7-15-2011

Exact Workspace Boundary by Extremal Reaches

Ciprian S. Borcea
Rider University

Ileana Streinu
Smith College, istreinu@smith.edu

Exact Workspace Boundary by Extremal Reaches*

Ciprian S. Borcea
Department of Mathematics
Rider University
Lawrenceville, NJ 08648, USA
borcea@rider.edu

Ileana Streinu
Department of Computer Science
Smith College
Northampton, MA 01063, USA
istreinu@smith.edu

ABSTRACT

We present the first exact, combinatorial, polynomial time algorithm for computing the description of the workspace boundary for the class of revolute jointed robot arms arising from polygonal orthogonal chains in 3D.

Categories and Subject Descriptors

F.2.2 [**Analysis of algorithms and problem complexity**]: Non-numerical algorithms and problems—*Geometrical problems and computations*

General Terms

Algorithms, Theory

Keywords

robot arm, reach problem, workspace, revolute joint, pseudo-triangulation

1. INTRODUCTION

Let $p = \{p_0, p_1, \cdots, p_{n+2}\}$ be a 3D polygonal chain with $n+1$ internal vertices $p_1, \cdots, p_{n+1}$ and n internal edges. The lengths of the edges $e_i = p_i p_{i+1}$, as well as the angles between two consecutive edges e_i and e_{i+1} are fixed, but in general they are not equal. An *orthogonal chain* has all the fixed angles equal to $\frac{\pi}{2}$, as in Fig. 1(a).

The fixed lengths of two consecutive edges, together with the fixed angle between them, turn triplets of consecutive vertices $p_{i-1}p_ip_{i+1}$ into rigid triangles, called *panels*, as in as in Fig. 1(b). A *panel-and-hinge chain* is a collection of flat rigid bodies ("panels"), serially connected by hinges $A_1, \cdots, A_n$ rigidly attached to them. Each hinge A_i is incident to two consecutive bodies and allows their relative rotation around the axis A_i. In the polygon case, the hinges A_i are the n internal edges p_ip_{i+1}, $i = 1, \cdots, n$, conceived as entire *lines*. More generally, we have *body-and-hinge chains*, with rigid bodies consecutively connected to each other by hinges. The hinges are also referred to as *revolute joints*, or *axes*. The specificity of a panel-and-hinge chain is that any two consecutive hinges lie in the same plane (the plane of the panel), whereas in a body-and-hinge chain, two consecutive hinges are in general skew. A polygonal chain with fixed edge lengths and angles is thus a (generic) panel-and-hinge chain. In robotics, *revolute-jointed robot arm* or *serial manipulator* are the names used for our body-and-hinge chain, the chain bodies are referred to as *links* and the hinges as *joints*. The chain adopts various spatial configurations as the relative angle of rotation around each hinge is varied . *In this paper we assume that there are no rotational limitations around the hinges and no self-collision prohibitions.* These theoretical assumptions are fairly common and methodologically useful. Quoting from [31], Part B, page 231: "we should distinguish between the workspace of the kinematic chain, regardless of the physical implementation of the chain, and that of the physical robot. ... In the early stages of robot design, joint limits need not be considered, the workspace thus exhibiting symmetries that are proper of the type of joints of the regional structure." See also aspects discussed in our related paper [5].

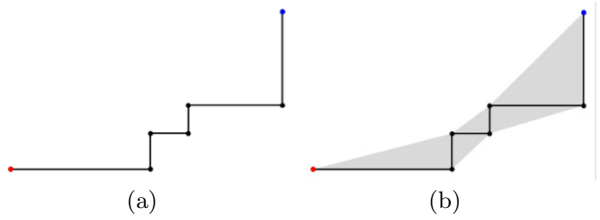

Figure 1: (a) An orthogonal chain with 4 hinges, in a flat zigzag position, and (b) The panel-and-hinge chain associated to it.

The first link of a robot arm (the first body or panel) is a grounded *base*, rigidly attached to the ambient space. The last link carries an *end-effector* or *hand* which is abstracted to a marked point T (the *Terminus*) . The *Workspace* of the chain is defined as the set of points in the ambient space R^3 which can be reached by T . For a polygonal chain, the base link is the triangle $p_0p_1p_2$ and the terminus is $T = p_{n+2}$.

When a *Start* point S is also selected on the first link ($S = p_0$ for polygonal chains), its distance to the terminus T (the *endpoint distance*) takes a continuum of values between two

*Research supported by NSF DMS-0714934, CCF-1016988 and DARPA HR0011-09-1-0003 grants.

SCG'11, June 13–15, 2011, Paris, France.

extremal values called the *Maximum* and *Minimum Reach*, achieved in *Max Reach* and *Min Reach* configurations.

Exact description and computation of the workspace is of fundamental importance for robot manipulator design, placement in the environment and performance evaluation [29] [20] [4] [31] and has received extensive attention. Several approaches and procedures have been used for workspace determination: numerical methods [34] [33] [21], recursive algebraic descriptions [12], analysis of Jacobian singularities [22] [3], tracing of extremal reaches [23] [30] or probabilistic techniques [27]. Many papers addressed the workspace or, equivalently, the workspace boundary determination problem for robots *with very few joints* [17] [24] [35] [14] [26]. In spite of its increased relevance for understanding nanorobots and protein structures [11], the case of an *arbitrarily large number of joints* remained challenging and elusive. The recursive algebraic procedure of [12], although general in principle, has been explicitly used only for instances with a small number of revolute axes [13]. The Jacobian rank-deficiency method becomes impractical as soon as one works with a larger number of joints [2]. Numerical methods are fragile and lack any guarantee of sampling more than a small region of the workspace boundary [1]. We note here the fact that the workspace, as a subset of R^3, may be topologically fairly intricate, due to the possible presence of holes and inner voids.

Overview of our results. In this paper we give the first exact, combinatorial, polynomial time algorithm for computing the description of the workspace boundary of any 3D orthogonal chain. The algorithm builds upon two recent results: a combinatorial criterion characterizing the extremal reaches (maximum and non-zero minimum configurations) of arbitrary panel-and-hinge chains [7], and a linear time algorithm for computing the maximum reach of orthogonal chains [8].

We present: (a) a complete geometric-combinatorial characterization of the workspace boundary for panel-and-hinge chains; (b) a proof that the complexity of the boundary, which may be exponentially large for arbitrary polygons, has polynomial size for orthogonal ones; (c) a linear time algorithm for the minimum reach of orthogonal polygons; and, finally, (d) an exact and efficient algorithm for computing the description of the workspace boundary, in the orthogonal case.

Overview of Techniques. On the **geometric** side, we show that the intersection of the workspace boundary with a plane passing through the first axis is always made of *circular arcs*. An example is shown in Fig. 2. We prove that the entire workspace boundary is obtained by tracing extremal reaches relative to a base-point S which sweeps the first revolute axis $A_1 = p_1p_2$. The boundary thus becomes naturally divided into a *max-boundary* and a *min-boundary*, as traced by the terminus point T in maximum, resp. minimum reach configurations. On the **combinatorial** side, we prove that the articulation of the circular arcs and the positions of the corresponding centers are controlled through a discrete structure based on the notion of *fold points* in extremal reach configurations. Arbitrary polygons can have exponentially many such circular arcs on their workspace boundary, but we show that for orthogonal ones, this number is polynomial. On the **algorithmic side**, our main result is an efficient, polynomial time algorithm for computing *exactly* the workspace boundary of orthogonal polygons. This requires one more step: an algorithm for minimum reaches. Based on our previous linear time algorithm for orthogonal Max Reaches, and the combinatorial characterization of minima, we could easily devise a polynomial time (roughly $O(n^3)$) algorithm for minimum reaches of orthogonal chains. Matching the optimal linear time complexity of the max reach algorithm needs more work. Indeed, the minimum reach appears to be a more challenging problem and the behavior and properties of maxima and minima are not symmetric in full generality (see also the discussion in [5]). Our optimal, linear time algorithm for finding the Minimum Reach of orthogonal chains combines the Maximum Reach approach with two new ideas: a refined structural characterization of minimum reach configurations, and an ingredient from the theory of pseudo-triangulations.

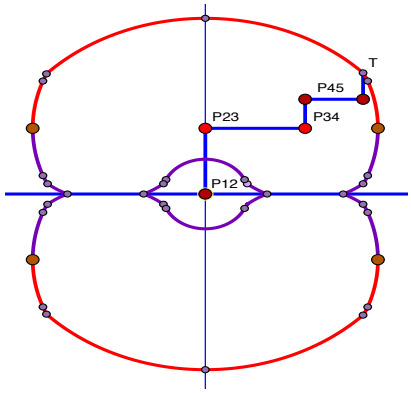

Figure 2: Planar section $\partial\mathcal{W}_2$ through the workspace boundary $\partial\mathcal{W}$ of an orthogonal chain with $n = 5$ hinges. The max-boundary coincides with the convex hull of the workspace boundary, and the min-boundary is the rest. The first hinge is horizontal.

2. PRELIMINARIES

Comment on notation and general assumptions. To reason about panel-and-hinge chains, we may use either the *polygon notation* $p_0, \cdots, p_{n+2}$ (with $T = p_{n+2}$), or the *axes notation* $A_1, \cdots, A_n$, plus the terminus point T. In this latter case, we refer to the point of intersection of axes A_i and A_{i+1} by $p_{i,i+1}$, as in Fig. 2. Technical details regarding parallel axes, multiply concurrent hinges and other *non-generic cases* are not addressed in this extended abstract.

Flat and Zigzag configurations. When all the panels (and thus all hinges, plus S and T) are coplanar, we say that the panel-and-hinge structure is *in a flat configuration* or simply *flat*. If the panels arise as triangles from a revolute-jointed polygonal chain, a special *zigzag* configuration is distinguished, where consecutive triangle do not overlap: the polygonal chain is laid flat in such a way that it alternates making left and right turns at consecutive vertices (as in Fig. 1).

The **endpoint axis** of the chain is the line through S and T. It is divided into two pieces: the finite segment $[ST]$ (the *endpoint segment*) and its *projective complement* $]ST[$, consisting of two infinite rays, "connected" by an projective *point at infinity* (see Fig. 4). The pieces are oriented: $[ST]$ in the usual way, from S to T, and the projective complement

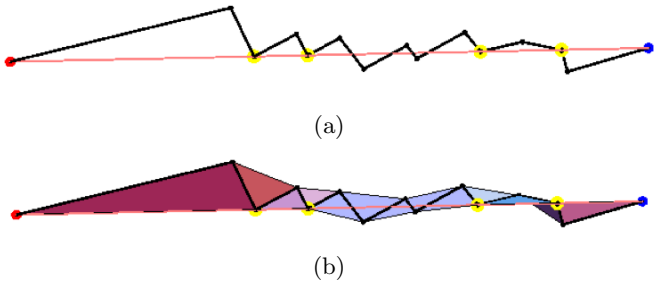

Figure 3: (a) A 3D Max Reach configuration of an orthogonal chain. (b) Shown with panels and fold points. The endpoint axis meets all hinges in natural order.

$]ST[$ is oriented from S, away from T towards infinity the ray incident to S, and then on the other ray from infinity towards T.

The *end-to-end* or **endpoint distance function** assigns a real non-negative value (the distance between the endpoints S and T) to each spatial configuration of the chain. The endpoint distance varies between two extreme values, the *global* minimum and maximum. As a function, it may have local minima, maxima and other critical values, attained in *critical configurations*. It is known that in all critical configurations, the endpoint *axis* meets all the hinge *axes* [23] [30].

Two or more consecutive hinges cut by the endpoint axis away from their intersection point must be coplanar: the panels between them are folded over in a flat configuration. This leads to a **structural decomposition** of a critical configuration into **flat pieces** and **fold points**. The flat pieces arise from contiguous segments of the chain in which several coplanar consecutive hinges are cut (simultaneously, in their common plane) by the endpoint axis (as in Fig. 3). The flat pieces are connected at fold points, which are those vertices of the polygon which *meet the endpoint axis* (see Fig. 3). The two hinges incident at each fold point determine, in addition, a simpler "triangular" **folding panel**, which is met by the endpoint axis only at the fold point.

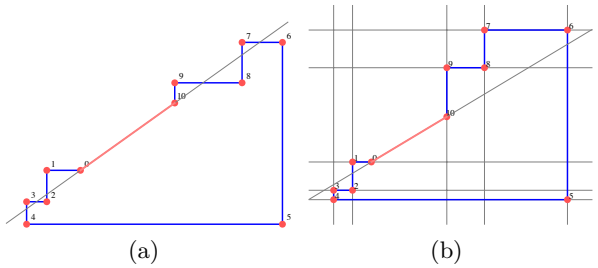

Figure 4: (a) A flat orthogonal chain in its global non-zero minimum position. (b) A flat position is always a critical (not necessarily extremal) configuration. The hinges (in light gray) are all crossed, but not in the natural order.

The sequence $1, 2, 3, \cdots$ of hinge indices, in the order in which they appear along the chain, i.e. the identity permutation, is referred to as the *natural order*. We proved in [7] that a body-and-hinge chain is in a *global maximum configuration* if and only if the *oriented segment* $[ST]$ intersects all hinge axes *in their natural order*. For a non-zero minimum reach configuration we have a similar characterization only in the *panel-and-hinge* case: the oriented projective complement $]ST[$ of $[ST]$ meets the hinge axes in the natural order. For the example in Fig. 1, the flat zigzag chain is *not* in a global max reach position. Fig. 4 illustrates a flat minimum reach configuration next to a critical configuration which is not extremal.

The global maximum has a *dual characterization* [7]: *it coincides with the length of the* shortest path *from S to T which meets all hinges in their natural order.* In [8], we gave an algorithm to compute this constrained shortest path in the case when it meets the hinge axes *inside the polygon segments*. The idea is simple and illustrated in Fig. 5: lay the chain flat in a zigzag position, compute the *paneled polygon*, defined as the union of all the triangular panels of the chain (the gray area in Fig. 1(b)), and then compute the shortest path from S to T *inside* the paneled polygon, as in Fig. 5(a). The turning points of this path, as in Fig. 5(b), correspond to the fold points of the chain in maximum reach position. The class of polygonal chains where this approach yields the global maximum is characterized by a relationship (related to the spherical triangle inequality) between three angles associated to each fold point. For orthogonal chains, these conditions are always satisfied.

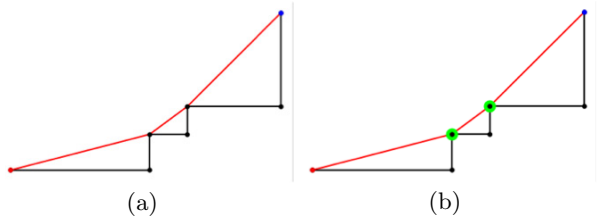

Figure 5: (a) The endpoint segment, (b) The constrained shortest path from S to T, and (c) the fold points. This chain has a 3D Max Reach configuration.

3. THE STRUCTURE OF THE WORKSPACE BOUNDARY

We turn now to our main theoretical result: the general structure theorem for the workspace boundary.

Workspace boundary and singularities of the end-point map. The considerations in this paragraph apply to general revolute-jointed robot arms. A body-and-hinge chain with $n+1$ bodies is given by its n hinges $A_1, ..., A_n$. The first body is fixed, and with it the first hinge A_1. The terminus T is marked on A_n.

The *workspace* is defined to be the locus of the end-point T in R^3, as the chain assumes all possible configurations. It is well known that the abstract configuration space for the chain can be parametrized by n dihedral angles and thus is the n-dimensional torus $(S^1)^n$. The workspace is therefore the image $\mathcal{W}$ of the *end-point map* $e : (S^1)^n \to R^3$, which takes a configuration $\theta = (\theta_1, ..., \theta_n) \in (S^1)^n$ to the corresponding position of the end-point $T(\theta) \in R^3$. There is an obvious symmetry of the possible configurations of the chain under the circle group S^1 given by rotations of the whole part from the second body on, around the fixed first hinge. Thus the workspace $\mathcal{W}$ is in fact determined by any planar section $\mathcal{W}_2$ which contains the first hinge (the sub-

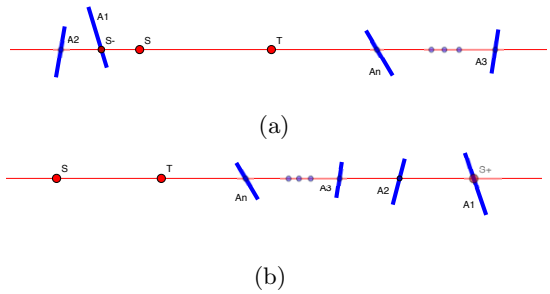

Figure 6: (a) A minimum reach configuration with the same T for both S and S_-. (b) A min reach configuration for S is a maximum reach configuration for S_+.

script serves as a reminder that we are referring to a planar region).

By definition, a *singularity* of the end-point map e is a configuration θ where the differential $De(\theta)$ has rank strictly less than three. The singularity locus is $\Sigma = \{\theta \in (S^1)^n : rank[De(\theta)] \leq 2\} \subset (S^1)^n$. A simple geometric characterization [9] [10] [6] is known for the singular configurations: $\theta \in \Sigma$ if and only if *there is a line through* $T = e(\theta)$ *which intersects projectively all hinges.*

A configuration θ where T reaches a point in the boundary $\partial\mathcal{W}$ of the workspace must be singular, for otherwise the image of e would cover a neighborhood of $T = e(\theta)$ in R^3. Therefore, *the boundary of the workspace is contained in the image of the singularity locus* $e(\Sigma)$: $\partial\mathcal{W}_3 \subset e(\Sigma)$. By considerations of algebraic-geometry and Sard's theorem, it follows that $e(\Sigma)$, the image of the singularity locus, consists of the real points of an algebraic surface. If we intersect this surface with a plane through the first axis we obtain the planar trace of the singularity locus image as an algebraic curve. Thus the planar boundary $\partial\mathcal{W}_2$ is a closed subset of a plane real algebraic curve, and it is a finite union of arcs situated on the planar trace of the singularity locus image. By symmetry considerations, it follows that the workspace boundary $\partial\mathcal{W}$ is a surface of revolution obtained by rotating the planar boundary $\partial\mathcal{W}_2$ around the first hinge A_1.

With these preliminary concepts in place, we turn to an overview of our characterization. From this point on, we work with panel-and-hinge chains. This means that any two consecutive hinges A_i and A_{i+1} are coplanar, that is, projectively incident. The intersection point is denoted by $p_{i,i+1} = A_i \cap A_{i+1}$. The spanned plane is denoted by $\pi_{i,i+1}$ and gives the $(i+1)$ panel of the chain. The first panel is a *fixed* plane containing A_1, which we take as the **reference plane**. The last panel is the plane spanned by A_n and the endpoint T.

The workspace boundary $\partial\mathcal{W}_2$ for panel-and-hinge chains. The main theoretical significance of panel-and-hinge chains, as specializations of general body-and-hinge chains, lies in the fact that the algebraic curve mentioned in the previous paragraph becomes highly *reducible* for this subclass. What we prove in this paper implies that all components of the complex curve are of degree two.

We use our previous characterization of extremal configurations as follows. We choose an arbitrary base-point or *start point* S on the fixed first panel. The configurations achieving the maximum distance between S and T must have T on the workspace boundary, and the oriented segment endpoint segment $[ST]$ intersects all hinges in natural order. We can move S along this segment until it reaches the first hinge A_1, with the same T as maximum reach. This means that *all maximum reach positions of* T *relative to a base-point* S *can be obtained with* S *on the first hinge* A_1. As we vary the position of S on A_1, the end-point T will trace a portion of the workspace boundary which we call the *max-boundary*. Since the sphere centered at S and passing through T obviously contains the whole workspace, the max-boundary is contained in the intersection of the workspace with the workspace convex hull boundary. Later in Theorem 3.4 we show that it actually coincides with this intersection.

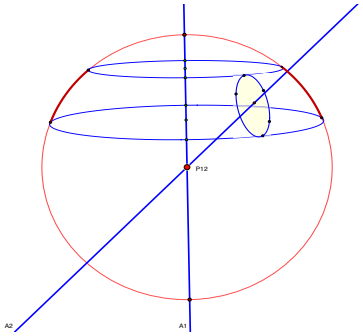

Figure 7: The circular arcs of the planar workspace boundary for two hinges. The reference plane is the plane through A_1 given by the red grand circle.

Non-zero minima allow a similar treatment, but the reason for restricting S to A_1 when tracing our min-boundary has to be explained upfront. For a given position of a base-point S with a non-zero value for the minimum reach, the end-point T must be on the workspace boundary and on the projective complement $]ST[$ of the endpoint segment. There are two possibilities, illustrated schematically in Fig. 6: (a) the intersection with A_1 is on the ray from S to the point at infinity; in this case we can move S until it reaches A_1 at a point denoted by S_-, with the same T as non-zero minimum reach, or (b) the intersection with A_1 (and thereby with all hinges) is on the ray from the point at infinity to T, when repositioning S on A_1 at a point denoted by S_+ produces a maximum reach configuration instead of a minimum one. This means that the corresponding position of T belongs to the max-boundary. *By restricting the choice of the base-point* S *to the first hinge* A_1*, the corresponding non-zero minimum reach positions of* T *trace all possible non-zero minimum reaches which are not already maximum reaches (relative to a repositioned* S*).* The closure of the locus obtained by varying S on A_1 and marking these non-zero minimal reaches of T gives our *min-boundary*.

Definition. The *max-boundary* is the subset of the workspace boundary that can be obtained as the maximum reach of the end-point distance function between a base-point S on the first joint axis and T. The *min-boundary* is the closure of the subset of the workspace boundary that can be obtained as a non-zero minimum reach of the end-point distance function between a base-point S on the first joint axis and T.

The rest of this section is devoted to the proof of our main structural theorem.

THEOREM 3.1. *The workspace boundary $\partial\mathcal{W}_2$ of a panel-and-hinge chain is the union of the max-boundary and the min-boundary.*

We start by emphasizing the role of two structural aspects which intervene repeatedly in our considerations, namely the basic case of two intersecting hinges and the reduction to this case resulting from the presence of fold points in extremal reaches.

Case of two concurrent hinges. The basic case of a chain with three panels and two hinges A_1 and A_2 meeting at p_{12} is illustrated in Figure 7. As the third panel revolves around A_2, the end-point T traces a circle. Since T remains at the same distance from p_{12}, this circle lies on the depicted sphere centered at p_{12}. As the second and third panels revolve around A_1, the circle sweeps a region of the sphere contained between two parallel planes which are normal to A_1. This *spherical region is the entire workspace* of the chain. When intersected with the plane of the first panel, the workspace traces two *circular arcs*, symmetric by reflection in A_1. It is **important** to observe that the end-points of these two circular arcs are reached precisely when all three panels are in the reference plane, that is, in a *flat* configuration.

Fold points and pivoting. A maximum or non-zero minimum reach configuration has the endpoint axis projectively incident with all hinges. By rotational symmetry around A_1, we may assume T in the reference plane. Thus the line ST will be in the reference plane as well, intersecting all hinges or possibly parallel to some of them (i.e. intersecting them at infinity). If the endpoint line ST avoids all points $p_{i,i+1} = A_i \cap A_{i+1}$, then all panels, one after another, must be in the same plane, namely the reference plane, and we have a flat configuration. Thus, non-flat extremal reaches must have one or more points $p_{j,j+1}$ on the end-point line ST. These are the *fold points* of the extremal configuration. (We treat flat limits of non-flat extremal configurations as still possessing their fold points.)

If $p_{j,j+1}$ is such a fold point, the corresponding panel $\pi_{j,j+1}$ may have a different normal direction from its two neighbors $\pi_{j-1,j}$ and $\pi_{j+1,j+2}$, but all panels $\pi_{j+1,j+2}, \cdots, \pi_{k-1,k}$ between one fold point $p_{j,j+1}$ and the next fold point $p_{k,k+1}$, with $j+1 < k$, must be coplanar. All panels up to the first fold point must be in the reference plane and all panels after the last fold point are coplanar. This was illustrated in Fig. 3. One may imagine now all hinges as 'locked', except for the two meeting at $p_{j,j+1}$, that is A_j and A_{j+1}. This creates a short (2-hinge) chain, which was treated above: when T is near the given initial position, it is tracing a circular arc with center at $p_{j,j+1}$ in the reference plane. For some time, the line through $p_{j,j+1}$ and T will intersect the first two axes in the same order as before. This scenario will be called **pivoting** (at the fold point $p_{j,j+1}$) and will be used repeatedly for tracing the max- and min-boundaries.

The max-boundary. It suffices to restrict the analysis to the plane of the first panel (the *reference plane*). The hinge A_1 induces two half-planes, and the reflection in A_1 is as induced symmetry. Thus, it suffices to place the starting point S in one half-plane, and T in the other. We refer to them as the S- and T-halfplanes.

By our maximum reach criterion, the oriented endpoint segment $[S, T]$ intersects all hinges in their natural order. Thus, in a maximum reach configuration, all intersections with hinges are in the T-halfplane. For all other points S' placed in the S-halfplane on the endpoint axis, the maximum reach configuration (with endpoints S' and T) is attained for the same configuration of the chain, and witnessed by the same naturally ordered intersections of the line with the hinges. This means that we obtain the trace of T in its half-plane by varying S *only along the first hinge* A_1. Particular attention is needed when it coincides with the intersection $p_{12} = A_1 \cap A_2$ of the first two axes. There, T has a whole arc of positions on the circle centered at p_{12} and with radius given by the maximum span between p_{12} and T for the panel-and-hinge chain made of all panels from the third on. The following proposition summarizes the situation:

LEMMA 3.2. *The planar workspace region $\mathcal{W}_2$ is contained between two circles centered at $p_{12} = A_1 \cap A_2$ and with radii $m = min||T - p_{12}||$, respectively $M = max||T - p_{12}||$, where the indicated extrema are for the panel-and-hinge chain made of all panels from the third on. The intersection of the workspace boundary $\partial\mathcal{W}_2$ with the larger circle consists of one or two arcs, symmetric with respect to the hinge A_1.*

PROOF. What is left to explain is the last part of the statement. For that, we imagine the panel-and-hinge chain made of panels 3 to $(n + 1)$, in a configuration achieving $M = ||T - p_{12}||$. Then, with this configuration 'frozen', we have to see what T traces in our plane when using rotations only around the first two hinges. This is the elementary case of the well understood workspace of a chain with two hinges. □ We follow now T along the max-boundary as we move S away from p_{12} on A_1. Since the segment $[S, T]$ intersects all hinges in order, all consecutive panels are necessarily in one and the same plane until $[S, T]$ passes through the intersection $p_{i,i+1} = A_i \cap A_{i+1}$ of two consecutive hinges. This is a *fold point*. At a fold point, the $(i + 1)$ panel may have a different direction than its predecessor and its successor, but a new panel direction after that can occur only at a new fold point.

If there is no fold point, all panels are in the plane of the first panel and we can move S along A_1 until we create a maximum reach configuration with a fold point. The flat maximum reach should be recorded as a source of a discontinuity for tangents along the max-boundary since the adjoining circular arcs won't meet tangentially there. Let us examine then what happens for a maximum reach configuration with fold points. Let $p_{f,f+1} = A_f \cap A_{f+1}$ denote the *first* fold point from S to T. Then, we may consider the two-hinge chain with hinges A_f and A_{f+1} consisting of the one panel made of the first f coinciding panels 'frozen', then the $(f + 1)$ panel $\pi_{f,f+1}$ and finally the panel defined by A_{f+1} and T as a flat substitute for the 'frozen' remaining part. The two hinge case shows that T moves on a circular arc until (i) the segment $[S, p_{f,f+1}]$, which keeps intersecting the 'frozen' hinges $A_1, ..., A_f$ in their natural order, passes through some intersection $p_{e,e+1} = A_e \cap A_{e+1},\ e < f$ or (ii) we reach the end of the circular arc. In case (i), we are again in a maximum reach configuration, only with an anterior fold point $p_{e,e+1},\ e < f$ to use as center. In case (ii), our 2R chain is in a flat configuration, which means that either the whole chain is in a flat maximum configuration or we have a posterior fold point $p_{g,g+1},\ f < g$ and we connect with its arc.

Thus T moves from circular arc to circular arc with no discontinuity in the direction of the tangent as long as no

flat maximum occurs. Indeed, at each event, the center of the circular arc is shifted on $[S,T]$ from $p_{f,f+1}$ to $p_{e,e+1}$, with the old and new arc tangent at T. We summarize these arguments as:

THEOREM 3.3. *When S moves away from p_{12} on A_1, the corresponding maximum reach configuration makes T trace circular arcs while receding from the outer circle of center p_{12} and radius M towards the inner circle of radius m. The resulting max-boundary is a differentiable curve with continuous derivative, except for points related to a flat maximum reach configuration. The max-boundary is either connected or has two connected components, symmetric with respect to A_1. At its end-points, the max-boundary curve has tangents perpendicular to A_1.*

The last claim is explained as follows: T approaches an end-point of the max-boundary when S approaches infinity along A_1. At all times, the tangent to the max-boundary curve is perpendicular to the corresponding $[S,T]$ line. As S approaches infinity, this line approaches a parallel to A_1.

THEOREM 3.4. *The max-boundary is the intersection of the workspace boundary with the workspace convex hull boundary.*

PROOF. One inclusion was proven above. For the other, we observe that, as a convex set, the workspace convex hull will always be on one side of any tangent to its boundary. Since the tangents at the end-points of the max-boundary are perpendicular to A_1, the theorem follows from the reflective symmetry with respect to A_1. □

An orthogonal example: the max-boundary tracing. We reconsider here the orthogonal chain with five hinges illustrated in Fig. 2, which gives the full image of the planar workspace boundary $\partial\mathcal{W}_2$. Note that the workspace has an additional symmetry given by reflection in the plane perpendicular to the horizontally placed A_1 at $p_{12} = A_1 \cap A_2$. It is enough to observe what happens in the quadrant shown in Fig. 8.

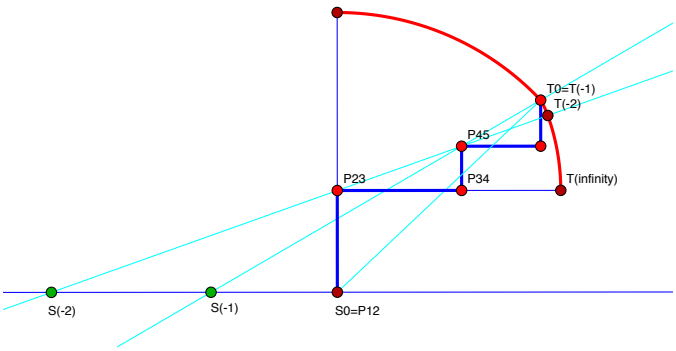

Figure 8: One quadrant of the workspace max-boundary with marked events.

The dots marked on the boundary are positions of the end-point T corresponding to *events* which are described below, as we follow the tracing of the max-boundary. From one dot to the next, we always have a circular arc traced by *pivoting* at a well determined fold point.

With the origin assumed at p_{12}, we shall follow now the tracing of the max-boundary in the first quadrant of the figure. We let the base-point S sweep the 'negative' half-axis of A_1, marking the positions corresponding to events by $S_0 = p_{12}, S_{-1}, ...$ As described in Lemma 3.2, for $S_0 = p_{12}$, we look for the maximum reach M for the distance between p_{12} and T for the chain with joint axes A_3, A_4, A_5. In our case, this maximum is flat and already depicted in Figure 2. Note that the segment from p_{12} to T meets the last three hinges in their sequential order. Thus, $S_0 = p_{12}$ gives in the first quadrant the circular arc of this radius M and center p_{12} going from this initial position to the vertical axis.

When S moves away from the origin, we have flat maximum reach configurations for the full chain in the initial position, until the segment $[ST]$ encounters p_{45}, which for S is the event marked S_{-1}. At this event, we *pivot* at the fold point p_{45} just produced. This means (see Section 3) that the tracing continues with a circular arc with center at p_{45}. This arc stops at $T = T_{-2}$, when the line $p_{45}T$ hits p_{23}. For S, this is the event S_{-2}. Now, we pivot at the fold point p_{23} and the arc with this center continues to T_∞, the junction point with the min-boundary.

Note that at $T = T_{-1}$ the arcs have different tangents because of the flat maximum occurring for $S \in [S_{-1}S_0]$, while at $T = T_{-2}$ the arcs connect with the same tangent.

The min-boundary. The case of the min-boundary is, in most respects, similar to that of the max-boundary. One distinction is that we must work with non-zero minimum reach configurations, that is $S \neq T$, in order to infer that $T \in \partial\mathcal{W}_2$. Configurations with $S = T \in \partial\mathcal{W}_2$ belong to the closure of this non-zero minimum reach locus. The following definition will be used for referring to such closure points:
Definition. An *axial configuration* is a configuration which has the end-point T on the first hinge A_1 and a projective line through T which (for the adequate orientation) intersects all hinges in their natural order.

Indeed, $S = T$ limits of non-zero minimum reach configurations will be precisely of this type.

From now on, we'll assume that we have non-zero minimum reach configurations, with S a base-point on the fixed first hinge A_1. According to the criterion of [7], the open segment (S,T) intersects no hinge, while its complement relative to the projective line through S and T, meets all hinges in order for the orientation from S to T on this projective arc. By reflective symmetry in A_1, tracing the min-boundary amounts to discussing what happens to T in one of the half-planes determined by A_1, as S moves along A_1.

If $m = 0$ in Lemma 3.2, then, in general, $p_{12} = A_1 \cap A_2$ will be in the interior of the workspace and the min-boundary tracing begins only with S in the complement of a closed neighborhood of p_{12}. It will be tacitly assumed that limit cases, such as, in this instance, when p_{12} would actually be on the workspace boundary, can easily be settled after the generic situation is presented and we won't pursue such degenerations.

Let us consider now, with $S \in A_1$ sufficiently distant from p_{12}, that we have a non-zero minimum reach configuration. By symmetry in A_1, the end-point T may be assumed to lie in a fixed (affine) half-plane. Thus, the projective segment from S to T through the point at infinity will go first through the opposite half-plane and only past the point at infinity through the half-plane of T. As in the case of the max-boundary, we look along this oriented segment for the

first fold point of the configuration, say $p_{k,k+1} = A_k \cap A_{k+1}$. (When there's no fold point, we can move S until one is formed. We must register however the discontinuity for tangents caused by the flat minimum reach configuration.)

We envisage now the panel-and-hinge chain with two hinges A_k and A_{k+1}, with the first panel made of the 'frozen solid' first k coplanar panels, then the $(k+1)$ panel $\pi_{k,k+1}$, and as the third and last panel the plane spanned by A_{k+1} and T, as a flat substitute for the 'frozen solid' remaining part of the original chain. T will trace a circular arc of the min-boundary centered at $p_{k,k+1}$ as long as the line through T and p_{k+1} keeps intersecting the first k axes n their natural order. The arc stops for one of the following reasons: (i) a minimum reach configuration with an anterior fold point, say $p_{j,j+1} = A_j \cap A_{j+1},\ j < k$, is obtained for the full chain *or* (ii) the arc reaches its end with a flat configuration for the 2R chain *or* (iii) an axial configuration is produced.

Cases (i) and (ii) are now familiar from the max-boundary version, while (iii) explains the possibility of axial voids. Obviously, the number α of axial configurations must be *even*, since A_1 enters and exits the workspace the same number of times. When $\alpha = 0$ we have an axial hole, while for $\alpha \geq 4$, we have $(a-2)/2$ axial voids.

When we move S towards infinity, we approach the max-boundary since in the limit, with S at infinity, we have a case which can be interpreted both as a maximum and as minimum reach configuration.

We may add the remark that, when the first fold point jumps from one half-space to the other, we'll have an 'inflection point' of the min-boundary there, in the sense that the two circular arcs are joining from opposite sides of the common tangent. This examination gives the min-version of our Theorem 3.3:

THEOREM 3.5. *Any connected component of the min-boundary is a union of circular arcs. When both limit points are on A_1, these two axial configurations, together with the respective pair of symmetric articulations of arcs form a simple loop bounding an axial void. Otherwise, at least one of the end-points connects that component to the max-boundary.*

We note the fact that all voids must be *axial voids*, that is, reflection invariant connected regions which contain a segment of the symmetry axis A_1. Observing that the boundary of any void must contain points of the min-boundary, we conclude the proof of Theorem 3.1.

4. MINIMUM REACH OF ORTHOGONAL CHAINS: A LINEAR TIME ALGORITHM

It is by now clear that minimum reach configurations play an important role in the workspace boundary. Here, we develop their specific theory. For convenience, assume that $n \geq 5$, since otherwise the calculation is straightforward. Unless explicitly stated otherwise, we will refer to a *non-zero minimum* simply as a *minimum* configuration. We start with two general properties, valid for arbitrary panel-and-hinge chains. The first one is a direct and easy consequence of the structure theorem for minimum reach from [7].

COROLLARY 4.1. **(Reversal Point of a Minimum Configuration)** *There exists a unique vertex R of the chain, called the* reversal point*, with the property that the* point at infinity *intervenes on the endpoint axis between the two hinges incident with the reversal point.*

For example, vertex 5 is the reversal point R for the flat minimum configuration from Fig. 4(a). The reversal point splits the chain into two subchains: the SR**-chain**, from S to $R-1$, and the TR**-chain**, which is the chain from $R+1$ to T, taken, for convenience, in reverse. The following Lemma describes what we'll refer to from now on as the *standard min-reach decomposition* of a chain having a non-zero minimum reach configuration.

LEMMA 4.2. **(Structure of a Minimum Configuration)** *A panel-and-hinge chain with endpoints S and T in a non-zero minimum configuration can be uniquely decomposed by two vertices S_m and T_m into three subchains:* **(S-chain)** *S to S_m,* **(flat min-chain)** *S_m to T_m and* **(T-chain)** *T_m to T. The middle one is in a* flat minimum configuration *and contains the reversal point R, while the S- and T-chains (which may be trivial[1]) are in 3D* maximum configurations.

PROOF. By the general structure theorem for critical configurations, the chain is decomposed at fold points into flat pieces. First, we notice that the reversal point cannot be a fold point. The endpoint axis meets the two hinge incident to the reversal point either by cutting across them, or going through one or both of its endpoints. In either case, there is a unique flat piece containing the reversal point. We denote its endpoints by S_m and T_m. Since the endpoint axis meets the S-chain in natural order, it is in maximum configuration. Similarly for the T-chain. □

We switch now to orthogonal chains for additional structure. Assume without loss of generality that a flat minimum reach chain is placed with the reversal vertex R at the origin, the $(R-1, R)$ hinge placed horizontally (along the x-axis) towards $-\infty$ and the $(R, R+1)$ hinge placed vertically, towards $+\infty$, as in Fig. 4(a). A zigzag placement as in Fig. 4(a), with the two SR and TR subchains in zigzag configurations, and the turns at $R-1$, resp. $R+1$, being the only positions where the zigzag (alternating) property is violated, will be referred to as the *min-standard placement*. The min-standard placement of a chain may not necessarily be a min-reach configuration, as illustrated in Fig. 4(b). The next lemma points to the specificities of the flat min-reach configurations. Recall that a *hinge* is a line. For a polygonal chain, we refer to the line segment between two consecutive vertices as a *hinge-segment*. The proof of the following lemma, illustrated by the example in Fig. 4(a), is straightforward.

LEMMA 4.3. **(Orthogonal Flat Minimum Configuration)** *An orthogonal chain in a non-zero minimum flat configuration is always in a min-standard placement. The endpoint axis crosses all the hinges in the interior of the hinge-segments, except for the two axes incident with the reversal point, which are crossed outside the hinge-segment, namely on the side away from the reversal point R.*

The statement above is slightly incomplete due to the existence of special cases, omitted for lack of space. For instance, if the points $R-1$ and S_m coincide, the endpoint axis goes through them.

Polynomial Time Algorithm for Minimum Reach. These lemmas already lead to a polynomial time "naive"

[1]A trivial chain consists in exactly one vertex.

algorithm for computing the minimum reach: (1) "Guess" the two endpoints S_m and T_m of the flat minimum piece and the reversal point R. There are $O(n^3)$ possibilities. (2) Verify that the chain between them, when laid flat in the min-standard position, is indeed a flat minimum, by joining the endpoints and checking the order in which the endpoint axis crosses the hinges. This takes linear time. (3) The maxima of the two remaining subchains can be computed, each in linear time, by our Max Reach algorithm from [8] (see the Appendix). The local folding condition based on the triangle inequality on the sphere, holds at the candidate fold points S_m and T_m. (4) Finally, the maxima are aligned along the S_mT_m axis: if they do not overlap, we have found a non-zero minimum. Otherwise, proceed to the next candidate pair.

The algorithm is guaranteed to produce the non-zero minimum, if it exists, simply because it tries all the possibilities and because the final outcome satisfies Theorem 3.5. Since all the "local" Min Reach configurations, i.e. those satisfying the criterion in Theorem 3.5, attain the global minimum, it follows that at most one "guess" will work, in which case the minimum reach is non-zero. The total complexity of the algorithm is however large, roughly $O(n^4)$. Our next task is to improve this algorithm to linear time. For this, we need a faster "guessing" phase of the flat minimum, eliminating the need to search through all the possible pairs of vertices and potential reversal points. This is accomplished by using the following property, whose proof is not difficult.

LEMMA 4.4. **(Dominant edges for non-zero minima)** *For an orthogonal chain, at least one of the two edges incident to the reversal point has the property that, when laid in the standard flat configuration, its length exceeds the sum of all the other similarly aligned (horizontal or vertical) edges.*

This is illustrated, for example, in Fig. 4 (two edges are dominant in both directions) and Fig. 9 (only one dominant direction). This lemma reduces the complexity of the quest for the minimum flat piece: given an orthogonal chain, we first look for dominant vertical and horizontal edges. If none exist, there is no non-zero minimum. Otherwise, we may have at most one dominant edge in each direction (horizontal and vertical). Each one of their endpoints (2, 3 or 4 of them) is a potential candidate for being the reversal point, but this being a constant number of possibilities, we can try them all. At this point, we have reduced the complexity of the "naive" algorithm to $O(n^3)$, since we have eliminated the need to search for all possible reversal points.

The next task is to find the endpoints of the minimum flat part without searching through $O(n^2)$ possibilities.

Funnels, lids and pseudotriangles. Consider an orthogonal chain in a non-zero min-reach configuration. The endpoint axis joining S to T must pass through all the fold points. In particular, it passes through S_m and T_m, which identify the uniquely defined flat min-reach subchain. Let us mark in blue the part of the endpoint axis that contains the point at infinity, starting from S_m to T_m, and in red the part that goes from S to S_m, resp. T_m to T, i.e. the part that crosses the panels of the S- and T-subchains. In 3D, these segments are aligned. Now, let us lay flat, in standard position, these S- and T-subchains. The red line is now a planar broken line (bending at fold points), lying inside the *panel polygon* induced by the S and T-subchains. The blue line crosses all the hinge segments of the flat min-reach subchain, except for those incident to the reversal point. In particular it crosses the segments $(R-2, R-1)$ and $(R+1, R+2)$, which will be called, from now on, the *lids*.

Let us focus, for the moment, on the entire SR-chain (the chain from S to $R-1$). Lay it flat in standard position and compute the shortest paths from S to $R-1$ and to $R-2$. These two paths may have some common part up to a vertex V_S, then diverge to the two distinct endpoints. Using properties of shortest paths (geodesics) inside polygons and classical concepts from computational geometry (see e.g. [16, 18]) related to shortest paths inside polygons (funnels), it is easy to show that:

LEMMA 4.5. **(Funnel + lid = pseudotriangle)** *The two paths from V_S to $R-2$ and to $R-1$ form a* funnel. *If the lid edge $(R-2, R-1)$ is added, the resulting figure is a pseudotriangle.*

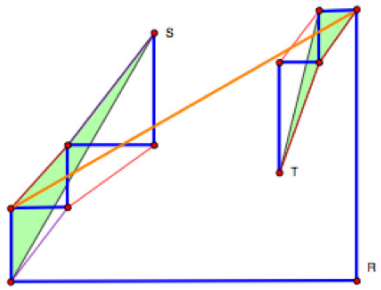

Figure 9: An illustration of the bitangent property from Theorem 4.6.

Applying the lemma to the S- and T-chains, we obtain two pseudo-triangles: the S- and the T-pseudotriangles. They is illustrated in Fig. 9. Pseudotriangles, introduced in [25], have been named so because of their relationship with pseudolines, to which they dualize. A *tangent* is a line tangent to one of the three convex chains defining a pseudotriangle. A *bitangent* is a common tangent to two pseudotriangles. Two generic pseudo-triangles (in the smooth setting) have a unique bitangent; their duals are two pseudolines intersecting exactly once. It is straightforward to extend this property to *generic* (polygonal) pseudotriangles. See also [28] (9.4).

We are now ready to state the *main structural property* for orthogonal minimum reach configurations.

THEOREM 4.6. **(Bitangent endpoints are flat min endpoints)** *The endpoints S_m and T_m of the flat minimum subchain are the endpoints of the unique bitangent to the S- and T-pseudotriangles.*

The proof follows from properties of minimum reach configurations for orthogonal chains: the crossing pattern on the endpoint axis, the fact that the red and blue segments cut through the hinge segments, and the relationship of these properties to simple observations on pseudotriangles.

Linear time algorithm for Minimum Reach. The algorithm takes $O(n)$ steps to find the $O(1)$ candidates for the reversal point R, linear time to compute the S- and T-pseudotriangles (using standard linear time algorithms for shortest paths inside simple triangulated polygons, such as [15, 19]), and linear time to compute the unique bitangent. Correctness follows immediately from the above theory.

5. ALGORITHM FOR WORKSPACE BOUNDARY DETERMINATION

Due to space constraints, we include here only the algorithm for max-boundary. The min-boundary is obtained by combining it with the linear time algorithm for Min Reach from the previous section.

Combinatorics of the boundary. Recall again the constrained shortest path from S to T (illustrated in Fig. 5), computed from the flat zigzag position of the orthogonal chain. Consider the chain in an initial zigzag position and assume that the first axis is horizontal. As S sweeps the axis A_1, the constrained shortest path from S to T undergoes combinatorial changes: fold points appear or disappear at specific event points. As in the previous section, define a *funnel* with apex at T. It is made of the two shortest paths from T to p_2 and p_3 (here we use the polygon notation). For disjoint paths, this is illustrated in Fig. 10(a). In the general case, the funnel has a *tail* (the common part of the two shortest paths that define it) and a *mouth* (the disjoint parts). If we extend the edges of the paths making the mouth of the funnel, and intersect them with the A_1 axis, we obtain the *event points* for S. See Fig. 10(b). The vertices of the funnel correspond to centers of the circles making up the max-boundary, and the total length of the shortest path from such a vertex, to T, will be the radius.

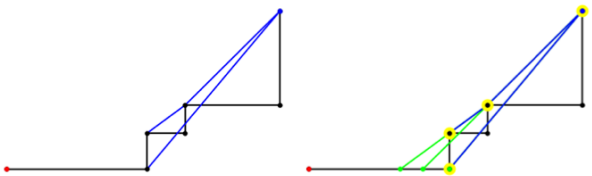

(a) (b)

Figure 10: (a) The funnel, (b) with extended segments, events and fold points. The Max boundary of this chain has 3 circular arcs per quadrant, corresponding to the three events.

This correspondence follows directly from the proof of the Max Boundary structure theorem in Sec. 3.

We are now ready to describe informally the algorithm for computing the max-boundary. Recall from [8] that the panel polygon is the union of all the triangle panels of the chain in zigzag placement, and from the previous section the fact that a funnel with a lid is a pseudo-triangle. Finally, recall that we only intend to describe a section through the plane of the first panel, which is a planar region. The max-boundary will be its intersection with its convex hull. Because of the four-fold symmetry of the orthogonal case (about the first and second axes), only the part of the boundary appearing in the first quadrant is described. The theory developed in Section 3 tells us that the max-boundary is made of circular arcs connecting smoothly. The algorithm provides, implicitly, a combinatorial interpretation for the max-boundary, by putting them in correspondence with segments of two shortest paths (the "funnel"), residing inside the panel polygon associated to the orthogonal chain.

ALGORITHM 1. *(1) Lay the chain in flat zigzag position. (2) Compute the funnel from the terminus vertex to the lid p_2p_3, using two calls to a shortest path algorithm inside the (already triangulated) panel polygon associated to the chain. Split these shortest paths into three parts: C (possibly empty) is their common part (the tail), $U = \{u_0, u_1, \cdots, u_a\}$ (upper path) and $L = \{l_0, l_1, \cdots, l_b\}$ (lower path) are what remains when C is eliminated, with $u_0 = l_0$ being their common vertex. U and L turn consistently only one way (left, for L, and right, for U); together, they form the mouth of the* funnel. *As paths, they are ordered from their common vertex $v = u_0 = l_0$ towards the mouth of the funnel, i.e. $u_a = p_3$ and $l_b = p_2$. See Fig. 10. (3) Compute the intersections with the first axis p_1p_2 of the funnel edge extensions. They are the* events *in the sweep by S of the first axis. (4) Let ℓ_U and ℓ_L be the ordered set of edge lengths in U and L, and c be the total length of the path C. Output a set of circular arcs determined as follows: (L) For the lower funnel path L: for all vertices $u_i, i = 1, \cdots, a-1$ on the path, in order, trace a circular arc by rotating counterclockwise, about vertex u_i, the edge u_iu_{i-1}, extended at u_{i-1} by a segment of length c, until it aligns with the next edge $u_{i+1}u_{i+2}$. The extended edge traces the desired arc. For the last edge on the path, rotate until the edge becomes vertical. (U) For the upper funnel path U: similarly, but rotate clockwise about vertices l_i, and, for the last edge on the path, stop when it becomes horizontal.*

The correctness of the algorithm is a consequence of all the theory developed in the paper. Using a linear time algorithm (e.g. [19]) for shortest paths inside triangulated polygons, its time complexity is linear. Because the arcs of the boundary are in correspondence with the vertices of the funnel, which is linear in size, we obtain:

COROLLARY 5.1. *The complexity of the workspace boundary for an orthogonal chain with n hinges is linear.*

We have examples (for arbitrary chains) where the complexity of the workspace boundary is exponential in size. This fact emphasizes the theoretical importance of the family of chains studied in this paper.

The min-boundary computation is a natural combination of this algorithm and the linear time algorithm for Min Reach. The complete workspace boundary is obtained by union with the min-boundary.

References

[1] K. Abdel-Malek, F. Adkins, H.-J. Yeh, and E. Haug. On the determination of boundaries to manipulator workspaces. *Robotics and Computer-Integrated Manufacturing*, 13 (1): 63–72, March 1997.

[2] K. Abdel-Malek, J. Yang, D. Blackmore, and K. Joy. Swept volumes: foundation, perspectives and applications. *Int. J. Shape Modeling*, 12 (1): 87–127, 2006.

[3] K. Abdel-Malek, J. Yang, and Y. Zhang. On the workspace boundary determination of serial manipulators with non-unilateral constraints. *Robotics and Computer-Integrated Manufacturing*, 24: 60–76, 2008.

[4] J. Angeles. *Fundamentals of Robotic Mechanical Systems: Theory, Methods and Algorithms*. Mechanical Engineering Series. Springer-Verlag New York, 2007.

[5] C. Borcea and I. Streinu. Extremal reaches in polynomial time. In these proceedings, SoCG, 2011.

[6] C. S. Borcea and I. Streinu. Singularities of hinge structures. In *Efficient Methods in Algebraic Geometry (MEGA'05)*, May 2005. arXiv:0812.1373.

[7] C. S. Borcea and I. Streinu. Extremal configurations of manipulators with revolute joints. In *Proc. ASME/IFToMM Conf. on Reconfigurable Mechanisms and Robots (ReMAR'09)*, King's College, London, J. S. Dai, M. Zoppi and X.Kong (eds.), pages 279–284. KC Edizioni, June 2009. arXiv:0812.1375.

[8] C. S. Borcea and I. Streinu. How far can you reach? In *Proc. ACM-SIAM Symp. Discr. Algorithms (SODA10)*, pages 928–937, January 2010. `http://www.siam.org/proceedings/soda/2010/SODA10_075_borceac.pdf`.

[9] J. W. Burdick. A classification of 3R regional manipulator singularities and geometries. *Mechanism and Machine Theory*, 30 (1): 71–89, January 1995.

[10] J. W. Burdick. A recursive method for finding revolute-jointed manipulator singularities. *Trans. ASME, Journal of Mechanical Design*, 117: 55–63, March 1995.

[11] J. F. Canny and D. Parsons. Geometric problems in molecular biology and robotics. In *Proc. 2nd Int. Conf. on Intelligent Systems for Mol. Biology, Stanford*, 1994.

[12] M. Ceccarelli. A formulation for the workspace boundary of general n-revolute manipulators. *Mechanism and Machine Theory*, 31(5):637–646, 1996.

[13] M. Ceccarelli and C. Lanni. A multi-objective optimum design of general 3R manipulators for prescribed workspace limits. *Mechanism and Machine Theory*, 39(2): 119–132, 2004.

[14] M. Ceccarelli and A. Vinciguerra. On the workspace of general 4R manipulators. *The International Journal of Robotics Research*, 14(2):152–160, 1995.

[15] B. Chazelle. A theorem on polygon cuttings with applications. In *Proc. 23rd Ann. Symp. on Foundations of Computer Science (FOCS)*, pages 339–349, 1982.

[16] B. Chazelle, H. Edelsbrunner, M. Grigni, L. J. Guibas, J. Hershberger, M. Sharir, and J. Snoeyink. Ray shooting in polygons using geodesic triangulations. *Algorithmica*, 12:54–68, 1994.

[17] F. Freudenstein and E. J. F. Primrose. On the analysis and synthesis of the workspace of a three-link, turning-pair connected robot arm. *J. Mechanisms Transm. and Automation in Design*, 106 (3): 365–370, 1984.

[18] M. T. Goodrich and R. Tamassia. Dynamic ray shooting and shortest paths via balanced geodesic triangulations. *Journal of algorithms*, 23:51–73, 1997.

[19] L. J. Guibas, J. Hershberger, D. Leven, M. Sharir, and R. E. Tarjan. Linear-time algorithms for visibility and shortest path problems inside triangulated simple polygons. *Algorithmica*, 2(4):209–233, 1987.

[20] K. C. Gupta and B. Roth. Design considerations for manipulator workspace. *Journal of Mechanical Design*, 104(4):704–711, 1982.

[21] E. Haug, C.-M. Luh, F. Adkins, and J.-Y. Wang. Numerical algorithms for mapping boundaries of manipulator workspaces. *Journal of Mechanical Design*, 118 (2):228–234, 1996.

[22] D. Kohli and J. Spanos. Workspace analysis of mechanical manipulators using polynomial discriminants. *Journal of Mechanisms Transmissions and Automation in Design*, 107(2):209–215, 1985.

[23] A. Kumar and K. J. Waldron. The workspaces of a mechanical manipulator. *Journal of Mechanical Design*, 103(3):665–672, 1981.

[24] E. Ottaviano, M. L. Husty, and M. Ceccarelli. Identification of the workspace boundary of a general 3R manipulator. *J.Mech. Design*, 128(1): 236–242, 2006.

[25] M. Pocchiola and G. Vegter. The visibility complex. *Internat. J. Comput. Geom. Appl.*, 6(3): 279–308, 1996.

[26] M. Raghavan and B. Roth. Kinematic analysis of the 6R manipulator of general geometry. *Proc. 5th Int. Symp. Robotic Res.*, H. Miura, S. Arimoto, eds, 1990.

[27] J. Rastegar and B. Fardanesh. Manipulation workspace analysis using the Monte Carlo method. *Mechanism and Machine Theory*, 25(2):233–239, 1990.

[28] G. Rote, F. Santos, and I. Streinu. Pseudo-triangulations: a survey. In J. E. Goodman, J. Pach, and R. Pollack, eds, *Surveys on Discrete and Computational Geometry - Twenty Years later*, *Contemporary Mathematics*, vol. 453, 343–410. AMS, 2008.

[29] B. Roth. Performance evaluation of manipulators from a kinematic viewpoint. *National Bureau of Standards NBS SP*, 495: 39–61, 1976.

[30] R. G. Selfridge. The reachable workarea of a manipulator. *Mech. Machine Theory*, 18 (2): 131–137, 1983.

[31] B. Siciliano and O. Khatib, editors. *Handbook of robotics*. Springer Verlag, 2008.

[32] I. Streinu. Pseudo-triangulations, rigidity and motion planning. *Discrete and Computational Geometry*, 34 587–635, November 2005. DOI `10.1007/s00454-005-1184-0`.

[33] Y. C. Tsai and A. H. Soni. An algorithm for the workspace of a general n-R robot. *J. Mech. Transm. Automation in Design*, 105(1):52–57, 1985.

[34] D. C. H. Yang and T. W. Lee. On the workspace of mechanical manipulators. *J. Mech. Transm. Automation in Design*, 105 (1): 62–69, 1983.

[35] M. Zein, P. Wenger, and D. Chablat. An exhaustive study of the workspace topologies of all 3R orthogonal manipulators with geometric simplifications. *Mechanism and Machine Theory*, 41 (8):971–986, August 2006.