_{1/n}(\mathcal{L}))} = \bigcup_{j=1}^{r} X_{j,n}$$

where $X_{j,n}$ is the connected component of $\overline{\mathrm{Annot}_{\mathcal{L}}(\mathrm{NConf}_{1/n}(\mathcal{L}))}$ containing Y_j. For a given n, two different $X_{j,n}$ are either disjoint or equal. Further, since $\overline{\mathrm{Annot}_{\mathcal{L}}(\mathrm{NConf}_{1/(n+1)}(\mathcal{L}))} \subset \overline{\mathrm{Annot}_{\mathcal{L}}(\mathrm{NConf}_{1/n}(\mathcal{L}))}$, $X_{j,n+1} \subset X_{j,n}$. By Lemma 3,

$$A(\mathcal{L}) = \bigcap_{n=1}^{\infty} \overline{\mathrm{Annot}_{\mathcal{L}}(\mathrm{NConf}_{1/n}(\mathcal{L}))} = \bigcap_{n=1}^{\infty} \bigcup_{j=1}^{r} X_{j,n} = \bigcup_{j=1}^{r} \bigcap_{n=1}^{\infty} X_{j,n}$$

where we can commute the intersection and union because for each j, the $X_{j,n}$ are a descending sequence. Now, $\bigcap_{n=1}^{\infty} X_{j,n}$ contains Y_j, since each $X_{j,n}$ does. Thus to prove our lemma, it suffices to show $\bigcap_{n=1}^{\infty} X_{j,n}$ is connected.

Notice that the $X_{j,n} \subset \overline{\mathrm{Annot}_{\mathcal{L}}(\mathrm{NConf}_{1/n}\,\mathcal{L})}$ are nonempty sets invariant under translation. By factoring out the translations via the choice of a point to place at the origin, we can write $X_{j,n} = \mathbb{R}^2 \times K_{j,n}$, where $K_{j,n}$ is nonempty, closed, and connected. Further, since $\mathcal{L}$ is connected, and our bars have bounded length, $K_{j,n}$ is bounded in $\mathbb{R}^N$ for $N = |V| + |E|^2 - 2$. It follows that $K_{j,n}$ is compact. Thus $\bigcap_{n=1}^{\infty} K_{j,n}$ is the intersection of a descending sequence of nonempty, compact, connected sets, and thus is a nonempty compact, connected set. Thus

$$\bigcap_{n=1}^{\infty} X_{j,n} = \bigcap_{n=1}^{\infty} \mathbb{R}^2 \times K_{j,n} = \mathbb{R}^2 \times \bigcap_{n=1}^{\infty} K_{j,n}$$

is the product of connected sets, and hence connected. Thus $A(\mathcal{L}) = \bigcup_{j=1}^{r} \bigcap_{n=1}^{\infty} X_{j,n}$ is a union of connected sets, each containing a connected component of $\mathrm{Annot}_{\mathcal{L}}(\mathrm{NConf}_0(\mathcal{L}))$, so $A(\mathcal{L})$ path-expands $\mathrm{Annot}_{\mathcal{L}}(\mathrm{NConf}_0(\mathcal{L}))$, as desired. $\square$

We are now ready for our main theorem.

Theorem 3 (Generalized Carpenter's Rule Theorem) *For any open chain linkage $\mathcal{L}$, $A(\mathcal{L})$ is connected. For any closed chain linkage $\mathcal{L}$, $A(\mathcal{L})$ has at most two connected components.*

Proof: Suppose $\mathrm{NConf}_0(\mathcal{L}) \neq \emptyset$. By the Carpenter's Rule Theorem, $\mathrm{NConf}_0(\mathcal{L})$ has one connected component if $\mathcal{L}$ is an open chain, and at most two if $\mathcal{L}$ is a closed chains. Applying Lemmas 8, and 9, we see that $A(\mathcal{L})$ also has this property.

If $\mathrm{NConf}_0(\mathcal{L}) = \emptyset$, then $\mathcal{L}$ must be a closed chain. If for some $\varepsilon > 0$, $\mathrm{NConf}_\varepsilon(\mathcal{L})$ is empty, then $A(\mathcal{L}) \subset \overline{\mathrm{NConf}_\varepsilon(\mathcal{L})}$ is empty, and so there are no connected components. The remaining case is that $\mathcal{L}$ has no nontouching configurations, but for any $\varepsilon > 0$, there are ε-related linkages to $\mathcal{L}$ with nontouching configurations. This happens precisely when $\mathcal{L}$ is a closed chain where one edge is equal in length to the sum of all the others, as in Figure 2. Such $\mathcal{L}$ have at most two noncrossing configurations (related by reflection), and thus at most two connected components. $\square$

6 Strictly Slender Polygonal Adornments

In [CDD+06], it is shown that chains with *slender adornments* satisfy an analogue of the Carpenter's Rule Theorem: every such open chain can be straightened and every such closed chain can be convexified. In this section, we show that strictly slender polygonal adornments of open chains satisfy a version of the self-touching Carpenter's Rule Theorem.

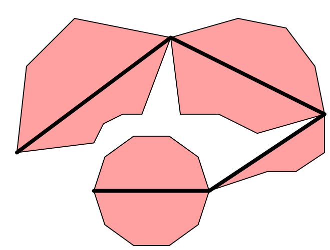

Figure 7: *A polygonally adorned chain with strictly slender adornments.*

Definition 17 *A* polygonal adornment R *is a compact, simply connected polygonal region, together with a* base B, *which is a distinguished line segment connecting two of its boundary points and contained in* R.

An inward normal *of a polygonal adornment is a ray X perpendicular to the boundary of the adornment starting from a point $x \in R \setminus B$ such that R contains a neighborhood of x in X.*

A polygonal adornment is strictly slender *if every inward normal intersects the relative interior of its base.*

A polygonally adorned chain *is set of polygonal adornments the bases of which form a chain.*

See Figure 7 for an example.

We can model a polygonally adorned linkage $\mathcal{L}$ as a linkage $\mathcal{L}'$ by replacing each polygonal adornment with a linkage triangulating it. This modeling is faithful because (1) any triangulation is rigid and (2) if two regions were to move from a nonoverlapping configuration to an overlapping one, the bars in the corresponding linkage defined by their boundaries would have to cross. We can thus reuse the topological machinery of Lemma 9 while replacing Lemma 8 with an analogue for linkages obtained by triangulating an open chain adorned by strictly slender polygonal adornments.

Lemma 10 *Suppose $\mathcal{L}$ is a linkage obtained by triangulating an open chain adorned by strictly slender polygonal adornments. Then there exists some $\delta > 0$ such that for all $\varepsilon < \delta$, any linkage ε-related to $\mathcal{L}$ is strictly slender.*

Proof: The property of being strictly slender is an open condition on the edge lengths of $\mathcal{L}$, and thus there exists a neighborhood of $\mathcal{L}$ that is contained in the set of strictly slender linkages. □

Lemma 11 *Suppose $\mathcal{L}$ is a linkage obtained by triangulating an open chain adorned by strictly slender polygonal adornments. Then for some $\delta > 0$, $\mathrm{NConf}_\varepsilon(\mathcal{L})$ path-expands $\mathrm{NConf}_0(\mathcal{L})$ for all $\varepsilon \leq \delta$.*

Proof: Our argument follows the paradigm of Lemma 8. The canonical configurations are those in which the chain is straight (if the chain has n adornments, there are potentially 2^n such canonical configurations, determined by the choices of reflection for each adornment). Thus, each configuration of $\mathcal{L}$ has a corresponding canonical configuration.

Suppose $\varepsilon < \delta$, for the δ defined in Lemma 10. Then for any $C \in \mathrm{NConf}_\varepsilon(\mathcal{L})$, C is connected to its canonical configuration by Theorem 8 of [CDD$^+$06]. Thus, if we can show C's canonical configuration is connected to an element of $\mathrm{NConf}_0\,\mathcal{L}$, it will follow that $\mathrm{NConf}_\varepsilon(\mathcal{L})$ path-expands $\mathrm{NConf}_0(\mathcal{L})$.

Since we have an open chain, there is a path linearly interpolating between the canonical configuration for C and the corresponding canonical configuration in $\mathrm{NConf}_0(\mathcal{L})$. This path is contained entirely in $\mathrm{NConf}_\varepsilon(\mathcal{L})$, since two different slender adornments with bases in straight configuration never touch except at the endpoints of the bases, and within each triangulated adornment, for sufficiently small ε there will be no crossings. Thus the canonical configuration for C is in the same connected component as $C' \in \mathrm{NConf}_0(\mathcal{L})$, as desired. □

Theorem 4 *Suppose $\mathcal{L}$ is a linkage obtained by triangulating an open chain adorned by strictly slender polygonal adornments. Then any configuration in $A(\mathcal{L})$ can straighten its base.*

Proof: The result follows from Lemmas 9 and 11, since any connected component of $\mathrm{NConf}_0(\mathcal{L})$ contains a configuration in which its base is straight by [CDD$^+$06]. □

Theorem 4 may be somewhat unsatisfying in that it has a number of restrictions on its applicability: "strictly slender", "polygonal", and "open chain", in particular. Closed chains are more complex because [CDD$^+$06] leaves open whether all convex configurations of closed adorned chains are reachable from each other. One should be able to handle arbitrary strictly slender adornments via a limit of polygonal adornment approximations.

The strictly slender restriction, however, seems fundamental to our argument, because Lemma 10 is false if one replaces "strictly slender" with "slender". For example, consider a linkage $\mathcal{L}$ obtained by triangulating a slender adornment which has 3 colinear vertices on its convex hull. The for any $\varepsilon > 0$, there are linkages ε-related to $\mathcal{L}$ that are not slender (one needs only shift the middle of the 3 colinear vertices inwards a bit while keeping the other vertices in place).

7 Extensions of the Definition Methodology

In this section, we outline how our methodology for defining self-touching configurations could be extended to other types of objects. The definitions suggested in this section are preliminary.

7.1 Polygonal Assemblies and Rigid Origami

A *polygonal assembly* is a set of polygons in 3-space and a relation indicating which polygonal edges are attached together. Polygonal assemblies arise in the study of *rigid origami*, where the edges correspond to creases.

Polygonal assemblies are a natural generalization of 2D linkages to 3D. Edges are replaced by polygons, and vertices are replaced by edges. We can define $\mathrm{Ord}(P_1, P_2)$ as the area of the projection of P_2 onto P_1, signed by which side of P_1 P_2 is on, in direct analogy with the linkage definition. With this definition of Ord, we can extend universal reconfigurability results for single-vertex rigid origami [SW05] to self-touching configurations (i.e., those in which two sheets are folded flat against each other). As in the case of Slender Adornments, we can only prove universal reconfigurability for self-touching configurations that are in an open subset of all configurations, in this case those for convex cones.

7.2 3D Linkages

It does not seem possible to model 3D linkages with our methodology. The difficulty is that the codimension of object elements is 2 for 3D linkages (compared with 1 for 2D linkages and polygonal assemblies). Consequently, there is a continuum of ways in which two bars can overlap (each relative direction is possible), and thus it seems no function has the necessary continuity properties to define the annotations.

7.3 Paper

Paper is a much more interesting challenge for our definition methodology than linkages or polygonal assemblies. Indeed, paper has an infinite-dimensional configuration space, so we have to worry about the right topology to use. Moreover, with paper, the individually movable pieces are infinitesimally small, so an order function that is zero when it is not directly above or below a piece would be zero everywhere.

We work with a unit-size n-dimensional closed sheet of paper in $(n+1)$-dimensional space. A (possibly self-crossing) configuration of order k of a sheet of paper is represented by a mapping f from $[0,1]^n$ to $\mathbb{R}^{n+1}$. The order k of the configuration indicates the regularity of the mapping; f must be piecewise C^k except along a finite set of C^k hyper-surfaces of finite hyper-area. To avoid stretching the paper, f must also be an isomorphism, i.e., wherever it is defined, its Jacobian must be an orthogonal projection of rank n.

A nontouching configuration is simply a configuration for which f is injective. We now consider an example order function; this is preliminary work.

7.3.1 "Distance with Obstacles" Order Function

First we consider the following order function that maps two points on the paper to a real number:

$$\mathrm{Ord}(a, b) = d_o(a_+, b) - d_o(a_-, b),$$

where $d_o(a_+, b)$ is the infinum of the lengths of the paths that start from the positive side of the paper at a and end at b without crossing the paper. This function is nice because it is continuous when the configuration is varied in a nontouching way, and when the two points a and b converge towards each other in a sequence of nontouching configurations, the order function converges to a limit that depends on the side of the paper from which b converges to a. The annotation function is produced by applying the order function to each pair of points on the paper. This defines a set of annotated nontouching configurations.

Before we can define noncrossing configurations, we need to specify a distance function that will define the topology we are using when taking limits. We define this topology over all functions like f, except that we do not impose the isomorphism constraint. The distance between f and g is defined by:

$$d(f, g) = \max(\sup |f(x) - g(x)|, \quad \sup \|D_f(x) - D_g(x)\|, \quad \sup |\,\mathrm{Ord}(f(x), f(y)) - \mathrm{Ord}(g(x), g(y))|)$$

where x and y are in $[0,1]^n$ and $D_f(x)$ is the Jacobian of f around x divided by the norm of the second order derivative of f around x. Essentially, d is the supremum norm applied to the functions, their derivatives and their annotations, except that the derivatives are scaled by a factor inversely proportional to the second order derivative of each function around that point (multiplied by zero at crease points). Without this scaling, points that should converge near creases would cause convergence problems.[3]

We can now define a noncrossing configuration as a limit of annotated noncrossing configurations for which the isometry constraint has been dropped. We drop the isometry constraint to ensure that we do not miss any self-touching configurations.

The order function we have chosen is somewhat tedious to work with because of its global nature. It would be nice to find a definition of the order function that only depends on f in a neighborhood of the points to which it is being applied. This is a difficult task because the paper can have very small features, so there is no canonical neighborhood size to take. One possibility might be to pick the largest neighborhood over which f behaves "nicely". This is all left for future work.

7.4 Conclusion

In this paper, we have introduced a new topological definition of self-touching 2D linkages. It is equivalent to the previously proposed definition, but is easier to work with. We showed how to use it to extend the carpenter's rule theorem to self-touching linkages.

This result demonstrates the advantages of our topological definition over the previously proposed combinatorial definition. The underlying topological methodology can be used to define the self-touching configurations of other classes of codimension 1 self-touching objects such as Origami and polygonal assemblies in terms of their nontouching configurations. Objects with a codimension of 2 or more seem more difficult to characterize.

References

[ABD+04] Esther M. Arkin, Michael A. Bender, Erik D. Demaine, Martin L. Demaine, Joseph S. B. Mitchell, Saurabh Sethia, and Steven S. Skiena. When can you fold a map? *Computational Geometry: Theory and Applications*, 29(1):23–46, September 2004.

[BCR98] J. Bochnak, M. Coste, and M-F. Roy, *Real Algebraic Geometry*, Springer, 1998.

[CDD+06] Robert Connelly, Erik D. Demaine, Martin L. Demaine, Sàndor Fekete, Stefan Langerman, Joseph S. B. Mitchell, Ares Ribó, and Günter Rote. Locked and Unlocked Chains of Planar Shapes. In Proceedings of the 22nd Annual ACM Symposium on Computational Geometry (SoCG 2006), Sedona, Arizona, June 5-7, 2006, pages 61-70.

[CDD+07] David Charlton, Erik D. Demaine, Martin L. Demaine, Gregory N. Price, Anastasios Sidiropoulos, Yaa-Lirng Tu. A Locked Orthogonal Tree. Manuscript, 2007.

[CDR02] Robert Connelly, Erik D. Demaine, and Günter Rote. Infinitesimally locked self-touching linkages with applications to locked trees. In Jorge Calvo, Kenneth Millett, and Eric Rawdon, editors, *Physical Knots: Knotting, Linking, and Folding of Geometric Objects in* $\mathbb{R}^3$, pages 287–311. American Mathematical Society, 2002. Collection of papers from the Special Session on Physical Knotting and Unknotting at the AMS Spring Western Section Meeting, Las Vegas, Nevada, April 2001.

[CDR03] Robert Connelly, Erik D. Demaine, and Günter Rote. Straightening polygonal arcs and convexifying polygonal cycles. *Discrete & Computational Geometry*, 30(2):205–239, September 2003.

[DD01] Erik D. Demaine and Martin L. Demaine. Recent results in computational origami. In *Origami*^{3}*: Proceedings of the 3rd International Meeting of Origami Science, Math, and Education*, pages 3–16, Monterey, California, March 9–11 2001. A K Peters.

[DDM00] Erik D. Demaine, Martin L. Demaine, and Joseph S. B. Mitchell. Folding flat silhouettes and wrapping polyhedral packages: New results in computational origami. *Computational Geometry: Theory and Applications*, 16(1):3–21, 2000. Special issue of selected papers from the 3rd CGC Workshop on Computational Geometry, 1998.

[DDMO04] Erik D. Demaine, Satyan L. Devadoss, Joseph S. B. Mitchell, and Joseph O'Rourke. Continuous foldability of polygonal paper. In *Proceedings of the 16th Canadian Conference on Computational Geometry*, pages 64–67, Montréal, Canada, August 2004.

[3]This problem could be solved in other ways such as by using an integral norm instead of the supremum norm (might allow kinks to remain in the paper), or by considering that the crease must form at its final location after a finite number of steps in the limit taking process.

[Dem00] Erik D. Demaine. Folding and unfolding linkages, paper, and polyhedra. In *Revised Papers from the Japan Conference on Discrete and Computational Geometry*, volume 2098 of *Lecture Notes in Computer Science*, pages 113–124, Tokyo, Japan, November 22–25 2000.

[DO05] Erik D. Demaine and Joseph O'Rourke. A survey of folding and unfolding in computational geometry. In Jacob E. Goodman, János Pach, and Emo Welzl, editors, *Combinatorial and Computational Geometry*, volume 52 of *Mathematical Sciences Research Institute Publications*, pages 167–211. Cambridge University Press, August 2005.

[DO07] Erik D. Demaine and Joseph O'Rourke. *Geometric Folding Algorithms: Linkages, Origami, and Polyhedra*. Cambridge University Press, 2007. To appear.

[Gal02] Britney C. Gallivan. *How to Fold Paper in Half Twelve Times*. Historical Society of Pomona Valley, Inc., Pomona, California, 2002. See http://www.pomonahistorical.org/12times.htm.

[Mor] Ares Ribó Mor. Realization and counting problems for planar structures. PhD Thesis, 2006.

[Mun00] James R. Munkres. Topology. Prentice Hall, 2000.

[O'R98] Joseph O'Rourke. Folding and unfolding in computational geometry. In *Revised Papers from the Japan Conference on Discrete and Computational Geometry*, volume 1763 of *Lecture Notes in Computer Science*, pages 258–266, Tokyo, Japan, December 1998.

[Str00] Ileana Streinu A Combinatorial Approach to Planar Non-Colliding Robot Arm Motion Planning In Proc. 41st ACM Annual Symposium on Foundations of Computer Science (FOCS) Redondo Beach, CA, Nov. 12-14, 2000, pp. 443-453.

[SW05] Ileana Streinu and Walter Whiteley. Single-Vertex Origami and Spherical Expansive Motions. Proc. Japan Conf. on Discrete and Comp. Geometry Tokai University, Tokyo, Oct. 8-11 2004. Lecture Notes in Computer Science 3742, pp. 161-173, Springer 2005.