1-1-2006

On the Maximum Span of Fixed-Angle Chains

Nadia Benbernou
Smith College

Joseph O'Rourke
Smith College, jorourke@smith.edu

Recommended Citation

Benbernou, Nadia and O'Rourke, Joseph, "On the Maximum Span of Fixed-Angle Chains" (2006). Computer Science: Faculty Publications, Smith College, Northampton, MA.
https://scholarworks.smith.edu/csc_facpubs/197

On the Maximum Span of Fixed-Angle Chains*

Nadia Benbernou† Joseph O'Rourke‡

January 1, 2008

Abstract

Soss proved that it is NP-hard to find the maximum 2D span of a fixed-angle polygonal chain: the largest distance achievable between the endpoints in a planar embedding. These fixed-angle chains can serve as models of protein backbones. The corresponding problem in 3D is open. We show that three special cases of particular relevance to the protein model are solvable in polynomial time. When all link lengths and all angles are equal, the maximum 3D span is achieved in a flat configuration and can be computed in constant time. When all angles are equal and the chain is simple (non-self-crossing), the maximum flat span can be found in linear time. In 3D, when all angles are equal to 90° (but the link lengths arbitrary), the maximum 3D span is in general nonplanar but can be found in quadratic time.

1 Introduction

Polygonal chains with fixed joint angles, permitting "dihedral" spinning about each edge, have been used to model the geometry of protein backbones [ST00] [DLO06]. Soss studied the *span* of such chains: the endpoint-to-endpoint distance. He proved that finding the minimum and the maximum span of planar configurations of the chain—the min and max *flat span*—are NP-hard problems [Sos01]. Protein backbones are rarely planar, so the real interest lies in 3D. Soss provided an example of a 4-chain whose maximum span (or *maxspan*) in 3D is not achieved by a planar configuration, establishing that 3D does not reduce to 2D. He designed an approximation algorithm, but left open the computational complexity of finding 3D spans.

Soss concentrated on the maxspan problem, and we do the same. We make progress on the 3D maxspan problem by focusing on restricted classes of chains, which are incidentally among the most relevant under the protein model.

Let a polygonal chain C have vertices $(v_0, v_1, \ldots, v_n)$. The fixed joint angle is $\alpha_i = \angle v_{i-1} v_i v_{i+1}$. Define an *$\alpha$-chain* as one all of whose joint angles are the

---

[*] Revised and expanded version of [BO06], based originally on [Ben06].

[†] Applied Mathematics, MIT. `nbenbern@mit.edu`.

[‡] Dept. Comput. Sci., Smith College, Northampton, MA 01063, USA. `orourke@cs.smith.edu`. Supported by NSF SGER Grant 0500290.

same angle α. Protein backbones can be crudely modeled as α-chains, with α obtuse, roughly in the range $[109^\circ, 122^\circ]$. Define a *unit chain*[1] as one all of whose link lengths are 1. Again roughly, protein backbones have equal-length links, because the bonds along the backbone lie roughly in the range [1.33Å,1.52Å].

We can summarize Soss's investigation in the first two lines of Table 1, and our results in the last three lines. We show that the 3D maxspan of a unit α-chain is achieved in a planar configuration, what we call the *trans-configuration*: a flat configuration in which the joint turns $\tau = \pi - \alpha$ alternate between $+\tau$ and $-\tau$. (The terminology is from molecular biology, which distinguishes between the trans- and cis-configurations of molecules.) We provide examples that show that, without the equal-length assumption, or without the equal-angle assumption, the maxspan configuration[2] might be nonplanar. For α-chains without the unit-length assumption, the simple flat maxspan is achieved by the trans-configuration, and can be found efficiently, in contrast to the arbitrary-α situation. Finally, we establish a structural theorem that characterizes the maxspan configuration of arbitrary fixed-angle chains in 3D, which permits the 3D maxspan of 90°-chains to be computed via a dynamic programming algorithm in $O(n^2)$ time.

Chain	dim	angles	lengths	complexity
fixed-angle chains	2	arbitrary	arbitrary	NP-hard
	3	arbitrary	arbitrary	?
unit α-chains	$2, 3$	$= \alpha$	1	$O(1)$
simple α-chains	2	$= \alpha$	arbitrary	$O(n)$
α-chains	3	$= 90^\circ$	arbitrary	$O(n^2)$

Table 1: Maxspan Computational Complexities.

2 Basic Lemmas for Arbitrary Chains

We start with two lemmas which hold for arbitrary joint angles and arbitrary link lengths.

Lemma 1 (3-Chain) *The maxspan of any fixed-angle* 3*-chain is achieved in a planar configuration.*

Proof: Let the chain be (v_0, v_1, v_2, v_3), and let β denote the angle between v_0v_2 and v_2v_3. Then the maximum distance between v_0 and v_3, $\max |v_0v_3|$, is achieved when β is largest, because the lengths $|v_0v_2|$ and $|v_2v_3|$ are already determined by the fixed edge lengths and fixed turn angles of the chain, leaving

---

[1] Our terminology is from [Poo06]. Also known as an *equilateral chain*.

[2] Whether or not there are several incongruent configurations that achieve the maxspan will not be relevant in this paper. We use "the" in referring to a maxspan configuration for convenience.

only β to vary. Now we just need to show that β is largest when v_3 is in the plane Π determined by $\{v_0, v_1, v_2\}$. See Fig. 1(a). Looking down on Π from above as in (b), it is clear that the segment that is the projection of the cone rim on which v_3 rides must cut the level curves transversely. For only if $\{v_0, v_1, v_2\}$ were collinear could it be parallel to the level curves, and then $\alpha = 0$ or π and the entire chain is contained in a line. Thus the v_3 projection intersects each level curve at most once, beginning at some intermediate β and ending at the maximum β in the plane Π. Hence max $|v_0 v_3|$ is achieved when v_3 lies in Π, and so the maximal configuration is planar. □

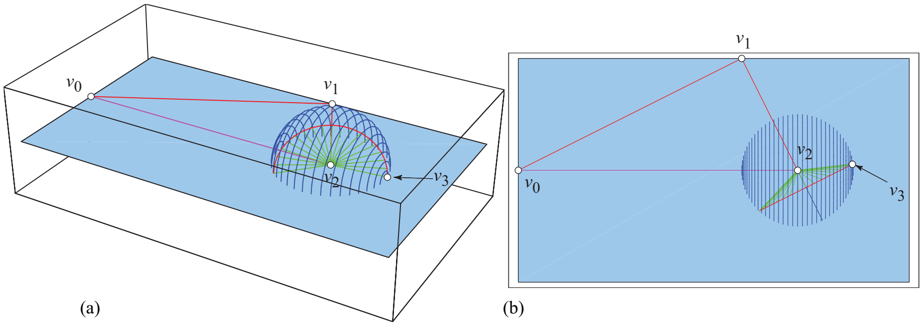

Figure 1: The maxspan of a fixed-angle 3-chain is achieved in a flat configuration. The rim of the cone is the locus of possible locations of v_3. The cone ribs specify all possible locations of edge $v_2 v_3$. The rings are the level sets for $\beta = \angle v_0 v_2 v_3$.

A near-immediate corollary is:

Lemma 2 (4-Vertex) *Let* $(v_0, v_1, \ldots, v_k)$ *be a fixed-angle* k*-chain. Then in any maximal configuration of the chain, vertices* $\{v_0, v_1, v_2, v_k\}$*, and vertices* $\{v_0, v_{k-2},\ v_{k-1}, v_k\}$ *are coplanar.*

Proof: We prove the latter claim; the former follows by relabeling the vertices in reverse. Let Π be the plane determined by $\{v_0, v_{k-2}, v_{k-1}\}$. As in the proof of the previous lemma, let β denote the angle between $v_0 v_{k-1}$ and $v_{k-1} v_k$. Any position of the three vertices $\{v_0, v_{k-2}, v_{k-1}\}$ in Π determine a "virtual" 3-chain $(v_0, v_{k-2}, v_{k-1}, v_k)$ whose span is maximized when v_k lies in Π (i.e., when β is largest) by Lemma 1. That is to say, for any such position, rotating v_k into the planar trans-configuration of the corresponding 3-chain yields the largest distance between v_0 and v_k for those particular positions of the vertices v_0, v_{k-2}, and v_{k-1}. This rotation is always possible because the cone on which $v_{k-1} v_k$ rides is centered on the line through $v_{k-2} v_{k-1}$, which lies in Π. Hence, in any maximal configuration, we must have $\{v_0, v_{k-2}, v_{k-1}, v_k\}$ coplanar; otherwise we could increase the distance between v_0 and v_k by rotating v_k into Π. □

Note that this lemma does not imply that the maxspan configuration of a 4-chain is planar, only that four of the five vertices lie in a plane.

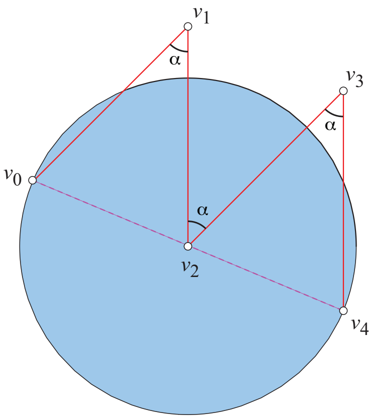

Figure 2: The maximal configuration of a unit, α-, 4-chain. The maxspan is $2\,|v_0 v_2|$.

3 Unit α-Chains

Now we specialize to unit α-chains. Our first lemma will serve as the base case in an induction proof to follow.

Lemma 3 *The 3D maxspan of a unit α-chain of 4 links is achieved by the trans-configuration.*

Proof: Let $(v_0, v_1, v_2, v_3, v_4)$ be such a chain. Let Π be the plane determined by $\{v_0, v_1, v_2\}$. Draw a sphere of radius $|v_0 v_2|$ centered at v_2. Because $|v_2 v_4| = |v_0 v_2| = 2 \sin \frac{\alpha}{2}$, v_4 must also lie on this sphere. By Lemma 2, we know that v_4 must also lie in Π. Hence v_4 must lie on the equatorial great circle that is the intersection of Π with the sphere. See Fig. 2. The maximum distance between v_0 and v_4 is just the diameter of this circle, i.e., $|v_0 v_2| + |v_2 v_4| = 2|v_0 v_2|$. And since the planar trans-configuration achieves this distance, we have that the trans-configuration is a maximal configuration. $\square$

This lemma is false without either the unit-length or the same-angle assumptions: See Fig. 3.

We now focus on unit α-chains of an arbitrary number of links. Our argument is easier for an even number of links than it is for an odd number of links.

Lemma 4 *The 3D maxspan of a unit α-chain, having an even number k of links, is achieved by the planar trans-configuration.*

Proof: We will prove this by induction. The base case $n = 4$ is achieved by the planar trans-configuration by Lemma 3 above. Assume it holds for all even $n \leq k - 2$ that a maximal configuration of a unit α-chain with n links is the planar trans-configuration, i.e., $\max |v_0 v_n|$ is achieved in the planar trans-configuration. Now we'll show that this is true for $n = k$ by using a "subadditive" argument.

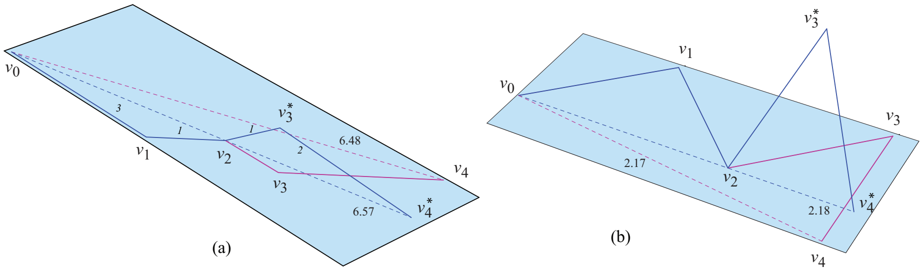

Figure 3: (a) Non-unit, 135°-chain whose maxspan configuration is nonplanar. (b) Unit chain with non-equal joint angles $(90°, 90°, 45°)$, whose maxspan configuration is nonplanar. (This latter example is effectively equivalent to Soss's example [Sos01, Fig. 6.9] mentioned in Sec. 1.)

Because the distance $|v_{k-2}v_k|$ is uniquely determined from the joint angle α,

$$\max |v_0 v_k| \leq \max |v_0 v_{k-2}| + |v_{k-2} v_k|$$

For if $\max |v_0 v_k|$ were larger than this quantity, the fixed distance $|v_{k-2}v_k|$ would imply that $\max |v_0 v_{k-2}|$ is not in fact maximal. By induction, $\max |v_0 v_{k-2}|$ is achieved in the planar trans-configuration. The planar trans-configuration of the full k-chain gives us equality in the above expression, so this must be a maximal configuration since $|v_0 v_k|$ can be no larger. See Fig. 4(a). $\square$

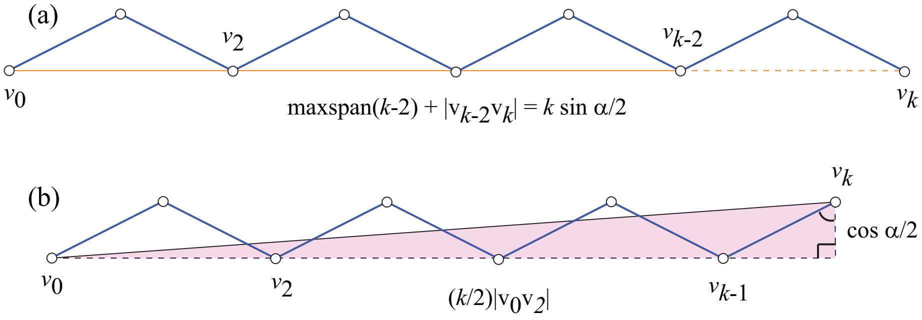

Figure 4: Maximal configurations of unit α-chains: (a) even; (b) odd.

Lemma 5 *The 3D maxspan of a unit α-chain, having an odd number k of links, is achieved by the planar trans-configuration.*

Proof: Proving this result for odd k is significantly more difficult. We will again use induction. Our base case is a unit 3-chain, which we know we know has the planar trans-configuration for its maximal configuration by Lemma 1. Assume it is true for all odd $n \leq k-2$ that a maximal configuration of a unit

n-chain with turn angles α is the planar trans-configuration. We will now show true for $n = k$.

Let Π be the plane determined by vertices $\{v_{k-2}, v_{k-1}, v_k\}$. We will show that the position of v_0 that maximizes $|v_0 v_k|$ is that of the planar trans-configuration. By the 4-Vertex lemma (Lem. 2), we know that v_0 must also lie in Π if we are to achieve a maximal configuration. Let maxspan $|m|$ denote the max span of a unit α-chain with m links. Let transspan $|m|$ denote the span of the trans-configuration of such a chain.

Draw a circle C_{k-2} in Π of radius maxspan $|k-2|$ centered at v_{k-2}. We know by the induction hypothesis that this radius is just the span of the trans-configuration, that is, maxspan $|k-2|$ = transspan $|k-2|$. Similarly draw a circle C_{k-1} of radius maxspan $|k-1|$ centered at v_{k-1}. Now because $k-1$ is even, maxspan $|k-1|$ = transspan $|k-1|$ by Lemma 4. Finally, draw a circle C_k of radius transspan $|k|$ centered at v_k. It is clear that these three circles C_{k-2}, C_{k-1}, and C_k must intersect at a common point v^*, since any subchain of a trans-chain is itself trans, and all three circles are based on trans-configurations. This construction is displayed in Fig. 5.

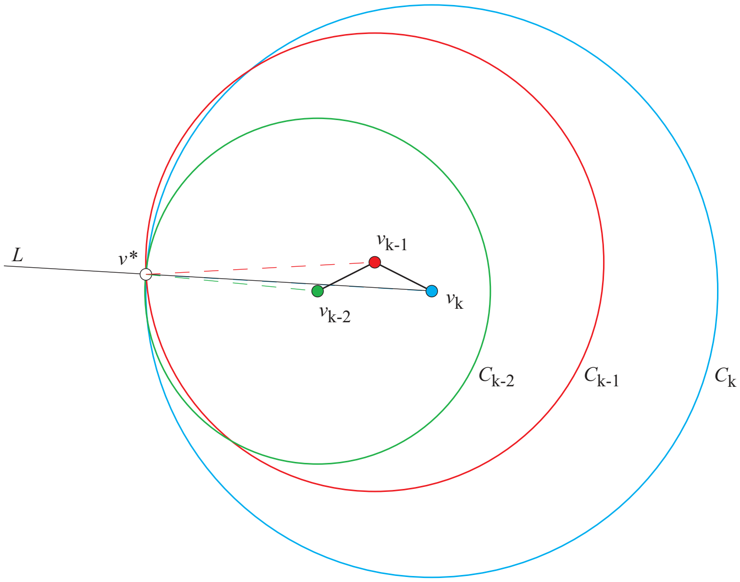

Figure 5: C_{k-2} is a circle of radius maxspan $|k-2|$ = transspan $|k-2|$ centered at v_{k-2}, C_{k-1} is a circle of radius maxspan $|k-1|$ = transspan $|k-1|$ centered at v_{k-1}, and C_k is a circle of radius transspan $|k|$ centered at v_k.

We aim to prove that the maxspan $|k|$ is achieved when $v_0 = v^*$, the position of v_0 when $(v_0, \ldots, v_k)$ is in the trans-configuration. Suppose for contradiction that there is a position of v_0 for which $|v_0 v_k| > |v^* v_k|$. Then v_0 is exterior to C_k Let L denote the line through v^* and v_k. If L also passes through v_{k-2}, then the last two links exactly extend the trans-configuration of the first $k-2$ links, and we are finished. Note this is because we have the upperbound

$$\operatorname{maxspan}|k| \leq \operatorname{maxspan}|k-2| + |v_{k-2}v_k| = \operatorname{transspan}|k-2| + |v_{k-2}v_k|,$$

and so if L passed through v_{k-2} we would achieve equality in that expression. So assume L misses v_{k-2}, and in particular, intersects $v_{k-2}v_{k-1}$.

That this intersection is without a loss of generality can be seen by the following reasoning. Orient the trans-configuration of the chain horizontally (that is to say, the x-axis bisects each link of the chain and so the y-coordinates of each vertex alternate between $+y$ and $-y$) as in Fig. 6, with v^* having y-coord $+y$. Both v_{k-2} and v_k have y-coord $-y$, and v_k lies to the right of v_{k-2}; hence the line L through v^*v_k is above the line v^*v_{k-2}. And the y-coordinate of v_{k-1} is $+y$, so the line determined by v^*v_{k-1} is horizontal. Hence L is sandwiched between the lines along v^*v_{k-2} and v^*v_{k-1} and must intersect $v_{k-2}v_{k-1}$ by continuity.

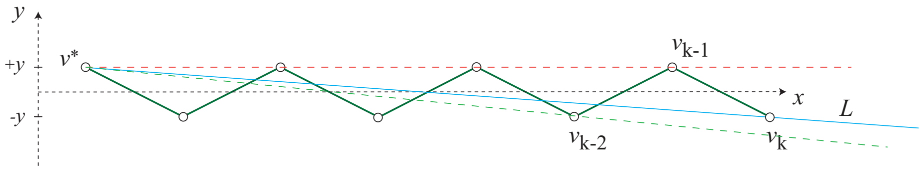

Figure 6: L must intersect $v_{k-2}v_{k-1}$.

Recall that we have supposed for contradiction that $|v_0v_k| > |v^*v_k|$, and hence that v_0 is exterior to C_k. We have two cases to consider.

Case 1: v_0 is above L and exterior to C_k. Because v_{k-2} lies below L and the radius $|v^*v_{k-2}|$ of C_{k-2} is smaller than that of C_k, C_{k-2} lies interior to C_k above L. Hence, v_0 is exterior to C_{k-2}, which contradicts our assumption that $\operatorname{maxspan}|k-2| = \operatorname{transspan}|k-2|$.

Case 2: v_0 is below L and exterior to C_k. Because v_{k-1} is positioned above L and the radius $|v^*v_{k-1}|$ of C_{k-1} is smaller than that of C_k, C_{k-1} lies interior to C_k below L. Hence, v_0 is exterior to C_{k-1}, which contradicts our assumption that $\operatorname{maxspan}|k-1| = \operatorname{transspan}|k-1|$.

Hence v_0 must lie interior or on the boundary of C_k. Thus we have

$$|v_0v_k| \leq |v^*v_k| = \operatorname{transspan}|k|$$

so the maximum of $|v_0v_k|$ is achieved by taking $v_0 = v^*$. And since v^* corresponds to the planar trans-configuration of the k-chain, we have that a maximal configuration of the k-chain occurs in the trans-configuration as desired. □

Putting Lemmas 4 and 5 together, we obtain:

Theorem 6 (Unit α-Chain) *The 3D maxspan of any unit α-chain is achieved in the planar trans-configuration.*

It now follows easily from Fig. 4 that computing the maxspan for unit α-chains takes constant time, the third row of Table 1.

4 Maximum Flat Span of α-Chains

Although Theorem 6 fails without the unit-length assumption, if we restrict an α-chain to the plane, then there are two conditions under which we can prove that the max flat span is still the trans-configuration: when the chain is simple, i.e., non-self-crossing; or when α=90°. The latter result is straightforward, and we establish that first.

4.1 Flat 90°-Chains

Let the 90°-chain be $C = (v_0, v_1, \ldots, v_n)$, and let ℓ_i be the length of link $v_i v_{i+1}$. Let $L_e = \ell_0 + \ell_2 + \ell_4 + \cdots$ be the sum of the even-indexed link lengths, and $L_o = \ell_1 + \ell_3 + \ell_5 + \cdots$ be the odd sum. Establish the convention that v_0 is the origin and v_1 is on the positive x-axis. Then, in the special case when $\alpha = 90°$, all the even links are horizontal, and all the odd links vertical, regardless of whether the angle turn is $+90°$ or $-90°$ at any joint. The trans-config yields (L_e, L_o) for the coordinates of v_n.

Call an edge of a chain a *cis-edge* if both turns at its endpoints are in the same direction. Now consider the same lengths ℓ_i forming a 90°-chain $C' = (v'_0, \ldots, v'_n)$ with at least one cis-edge. Then v'_n has either x- or y-coordinate strictly less than L_e or L_o respectively. Moreover, because of the convention that $v'_0 v'_1$ is horizontal to the right, the x-coordinate of v'_n is at least $-L_e + \ell_0$. Therefore, the absolute value of either the x- or y-coordinate of v'_n is strictly smaller than that of v_n, and the other coordinate is no larger. Therefore $|v_0 v'_n| < |v_0 v_n|$, and we have established the claim:

Lemma 7 *The max flat span of an α-chain with α=90° is achieved in the trans-configuration.*

4.2 Simple Flat α-Chains

Deviating from α=90° changes the analysis considerably. Our goal in this section is to prove this claim:[3]

Theorem 8 (Simple Flat Maxspan) *If C is an α-chain then the simple flat maxspan of C is realized by the trans-configuration.*

See Fig. 7 for an example. Fig. 8 shows that the qualifier "simple" is necessary: there exist self-crossing α-chains whose cis-configuration (all angle turns in the

---

[3] This corrects [BO06, Thm. 5], which erroneously claimed the result for all α-chains.

same direction) has a longer span than its trans-configuration.[4] We will return to this point below.

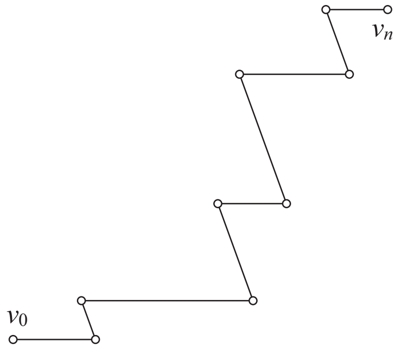

Figure 7: Planar simple trans-configuration of an α-chain with with acute α.

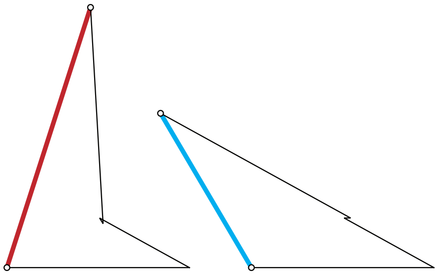

Figure 8: A 4-chain with link lengths $(7, 4, 0.2, 8)$ and $\alpha{=}30°$. The cis-span is approximately 10, while the trans-span is 7.

We prove Theorem 8 via two reflection techniques, which we establish in the next two lemmas (and which will be employed in Sec. 5 as well). Let the α-chain be $C = (v_0, v_1, \ldots, v_n)$, and let L be the line containing $v_0 v_n$, which we take to be horizontal for convenience.

Lemma 9 (Reflect) *If there is any edge of the chain $v_k v_{k+1}$ whose containing line $M \supset v_k v_{k+1}$ has $\{v_0, v_n\}$ strictly to the same side, then reflection of the suffix chain $(v_{k+1}, \ldots, v_n)$ across M creates a new α-chain C' with a larger span: $|v_0 v'_n| > |v_0 v_n|$, where v'_n is the reflected position of v_n.*

Proof: Note that it cannot be that either $k = 0$ or $k + 1 = n$, because then either v_0 or v_n would not be strictly to one side of M.

See Fig. 9(a) for a typical instance of the situation described in the lemma. The line M is the bisector of $v_n v'_n$, and so constitutes the Voronoi diagram of the two points $\{v_n, v'_n\}$. Because v_0 is to the same side of M as v_n, it is in v_n's Voronoi cell. Thus v_0 is closer to v_n than it is to v'_n, which is the claim of the lemma. The turn angle at v_{k+1} is negated, and otherwise all angles remain the same. Therefore, C' is an α-chain. $\square$

Lemma 10 (Reflect-Translate) *Suppose there is a pair of parallel edges in the chain C, $v_k v_{k+1}$ and $v_m v_{m+1}$, $k < m$, such that the line $M \supset v_k v_{k+1}$ does not have both $\{v_0, v_n\}$ strictly to the same side. Then reflection of the chain $(v_{k+1}, \ldots, v_m)$ across M, plus rigid attachment of $(v_{m+1}, \ldots, v_n)$, creates a new α-chain C' with a larger span than C.*

Proof: First, we may assume that $\{v_0, v_m\}$ are to the same side of M, for if instead $\{v_m, v_n\}$ are to the same side, relabeling C in reverse switches the roles

---

[4] In this example, α is acute, but we also found similar examples for obtuse α.

of v_0 and v_n. Second, it is also no loss of generality to assume v_0 is on or below M and v_n on or above, as in Fig. 9(b), for reflection about L switches above to below. In the illustrated situation, the reflection of the middle subchain $(v_{k+1}, \ldots, v_m)$ is upward and to the right. The suffix chain $(v_{m+1}, \ldots, v_n)$ is translated rigidly, maintaining the angle at v_{m+1}. The effect is to displace v_n upward and to the right. Now the bisector M' of $v_n v'_n$ has $\{v_0, v_n\}$ strictly to one side and v'_n to the other, and the argument in the Reflect lemma (Lem. 9) applies to show the span has increased. C' is an α-chain because only the turn angle at v_{k+1} changes, and that is negated. □

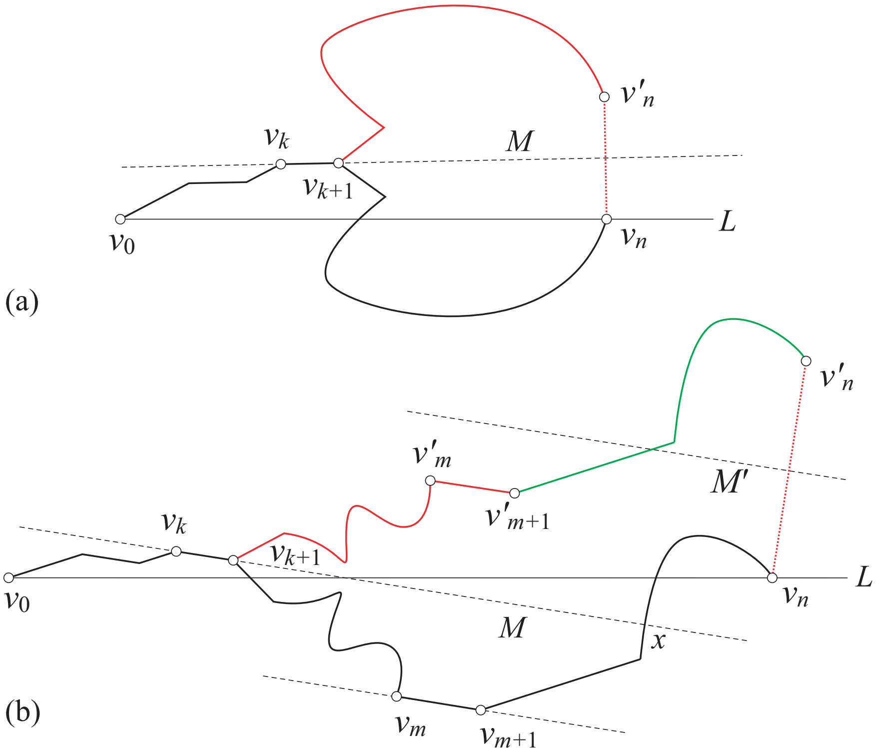

Figure 9: (a) Reflect Lemma 9; (b) Reflect-Translate Lemma 10.

We classify chains C into three types:

1. *extremity-crossing*: Chains that cross $L \setminus v_0 v_n$, i.e., cross L outside of $v_0 v_n$.
2. *self-crossing spirals*: a curve with at least one loop caused by a self-crossing.
3. All other chains.

A non-self-crossing "spiral" is necessarily extremity-crossing (but not all extremity-crossing chains are spirals, under any natural definition of "spiral.") We have already seen in Fig. 8 that the theorem does not always hold for self-crossing spirals. We next establish the theorem for extremity-crossing chains.

Lemma 11 (Extremity-Crossing) *No extremity-crossing α-chain can be in maxspan configuration.*

Proof: Assume C crosses L outside of the interval $[v_0, v_n]$ with edge $v_k v_{k+1}$. Then the line M determined by this edge has $\{v_0, v_n\}$ strictly to the same side, and so the Reflect lemma (Lem. 9) applies and establishes the claim. □

We are finally ready to prove the Simple Flat Maxspan theorem (Thm. 8).

Proof: In view of Lemma 11, we may assume that C does not cross $L \setminus v_0 v_n$. Let $v_k v_{k+1}$ be a cis-edge of C, with containing line $M \supset v_k v_{k+1}$. If M has $\{v_0, v_n\}$ to the same side, apply the Reflect lemma (Lem. 9). So now assume M has v_0 and v_n on opposite sides. Because $v_k v_{k+1}$ is a cis-edge, v_{k-1} and v_{k+2} lie to the same side of M. Assume without loss of generality that these two vertices are on the v_0 side of M (as in Fig. 9(b)); if they are on the v_n side, relabeling the chain in reverse puts them on the v_0 side. We now argue for the existence of another edge $v_m v_{m+1}$ parallel to M to the v_0-side.

We know that the suffix portion C_1 of C beyond v_{k+1} must eventually cross M to reach v_n on the other side. Because C is simple, C_1 cannot cross the prefix chain $(v_0, \ldots, v_{k+1})$, and it cannot cross L left of v_0 (since no extremity crossing chain can be in maxspan configuration). Therefore the first crossing of C_1 and M must be right of v_{k+1} on M, say at x. Now C_1 up to x plus the segment $x v_{k+1}$ forms a simple polygon P. Let the angle of $v_k v_{k+1}$ be μ, so that the angle of the edge $v_{k+1} v_{k+2}$ of P is $\mu - \tau$. The orientation of the edges of P must cycle counterclockwise past μ to close with the segment $x v_{k+1}$ at angle $-\mu$. Because all the angles are ±sums of the same τ, this angle must pass through μ exactly for some edge $v_m v_{m+1}$. For example, the supporting line for P parallel to M passes through such an edge.

Finally, we may apply the Reflect-Translate lemma (Lem. 10) to show that C is not in maxspan configuration. □

Note that the argument to conclude there is a parallel edge $v_m v_{m+1}$ fails for spirals, because the angle turns never need cancel out.

Although we know this theorem does not hold in general for self-crossing spirals, we know from Lemma 7 that it does in the special case of α=90°. We suspect there are other natural classes of chains, which we collectively call the *trans-family* of chains, for which the max flat span is always achieved by the trans-configuration, a point revisited in Sec. 7.

Lemma 7 and Theorem 8 permit, in these two cases, computation of the max flat span of an α-chain in $O(n)$ time, as in Table 1, in contrast to Soss's NP-hardness result for arbitrary angles. We will use this complexity result as part of the dynamic programming algorithm in Sec. 6.

5 3D Structure Theorems

Our results in the previous sections were all in 2D. Here we turn to 3D, and establish several “structure theorems” that characterize the structure of maxspan configurations in a variety of circumstances. The theorems all have the same form: a 3D maxspan is composed of “aligned” planar spans. Despite this unity, our proofs are neither simple nor as general as might be possible.

These structural results were suggested by an implementation of a gradient ascent approximation algorithm.[5] Typical output is shown in Fig. 10.

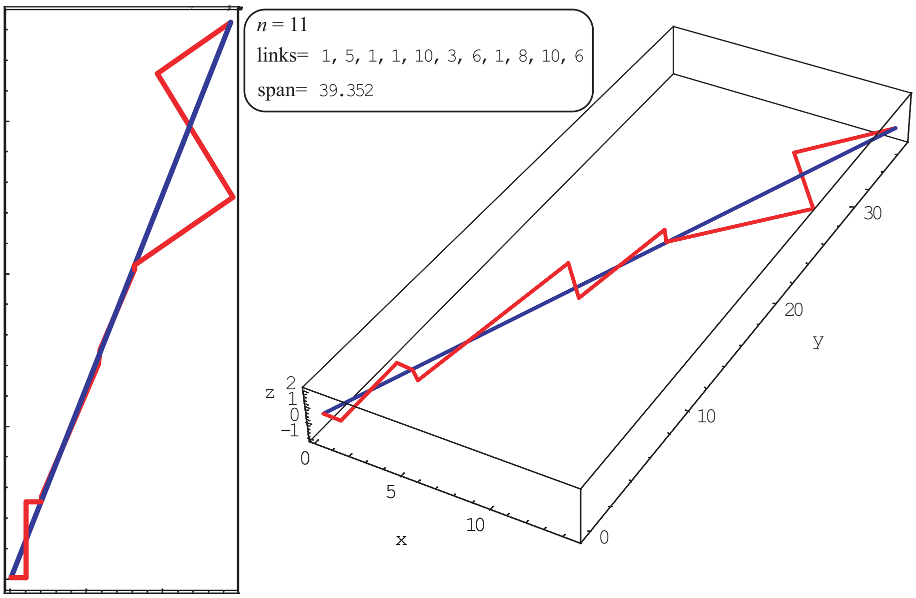

Figure 10: Two views of a 90°-chain of 11 links. 3-, 5-, and 3-link planar subchains align along the central line.

The structure theorems for n-chains (Sec. 5.3) relies on structure theorems for 4- and 5-chains described in the next two sections.

5.1 4-Chain Structure Theorem

Theorem 12 (4-Chain Structure) *The maxspan of a 4-chain is achieved in one of two configurations:*

1. *Alignment of the spans of the two 2-chains (v_0, v_1, v_2) and (v_2, v_3, v_4); or*
2. *The entire configuration is planar.*

---

[5] Our implementation is similar to that suggested by Soss [Sos01, p. 115].

We will illustrate our reasoning with the example of a 90°-chain whose link lengths are $(2, \frac{1}{4}, 1, 1)$. We place v_2 at the origin, v_1 at $(-\frac{1}{4}, 0)$, and v_0 at $(-\frac{1}{4}, -2)$ on the xy-plane. See Figs. 11 and 12.

Proof: We can fix the first two links and $\{v_0, v_1, v_2\}$ in the xy-plane without loss of generality. We seek the configuration that achieves $\max|v_0v_4|$. From the 4-vertex lemma (Lem. 2), we know that v_4 must also lie in the xy-plane. Now we examine the possible motions of the 2-link chain (v_2, v_3, v_4) anchored at the origin v_2. In our example, v_3 rotates on a circle centered on v_1v_2 of radius 1. Around each v_3 position is another circle of radius 1 where v_4 may lie. These sweep out a type of torus (see Fig. 11) with axis through v_1v_2; call it the v_4-*torus*. For arbitrary α, the situation is qualitatively the same (although with less symmetry); and depending on the link lengths, the hole of the torus may close up. We seek the intersection of the v_4-torus with the xy-plane.

This intersection is simply two arcs $A = ab$ and $A' = a'b'$ of a circle of radius $|v_2v_4|$ centered at v_2, symmetrically placed on opposite sides of v_1v_2. See Fig. 12. The endpoints of these arcs correspond to planar configurations of the 4-chain (i.e., when v_3 also lies in the xy-plane). In general, only one arc can possibly contain the maximum, in our example, A.

Now, it is clear that if v_4 is on the relative interior of an arc, then v_0v_4 must be orthogonal to that arc; otherwise we could increase $|v_0v_4|$ by moving towards orthogonality. Hence v_0v_4 passes through the center v_2 of the circle containing the arcs, and we have alignment of the spans of the two 2-chains (v_0, v_1, v_2) and (v_2, v_3, v_4). This is the first option of the lemma claim.

If v_4 coincides with one of the endpoints a or b, then v_3 lies in the xy-plane and hence the maximal configuration is planar, the second option of the lemma claim. $\square$

Note that an implication of this lemma is that the maxspan configuration of a 4-chain is never achieved by a planar 3-chain attached to one link not in that plane, which would, in any case, violate the 4-vertex lemma (Lem. 2).

5.2 5-Chain Structure Lemma

Although we believe the analog of Theorem 12 holds for 5-chains with (in general) different α_i at each joint, we only establish a more narrow result for 5-chains whose two central angles are equal. In some sense this "5-Chain cis" lemma is the heart of the 3D proofs, which are ultimately reduced to it.

Let K_i be the cone on whose rim v_i must lie in order for the angle at v_{i-1} to be its given fixed value, $\angle(v_{i-2}, v_{i-1}, v_i) = \alpha_{i-1}$. The axis of K_i is the line through the chain link (v_{i-2}, v_{i-1}). In general, K_i moves in space as this link moves.

Lemma 13 (5-Chain cis) *Let* $C = (v_0, v_1, v_2, v_3, v_4, v_5)$ *be a 5-chain,*

1. *with its two central angles equal,* $\alpha_2 = \alpha_3 = \alpha$ *(the two extreme angles,* α_1, α_4*, are arbitrary);*
2. *with the first two links* $C_1 = (v_0, v_1, v_2)$ *lying in plane* Π_1 *and the last three links* $C_2 = (v_2, v_3, v_4, v_5)$ *lying in plane* Π_2*, with* $\Pi_1 \neq \Pi_2$*;*

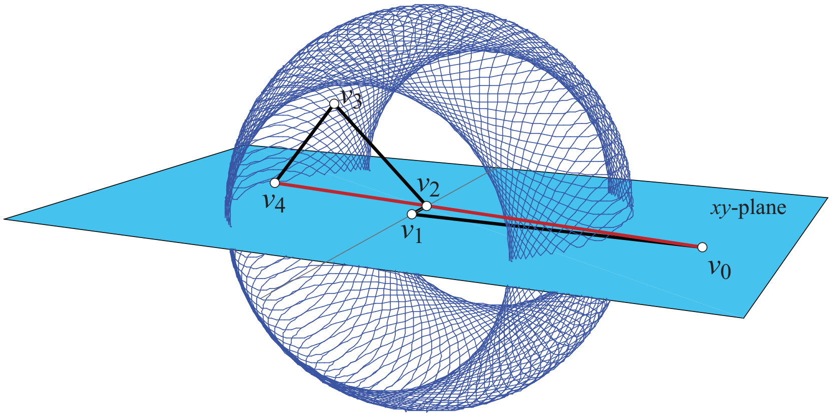

Figure 11: The v_4-torus is centered on an axis through $v_1 v_2$. The chain is shown in its maxspan configuration.

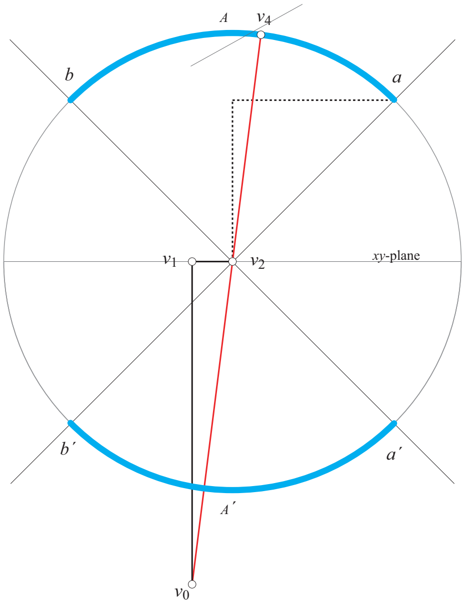

Figure 12: The locus of v_4 positions in the xy-plane, corresponding to an overhead view of Fig. 11. The maxspan shown aligns two 2-chains.

3. with C_1 and C_2 aligned: v_5 lies on the line L through $\{v_0, v_2\}$.

Then, if C is in maxspan configuration, C_2 cannot be in cis-configuration.

Proof: L lies in Π_1, and by assumption, v_5 is on L and so on Π_1. We have $\Pi_1 \cap \Pi_2 = L$. We can fix C_1 and just consider the joints of C_2 free to move. We aim to show that v_5 can move to project further out on L, thereby establishing a contradiction to the assumption that C is in maxspan configuration. Although v_3 is free to move on K_3, we freeze this degree of freedom to simplify the argument, and only permit v_4 and v_5 to move. Fig. 13 illustrates the situation and establishes notation; only K_4 is shown.

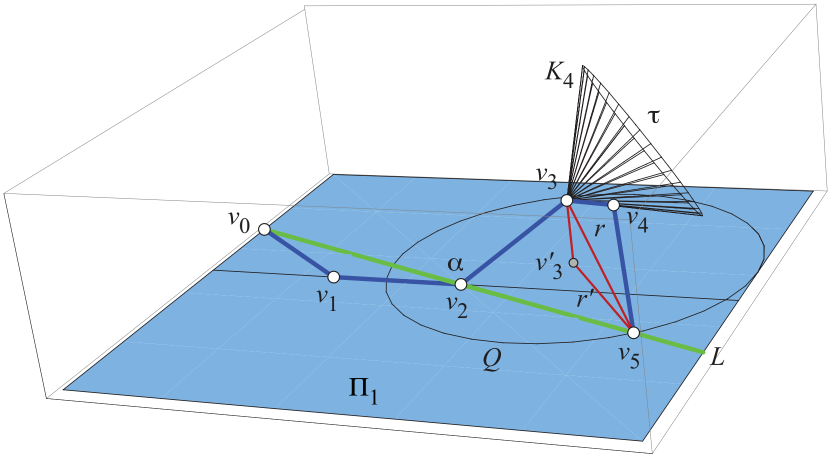

Figure 13: v_5 must lie on the circle Q of radius r' centered on v'_3. K_4 has half-angle $\tau = \pi - \alpha$.

Now, with α_4 fixed, the two link chain (v_3, v_4, v_5) has a fixed span $r = |v_3 v_5|$. Let v'_3 be the projection of v_3 onto Π_1, and let r' be the projection of r to this plane. Then, because r is a fixed distance from v_3, the vertex v_5 must lie on the circle Q of radius r' centered on v'_3. (One can view Q as the intersection of a sphere of radius r centered on v_3 with Π_1.) In general v_5 cannot be located anywhere on this circle, but only on a subarc $q \subset Q$ of it. (There may be two subarcs but only one, which we call q, is relevant for maxspan.) We now determine this arc q.

For each position of v_4 on the rim of K_4, $v_4 v_5$ lies on the surface of K_5 and v_5 lies on its rim. Thus the positions for v_5 given a fixed v_4 lie at the intersection of the circle that is the rim of K_5 with Π_1. This circle intersects Π_2 in 0, 1, or 2 points; or it may lie entirely in Π_1. (In this last degenerate case, $q = Q$.) Fig. 14 illustrates the resulting set of v_5 positions. A generic position of v_4 results two solutions, shown in (b). There is a one-point intersection when the K_5 rim is tangent to Π_1, which occurs when $\triangle v_3 v_4 v_5$ lies in a vertical plane perpendicular to Π_1. This occurs at two symmetric positions of v_4 on K_4, as shown in (a) and (c) of the figure, symmetric about the line through v_2v_3', which line of symmetry lies in the same vertical plane as the lowest rib of K_4 (see ahead to Fig. 15 for an overhead view.) As v_4 moves around the remainder of the K_4 rim, K_5 rises above the Π_1 plane and does not intersect it. Therefore, in general, the set of v_5 solutions constitutes an arc q of Q whose endpoints are determined by the orthogonality of $\triangle v_3v_4v_5$.

Two remarks are in order. First note that, because we know $v_5 \in \Pi_1$, we know that q is not empty. However, it could degenerate to a single point, a case to which we will return below. Second, it is possible for the set of v_5 solutions to constitute two arcs, which occurs, for example, when $\alpha = 90°$ and $v_3' = v_2$, and the K_4 cone is vertical. However, only one arc is relevant, the one whose intersection with L yields a longer span, and it is this one that we label q.

Now we examine how L intersects $q \subset Q$.

1. $p = L \cap q$ is a point in the relative interior of q.
   (a) L is not orthogonal to Q at p. This is the generic case. Moving v_5 to one side or the other on q in a neighborhood of p lengthens its projection onto L, contradicting the assumption that C is in maxspan configuration.
   (b) L is orthogonal to Q at p. Then L must pass through v_3', the center of Q. Because Π_2 includes L (recall that $\Pi_1 \cap \Pi_2 = L$), and $C_2 \subset \Pi_2$ by hypothesis, we know that Π_2 includes both v_3' and v_3, and so includes the vertical segment v_3v_3'. Therefore, Π_2 must be orthogonal to Π_1; see Fig. 15. Define θ to be the angle between L and v_2v_3. As illustrated in Fig. 16, when Π_2 is orthogonal to Π_1, we must have $\theta \le \tau$, because v_3 rides on the rim of cone K_3, whose half-angle is $\tau = \pi - \alpha$. In fact, we can claim strict inequality, $\theta < \tau$, for the following reason. Suppose $\theta = \tau$. This can only occur when v_3 is at the highest point of K_3, in which case L aligns with v_1v_2. But we also know that $v_0 \in L$, so the first two links of C are collinear. In this case, our 5-chain reduces to a 4-chain and the 4-chain structure theorem (Thm. 12) applies to establish the claim of this lemma. So we henceforth assume that $\theta < \tau$.

      Now we use the assumptions that $\alpha_2 = \alpha_3 = \alpha$ and that C_2 is in cis-configuration. As Fig. 17 shows, it must be that v_3v_4 slants downward in Π_2 toward L. Define the reflection vector R to connect v_5 to its reflection v_5^r across the line containing v_3v_4. Then it must be that $R \cdot L > 0$, and v_5^r projects beyond v_5 onto L. Therefore, this reflection, which changes C_2 from cis to trans, increases the span of C, contradicting the assumption that it is in maxspan configuration.

2. $p = L \cap q$ is an endpoint of q. Recall that $\triangle v_3v_4v_5$ lies in a vertical plane (orthogonal to Π_1) at the endpoints of q (Fig. 14(a,c)). Because C_2 is planar and lies in Π_2, this means that again we must have L through v_2

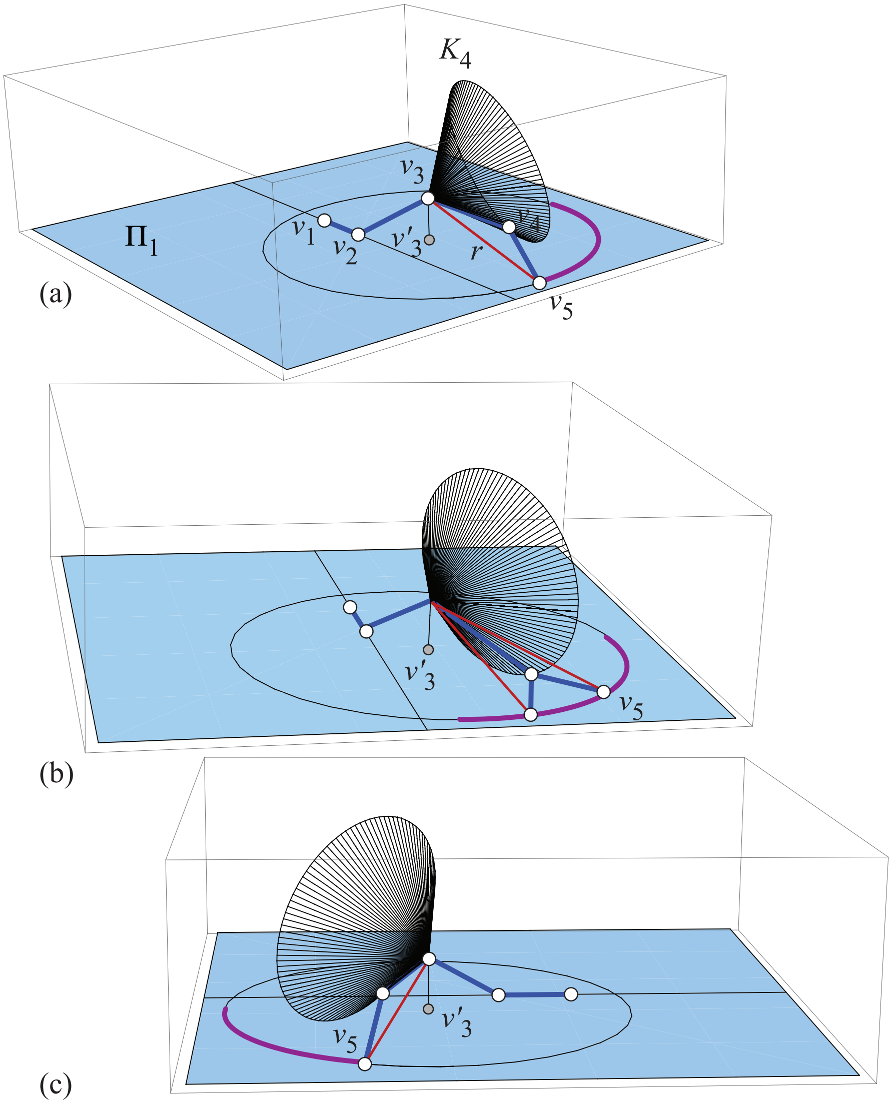

Figure 14: Arc endpoints (a) and (c) are determined when $\triangle v_3v_4v_5$ is orthogonal to Π_1. In between, there are two solutions (b).

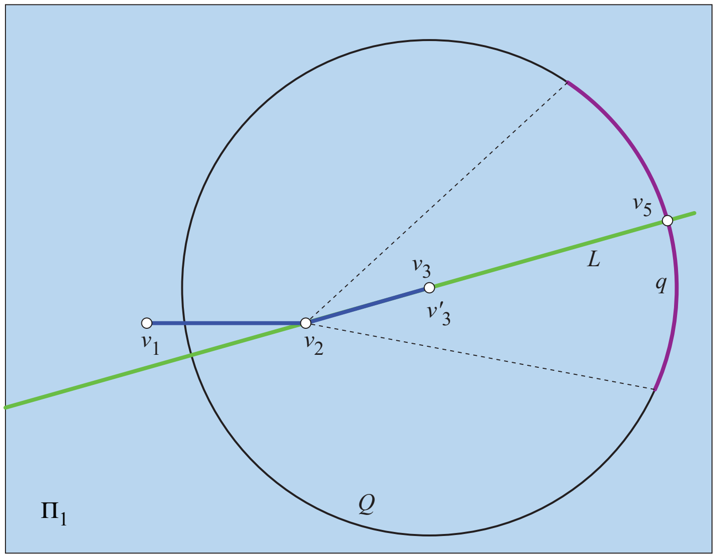

Figure 15: Overhead view when Π_2 is orthogonal to Π_1.

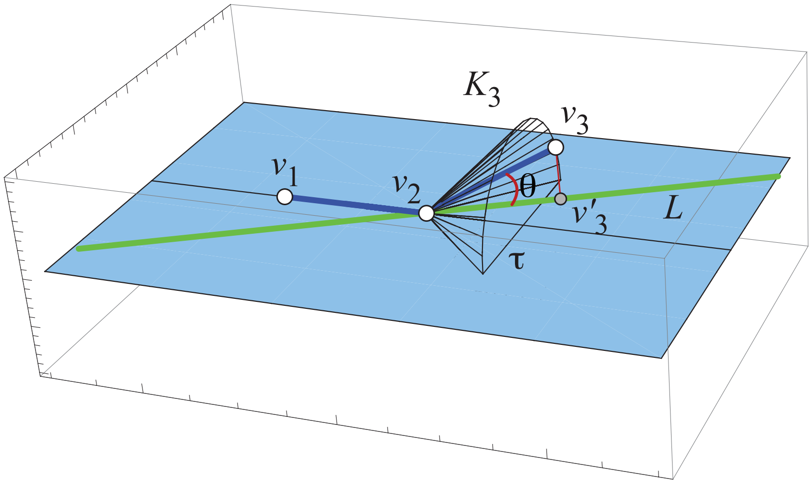

Figure 16: $\theta \leq \tau$ when L passes through v'_3.

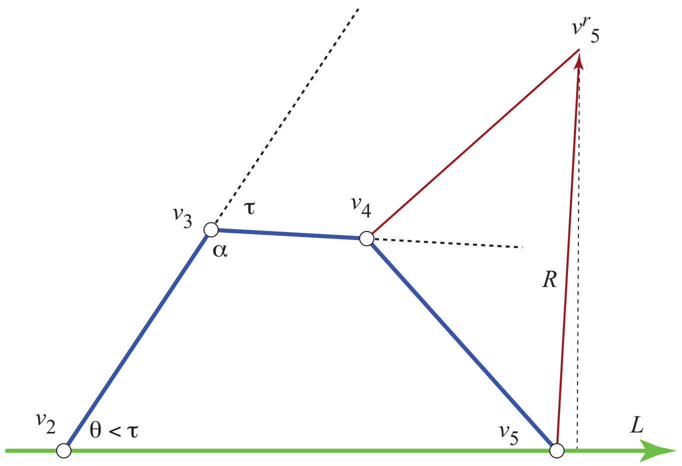

Figure 17: $\theta \le \tau$ implies that $R \cdot L > 0$.

> and v_3', just as in Case 1(b) above. But now v_5 sits at an endpoint of q. Recalling that these endpoints are symmetric about the line through $v_2 v_3'$ shows that in fact in this case, q reduces to a single point. (Note that this is no contradiction to e.g., Fig. 14(c), because C_2 is not planar at this q endpoint.) Now we can apply the exact same reflection argument as above to conclude that $R \cdot L > 0$ and C could not have been in maxspan configuration.

□

We will have occasion to use the above lemma in a slightly more general context:

Corollary 14 (n-Chain cis) *The 5-Chain cis lemma above (Lem. 13) holds for an arbitrary chain replacing* $C_1 = (v_0, v_1, v_2)$*.*

Proof: C_1 remains fixed throughout the argument. We only need that $\alpha_2 = \alpha_3 = \alpha$, so any C_1 that meets C_2 at the same angle would serve as well. □

As mentioned, we believe the assumption that $\alpha_2 = \alpha_3 = \alpha$ in this lemma is not needed. However, it is this assumption that permits the reflection argument to work, and that permits fixing v_3 throughout the argument. For $\alpha_2 \neq \alpha_3$, it seems necessary to argue that moving v_3 lengthens C, and that introduces another level of complexity in an already long proof. Because it suffices for our purposes to assume the two central angles are equal, we have opted for this weaker lemma.

5.3 n-Chain Structure Theorems

We now turn to n-chains, and capture what was empirically observed in Fig. 10 in Theorems 17 and 18 below.

5.3.1 Planar Partition and Alignment

For the n-chain structure theorems, we will partition a chain $(v_0, \ldots, v_n)$ into planar sections by executing the following procedure. Group $\{v_0, \ldots, v_i\}$ into one section if they lie in plane Π_1, but v_{i+1} does not lie in this plane. Then group $\{v_{i+1}, \ldots, v_j\}$ into a second section if they lie in plane $\Pi_2 \neq \Pi_1$, and v_{j+1} does not lie in Π_2. And so on. See Fig. 19(a). Each section, except perhaps the last, contains at least two links (because three vertices determine a plane). Although the partition could be different if the indices are reversed, this ambiguity will not be relevant. It is, however, important that the last section not contain one link, established as part of Theorem 17 below.

We will need two simple technical lemmas in the sequel.

Lemma 15 (Nested Cones) *Let $C = (\ldots, v_{i-1}, v_i, v_{i+1}, v_{i+2}, \ldots)$ be a n-chain with $n \geq 4$, and $C' = (\ldots, v_{i-1}, v_i, v_{i+2}, \ldots)$ the same chain but with v_{i+1} shortcut by $v_i v_{i+2}$. Then, for any fixed-angle configuration of C', there is a fixed-angle configuration of C that matches at the corresponding vertices.*

Proof: Let $\alpha_i = \angle v_{i-1} v_i v_{i+1} = \alpha$ and $\angle v_{i-1} v_i v_{i+2} = \alpha'$, and let link $v_i v_{i+1}$ lie on cone K_{i+1} and $v_i v_{i+2}$ on cone K'_{i+2}; see Fig. 18. Note these two cones share a common axis through $v_{i-1} v_i$, and so are "nested." For any configuration of C',

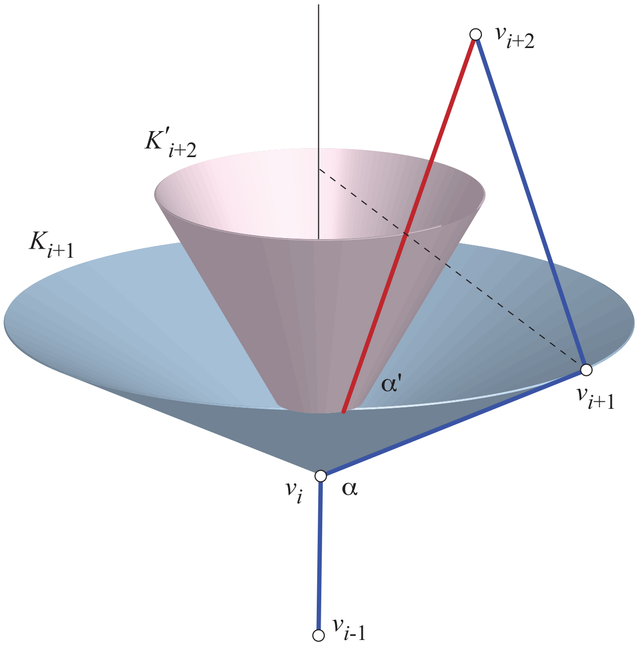

Figure 18: $\triangle v_i v_{i+1} v_{i+2}$ lies in a plane through the common axis of the cones.

there is a placement of v_{i+1} on the rim of cone K_{i+1} so that $\angle v_{i-1} v_i v_{i+1} = \alpha$, determined by the intersection of the plane containing (v_{i-1}, v_i, v_{i+2}) with K_{i+1}, as illustrated in the figure. In this position, $\triangle v_i v_{i+1} v_{i+2}$ is orthogonal to K'_{i+2}. The remainder of C matches C'. $\square$

Lemma 16 (Coplanar) *Let L be a line. If both the three points $\{a, b, c\}$, and the three points $\{b, c, d\}$ are coplanar with L, then, either both b and c are on L, or the four points $\{a, b, c, d\}$ are coplanar with L.*

Proof: Let Π_1 be the plane containing $\{a, b, c\}$ and let Π_2 be the plane containing $\{b, c, d\}$. If $\Pi_1 = \Pi_2$, the claim is established. If $\Pi_1 \neq \Pi_2$, then because two distinct planes meet in one line, $\Pi_1 \cap \Pi_2 = L$. But also we must have that $\Pi_1 \cap \Pi_2 \supset \{b, c\}$. Therefore b and c lie on L, the alternative claim of the lemma. $\square$

We are now ready to prove the alignment claim for maxspan n-chains.

Theorem 17 (n-Chain Partition) *The above-defined planar partition for an n-chain C (with arbitrary α_i) in maxspan configuration has the following two properties:*

1. *The vertices shared between adjacent planar sections all lie along the line L through $v_0 v_n$.*
2. *The last planar section cannot contain just one link $v_{n-1} v_n$.*

Proof: Let

$$C = (v_0, \ldots, v_{k-1}, v_k, v_{k+1}, \ldots, v_{n-1}, v_n) \;.$$

The proof reduces C to various 4-chains C':

$$C' = (v_0, v_{k-1}, v_k, v_{k+1}, v_n)$$

We will then apply the 4-chain Structure Theorem (Thm. 12) to C', obtaining that lemma's conclusion, which we will abbreviate Lem4C(k).

Now we justify why that lemma is applicable, for any $k = 2, \ldots, n-2$. First, this range of k ensures that C' will indeed be a 4-chain; see Fig. 19(b). That the lemma is applicable follows from applying the Nested Cones lemma (Lem. 15) twice, once to each end of C'. In one direction, v_k, v_{k+1}, v_n here play the roles of v_{i-1}, v_i, v_{i+1} in Lemma 15. The "shortcut" $v_{k+1} v_n$ in C' substitutes for the rigid chain $(v_{k+1}, \ldots, v_{n-1}, v_n)$ in C. In the other direction, v_k, v_{k-1}, v_0 here play the roles of v_{i-1}, v_i, v_{i+1} in Lemma 15. So a C' configuration yields a C configuration. Now, because C is in maxspan configuration, C' must be as well, for if it were not, the span $|v_0 v_n|$ of C', and therefore of C, could be increased.

Each application of the lemma yields Lem4C(k) $= P_k \vee A_k$ where P_k means that $\{v_0, v_{k-1}, v_k, v_{k+1}, v_n\}$ are coplanar, and A_k means that the 4-chain C' aligns its two sub-2-chains, and therefore $v_k \in L$. It will be more convenient to interpret P_k as the claim that $\{v_{k-1}, v_k, v_{k+1}\}$ is coplanar with $L \supset v_0 v_n$, which is clearly equivalent. This viewpoint separates out what is common to each application (v_0 and v_n) and what varies (the three central vertices of C').

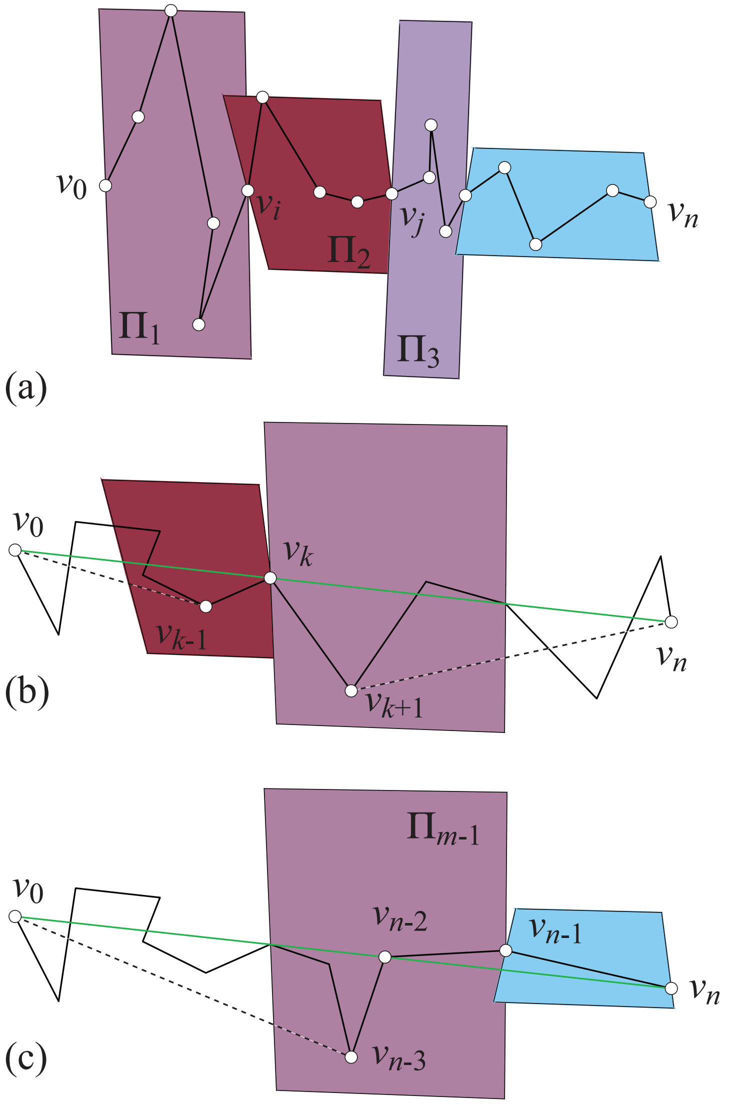

Figure 19: (a) Partition of a chain into planar sections; (b) $\{v_0, v_k, v_n\}$ are collinear; (c) The last section cannot contain just one link $v_{n-1}v_n$.

Let there be m planar sections. Let the first three planes for these sections (assuming there are that many) be Π_1, Π_2, Π_3, with v_i the vertex at the join of the first two sections, and v_j the vertex at the joint of the second two sections. See Fig. 19(a). We partition the argument into three parts: Beginning, Middle, and End.

Beginning. We start at v_i. Lem4C(i) $= P_i \vee A_i$. If P_i, then $\{v_{i-1}, v_i, v_{i+1}\}$ is coplanar with $L \supset v_0 v_n$. Because $\Pi_1 \supset \{v_0, v_{i-1}, v_i\}$, we have two possibilities to consider. Either $\{v_0, v_{i-1}, v_i\}$ fully span Π_1, in which case P_i would imply that $v_{i+1} \in \Pi_1$, which contradicts the fact that v_i is the last vertex of the subchain in Π_1. Or $\{v_0, v_{i-1}, v_i\}$ degenerates to a line in Π_1. But then the 4-Vertex lemma (Lem. 2) implies $v_n \in \Pi_1$, so we have alignment as desired (i.e., $v_i \in L$)

The other possibility is that A_i holds, and $v_i \in L$. So now we know that $L \supset \{v_0, v_i, v_n\}$.

Middle. We next examine Lem4C($i+1$) $= P_{i+1} \vee A_{i+1}$. If A_{i+1}, then $v_{i+1} \in L$. But $L \subset \Pi_1$, which would mean that $v_{i+1} \in \Pi_1$, a contradiction to the assumption that v_i is the transition between Π_1 and Π_2. Therefore, it must be that P_{i+1} holds, and $\{v_i, v_{i+1}, v_{i+2}\}$ are coplanar with L; they lie in the plane Π_2.

Consider now Lem4C(j) $= P_j \vee A_j$. If P_j, then $\{v_{j-1}, v_j, v_{j+1}\}$ is coplanar with L. Lemma 16 then says that, either both v_{j-1} and v_j lie on L, or $\{v_{j-2}, v_{j-1}, v_j, v_{j+1}\}$ are coplanar with L. The latter cannot hold, for that would place $v_{j+1} \in \Pi_2$ when we know that v_{j+1} must lie in $\Pi_3 \neq \Pi_2$. So, if P_j, it must be that both v_{j-1} and v_j are on L, the latter of which is the alignment claim (1) of the lemma. Now we consider the possibility that A_j holds instead of P_j. This immediately implies that $v_j \in L$. So we obtain alignment either way.

Clearly this line of argument can be continued. Studying Lem4C($j+1$) establishes that Π_3 is coplanar with L, and the argument proceeds just as before.

The conclusion is that each planar section is coplanar with L, and that the vertex joins between the planar sections lie on L: claim (1) of the lemma.

End. Assume that the last planar section contains just one link $v_{n-1}v_n$. Let Π_{m-1} be the plane containing the penultimate planar section, containing (at least) $\{v_{n-3}, v_{n-2}, v_{n-1}\}$. From the argument above, we know that Π_{m-1} is coplanar with L. Because v_n lies on L, this says that $v_n \in \Pi_{m-1}$. But this contradicts the assumption that a last planar section was created by the partitioning procedure. Therefore, the last planar section contains at least two links, claim (2) of the lemma. $\square$

5.3.2 Trans-Structure Theorem

The n-Chain Partition theorem (Thm. 17) is our most general structural result; it holds for any fixed-angle chain. The next result we establish only for

90°-chains.[6] Because the α=90° assumption is only used at one point of the argument, and because we believe the theorem may hold more widely, we phrase the proof in terms of α-chains except to highlight when α=90° is employed.

Theorem 18 (Trans-Structure) *If C is a 90°-chain in 3D maxpsan configuration, then each planar section is in trans-configuration.*

Proof: The proof has two main cases, depending on the number m of planar sections in a planar partition of the chain: $m = 2$ and $m > 2$. The case $m = 1$ is settled by Lemma 7.

Case m=2. Assume that $C = (v_0, v_1, \ldots, v_n)$ partitions into two planar sections at vertex v_k, with $C_1 = (v_0, v_1, \ldots, v_k)$ in plane Π_1 and $C_2 = (v_k, v_{k+1}, \ldots, v_n)$ in Π_2, with $\Pi_1 \neq \Pi_2$ and $v_{k+1} \notin \Pi_1$. We know from Theorem 17 that $n > k+1$ and that the subchains align along line $L \supset \{v_0, v_k, v_n\}$. Because $\Pi_2 \supset L$, we have that $v_0 \in \Pi_2$.

Now suppose for contradiction that C_2 is not in trans-configuration. Let $v_r v_{r+1}$ be a cis-edge of C_2, and M the line containing this edge. If $\{v_0, v_n\}$ are strictly to the same side of M, then apply the Reflect lemma (Lem. 9) to increase the span of C by reflection. So assume instead that M places v_0 and v_n on opposite sides of M (or directly on M). We would like to apply the Reflect-Translate lemma (Lem. 10), which requires identifying an edge $v_m v_{m+1}$ parallel to $v_r v_{r+1}$. We cannot use the logic employed in the Simple Flat Maxspan theorem (Thm. 8), because we do not have the equivalent of the Extremity-Crossing lemma (Lem. 11) to exclude spirals. However, because α=90°, every other edge of the chain beyond v_{r+1} is parallel to $v_r v_{r+1}$, even if C_2 spirals. So if $n > r+2$, we are guaranteed such a parallel edge, and can apply Lemma 10 to lengthen C. So assume $n = r+2$. To avoid a parallel edge prior to v_r, we also need $r = k+1$, so that C_2 is a 3-chain, in cis-configuration, the last remaining case. See Fig. 20.

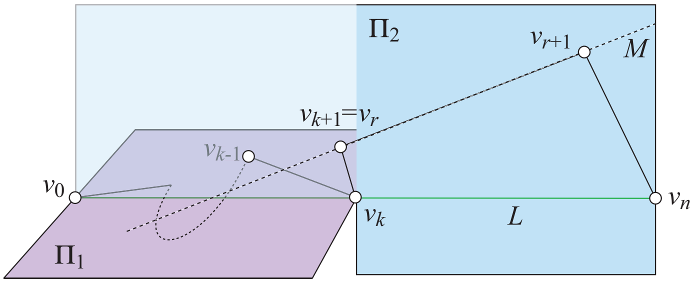

Figure 20: C_2 is a 3-chain with $v_{k+1}v_{k+2}$ a cis-edge, and M separating v_0 and v_n.

---

[6] The extension claimed in [BO06, Thm. 6] is now a conjecture.

Now the n-Chain cis corollary (Cor. 14) applies directly to C and shows it cannot be in maxspan configuration, because C_2 is in cis-configuration. This completes the m=2 case.

Case m>2. The $m > 2$ case parallels that for $m = 2$. Let $\Pi_1, \ldots, \Pi_m$ denote the planes of the m sections, and assume for contradiction that the subchain lying in Π_i, $C_i = (v_i, v_{i+1}, \ldots, v_j)$, is not in the trans-configuration. Replace the prefix chain, $(v_0, \ldots, v_i)$ with an α-chain C_1' that lies in Π_{i-1} and connects v_0 to v_{i-1}. This is easily accomplished with at most $\lceil \pi/\tau \rceil$ links, but because this replacement plays no substantive role in the proof, we do not present details of the replacement.

Now we treat $C' = C_1' \cup C_i$ as in the $m = 2$ case just examined. That proof shows we can increase the span of C'; call the resulting reconfigured chain C''. The plan is to rigidly reattach the suffix chain $C_s = (v_j, \ldots, v_n)$ to C'' to increase the span of the original C. This will succeed in the cases of that proof where reflection and/or rigid translation were used: when the cis-edge line M has $\{v_0, v_j\}$ to the same side, or we identify an edge parallel to the cis-edge. Clearly when the link is merely translated, rigid translation of C_s by the same translation vector maintains the α-angle at v_j, and therefore constitutes a valid reconfiguration of C. When the last link of C_i is reflected across M, we instead reflect C_s across the plane that contains M and is orthogonal to Π_i. Again this maintains the α-angle at v_j. So in these cases, we obtain a reconfiguration of the original chain C, whose span is increased because the span of C'' is longer than that of C'.

This leaves the case where $C_i = (v_i, v_{i+1}, v_{i+2}, v_j)$ is a 3-chain in cis-configuration whose M line separates v_0 from v_j (and therefore from v_n, because $\{v_0, v_j, v_n\} \subset L$). Here we do not see a way to rigidly attach C_s and maintain the α-angle at v_j, because the proof of the 5-Chain cis lemma (Lem. 13), on which this case relies, reconfigures the last link (at least potentially) in an arbitrary manner. Instead we repeat the $m = 2$ argument for this case but based on the 3-chain $(v_i, v_{i+1}, v_{i+2}, v_n)$. Note the angle at v_{i+2} is no longer α, but this 3-chain is still in cis-configuration, because of the slant of M (cf. Fig. 20).

Thus the n-Chain cis corollary (Cor. 14) applies and shows reconfiguration can lengthen the chain. The Nested Cone lemma (Lem. 15) ensures that the shortcut $v_{i+2}v_n$ can be replaced by the original $(v_{i+2}, \ldots, v_n)$, restoring the α-angle at v_{i+2} in the full chain C. Thus we conclude again that C cannot be in maxspan configuration.

This completes the proof of the theorem. $\square$

Although we have only proved this theorem for 90°-chains, we conjecture that an analogous claim holds for chains in the trans-family. And knowing that we are in the trans-family of chains, and so the flat maxspan is achieved in trans-configuration, leads to efficient computation, as we describe in the next section.

6 Dynamic Programming for 90°-Chains

As mentioned in Sec. 1, the complexity of computing the maxspan in 3D is not known. However, for any class of chains for which the Trans-Structure theorem (Thm. 18) holds, the maxspan can be computed in $O(n^2)$ time via a dynamic programming algorithm.

Let $C = (v_0, \ldots, v_n)$ be any α-chain for which the Trans-Structure theorem holds. Initially compute the trans-span of C, recording the coordinates of each vertex for future use. Then if we want to compute the trans-span of a subchain $(v_i, \ldots, v_j)$ we simply look up the coordinates of v_i and v_j and compute their distance in constant time. This is because any subchain of the trans-configuration is itself trans. The subproblems will be $(v_i, \ldots, v_n)$ for $i = n - 2$ down to 1. Hence there are $O(n)$ of these. To compute the maxspan of $(v_i, \ldots, v_n)$, guess the first partition point v_j (i.e., the first planar section) and recurse on $(v_j, \ldots, v_n)$. j will range from $i + 2$ to $n - 2$. For each partition point v_j, determine if the trans-configuration of $(v_i, \ldots, v_j)$ can align with the maxspan configuration of $(v_j, \ldots, v_n)$. We will show that checking alignment is a constant time computation below. If alignment is possible, then the maxspan is

$$\text{transspan}(v_i, \ldots, v_j) + \text{maxspan}(v_j, \ldots, v_n).$$

Store this value, and move onto to computing the maxspan of $(v_{i-1}, v_i, \ldots, v_n)$. If, however, alignment is not possible, then try the next j. If we have tried all possible partition points v_j, and none have lead to alignment, then maxspan of $(v_i, \ldots, v_n)$ is the trans-span, so store this value and move onto computing the maxspan of $(v_{i-1}, v_i, \ldots, v_n)$.

The number of subproblems $(v_i, \ldots, v_n)$ is $O(n)$, and we spend $O(n)$ time per subproblem guessing the partition point and checking whether alignment is possible. Hence the runtime of the algorithm is $O(n^2)$.

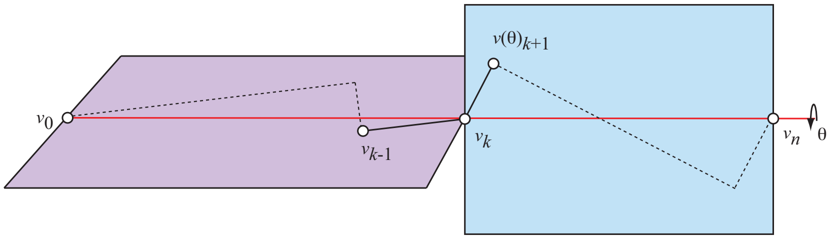

Figure 21: Spin the plane of C_2 about the line through $\{v_0,v_k,v_n\}$, and determine, if for any θ, $v_{k-1}v_k$ makes an angle α with v_kv_{k+1}.

We now show that checking for the possibility of alignment between two subchains $C_1 = (v_0, v_1, \ldots, v_{k-1}, v_k)$ and $C_2 = (v_k, v_{k+1}, \ldots, v_n)$ takes constant time. Attach C_2 to C_1 so that v_0v_k is collinear with v_kv_n. Then spin the plane of C_2 about the line through $\{v_0,v_k,v_n\}$, and determine if some rotation achieves $\angle v_{k-1}v_kv_{k+1} = \alpha$. See Fig. 21. If we parametrize the spin by θ, then this is equivalent to determining whether there exists a θ such that

$$\frac{(v_{k-1} - v_k)}{||v_{k-1} - v_k||} \cdot \frac{(v_{k+1}(\theta) - v_k)}{||v_{k+1}(\theta) - v_k||} = \cos\alpha \ ,$$

a constant-time computation.

7 Open Problems

1. We leave unresolved Soss's question of the complexity of computing the maximum 3D span of an arbitrary chain, line 2 of Table 1. We conjecture it is NP-hard.
2. The gradient ascent approximation algorithm seems not to be distracted by local maxima. Is there a theorem that explains the apparent efficacy of this algorithm?
3. The dynamic programming algorithm runs in polynomial time under two conditions: (a) there is a structure theorem analogous to the Trans-Structure theorem (Thm. 18) that identifies the structure of each planar section, and (b) this structure can be computed in polynomial time. It therefore would be useful to extend our understanding of the trans-family class beyond the two cases we identified in Sec. 4. For example, perhaps all fixed-angle chains whose link lengths fall in the range $[1, 2]$ are in this class? We note that Soss's NP-completeness proof employs links of widely different lengths.
4. Along the same lines, it would be useful to understand when the max flat span is achieved by a self-crossing configuration (Fig. 8).
5. Characterize the class of chains whose maximum 3D span is achieved in a planar configuration, extending the Unit α-chain theorem (Thm. 6).
6. There is every reason to expect that the structure theorems (e.g., Thm. 17) hold in arbitrary dimensions, but we have not pursued this.

Acknowledgment We thank Michael Albertson for suggesting the restricted link-lengths range question, Ruth Haas for the arbitrary dimension question, and Stefan Langerman for an observation that improved the running time of the algorithm from $O(n^3)$ to $O(n^2)$. We benefited from perceptive comments by a referee, whose questions led to the Simple Flat Maxspan theorem (Thm. 8).

References

[Ben06] Nadia Benbernou. Fixed-angle polygonal chains: Locked chains and the maximum span. Undergraduate thesis, Smith College, 2006.

[BO06] Nadia Benbernou and Joseph O'Rourke. On the maximum span of fixed-angle chains. In *Proc. 18th Canad. Conf. Comput. Geom.*, pages 93–96, 2006.

[DLO06] Erik D. Demaine, Stefan Langerman, and Joseph O'Rourke. Geometric restrictions on polygonal protein chain production. *Algorithmica*, 44(2):167–181, February 2006.

[Poo06] Sheung-Hung Poon. On straightening low-diameter unit trees. In P. Healy and N. S. Nikolov, editors, *Proc. 13th Internat. Sympos. Graph Drawing (2005)*, volume 3843 of *Lecture Notes Comput. Sci.*, pages 519–521, 2006.

[Sos01] Michael Soss. *Geometric and computational aspects of molecular reconfiguration.* Ph.D. thesis, School Comput. Sci., McGill Univ., 2001.

[ST00] Michael Soss and Godfried T. Toussaint. Geometric and computational aspects of polymer reconfiguration. *J. Math. Chemistry*, 27(4):303–318, 2000.