On the Complexity of the Set of Three-Finger Caging Grasps of Convex Polygons

Mostafa Vahedi
Department of Information and Computing Sciences
Utrecht University
Utrecht, The netherlands
Email: vahedi@cs.uu.nl

A. Frank van der Stappen
Department of Information and Computing Sciences
Utrecht University
Utrecht, The netherlands
Email: frankst@cs.uu.nl

***Abstract*— We study three-finger caging grasps of convex polygons. A part is caged with a number of fingers when it is impossible to rigidly move the part to an arbitrary placement far from its initial placement without penetrating any finger. A convex polygon with n vertices and a placement of two fingers —referred to as the base fingers— are given. The caging region is the set of all placements of the third finger that together with the base fingers cage the polygon. We derive a novel formulation of caging in terms of visibility in three-dimensional space. We use this formulation to prove that the worst-case combinatorial complexity of the caging region is close to $O(n^3)$, which is a significant improvement of the previously known upper bound of $O(n^6)$. Moreover we provide an algorithm with a running time close to $O(n^3 \log n)$ that considerably improves the current best known algorithm, which runs in $O(n^6)$ time.**

I. Introduction

The caging problem (or: capturing problem) was posed by Kuperberg [7] as a problem of finding placements for a set of fingers that prevent a polygon from moving arbitrarily far from its given position. In other words, a polygon is caged with a number of fingers when it is impossible to continuously and rigidly move it to infinity without intersecting any finger. A set of placements of fingers is called a grasp.

Caging grasps are related to the notions of form (and force) closure grasps (see e.g. Mason's text book [8]), and immobilizing and equilibrium grasps [13]. A part is immobilized by a number of fixed fingers (forming an immobilizing grasp) when any motion of the part violates the rigidity of the part or the fingers. An equilibrium grasp is a grasp whose grasping fingers can exert wrenches (not all of them zero) through grasping points to balance the object.

Rimon and Blake [12] introduced the notion of the *caging set* (also known as *inescapable configuration space* [16, 20], and recently regularly referred to as the *capture region* [4, 10, 11]) of a hand as all hand configurations which maintain the object caged between the fingers. They proved that in a multi-finger one-parameter gripping system, the hand's configuration at which the cage is broken corresponds to a frictionless equilibrium grasp.

Caging has been applied to a number of problems in manipulation such as grasping and in-hand manipulation [16, 20], mobile robot motion planning [5, 6, 9, 18, 19], and error-tolerant grasps of planar objects [1, 2, 12]. Caging grasps are particularly useful in scenarios where objects just need to be transported (and not subjected to e.g. high-precision machining operations). The fact that the object cannot escape the fingers guarantees that —despite some freedom to move— the object travels along with the fingers as these travel to their destination. The set of caging grasps is significantly larger than the set of immobilizing grasps (as the latter forms a lower-dimensional subset of the former). The additional options for finger placements can be of great value when maneuvering the object amidst obstacles. Moreover, caging grasps are considerably less sensitive to finger misplacements.

In this paper we consider algorithms for caging grasps by robotic systems with two degrees of freedom. The first two papers by Rimon and Blake [12] and Davidson and Blake [1] consider systems with a single degree of freedom. Several other papers [11, 12, 14, 17, 23] study two-finger caging grasps as a special case of robotic systems with one degree of freedom. There are also papers [3, 16, 19, 21] that consider robotic systems with more degrees of freedom. All these papers present approximate algorithms for computing caging grasps. As such they differ from our work, as these algorithms compute a subset of the set of caging grasps. We consider the computation of *all* caging grasps for a given placement of the base fingers.

There are two papers [4, 23] on three-finger caging grasps of polygons that propose algorithms for robotic systems with two degrees of freedom that report the entire solution set. In these papers, a polygon with n edges and a placement of two fingers—referred to as the *base fingers*—are given. It is required to compute the *caging region* for the third finger, which is the two-dimensional set of all placements of the third finger that together with the base fingers cage the polygon. (Consider Figure 1.) Erickson et al. [4] provided the first algorithm for the exact computation of the caging region of convex polygons, running in $O(n^6)$ time. In their paper the base fingers were assumed to be placed along the boundary of the polygon. They also established an upper bound of $O(n^6)$ on the worst-case complexity of the caging region of convex polygons, where the caging region was shown to be the visible scene of $O(n^3)$ constant-complexity surfaces in a three-dimensional space. Vahedi and van der Stappen [23] proposed another algorithm generalizing the previous results to compute the caging region of arbitrary polygons for *any* given placement of the base fingers, that runs in $O(n^6 \log^2 n)$ time. They established the same $O(n^6)$ upper bound on the worst-case complexity of the caging region of non-convex polygons, where the caging region was shown to be a subset of the arrangement of $O(n^3)$ constant-complexity curves defined by equilibrium grasps. (The *arrangement* of a set X of two-dimensional curves is the set of maximally-connected zero-, one-, and two-dimensional subsets induced by the curves of X not intersecting any of the subsets.) However, in both cases the mentioned upper bound on the worst-case complexity of the caging region was due to the proposed algorithms, and it remained an open problem to establish a better upper bound for convex or non-convex polygons. In this paper, we tackle the problem for convex polygons. We prove that the worst-case complexity of the caging region of convex polygons is $O(\lambda_{10}(n^3))$, which significantly improves the already known upper bound of $O(n^6)$, as $\lambda_{10}(n^3)$ is known [15] to be $O(n^3 \log^* n)$[1] and thus very close to $O(n^3)$. To establish the upper bound, firstly we have narrowed down the types of surfaces introduced by Erickson et al. [4] that play a role in the caging region complexity. Secondly, we have formulated a new way to compute the caging region using those surfaces. In addition, we develop an efficient algorithm to compute the caging region in $O(\lambda_s(n^3) \log n)$ time using a divide-and-conquer technique.

In Section II we introduce some definitions and assumptions used in the paper, define a three-dimensional space called *canonical grasp space*, and explain a formulation to compute the caging region in canonical grasp space. In Section III, we use this formulation to prove an upper bound on the complexity of the caging region and present an efficient algorithm to compute it. We conclude the paper with a discussion of future work.

II. Definitions and Assumptions

A *convex polygon* is an intersection of a number of half planes; Throughout the paper P is a bounded convex polygon without parallel edges that has a fixed reference frame and n edges.

We assume that the fingers are points. The distance between the base fingers is d. We assume that vertices and edges of P are in general position, i.e. no two vertices are at distance exactly d from each other, no vertex has (shortest) distance exactly d to an edge, and the angle between the altitude lines drawn from any vertex to any pair of edges is not equal to the supplement of the angle between the corresponding pair of edges (i.e. they do not add up to π).

Instead of considering rigid placements of the base fingers around the fixed polygon P, we can equivalently fix the base fingers at $b_1 = (0,0)$ and $b_2 = (d,0)$ and consider possible placements $q \in \mathbb{R}^2 \times [0, 2\pi)$ of P (with respect to these fingers). Let $P[q]$ denote the set of points covered by P when placed at q. Let $P[\alpha]$ be the polygon P rotated by the angle α around its reference point for which both base fingers are in contact with it and the extensions of the edges touched by the base fingers intersect each other below the x-axis. We refer to $P[\alpha]$ as the *canonical placement* of any other placement q of P with orientation α. A canonical placement can be specified by a single parameter, which is orientation. In Figure 1, $P[\alpha]$ is the canonical placement of $P[q]$. The polygon $P[\alpha]$ may not be defined for all orientations α, which we explain about when we define the base-diameter. In this paper, we consider canonical placements of P because the caging region of any other placement $q = (x, y, \alpha)$ of P is a subset of the caging region of its canonical placement $P[\alpha]$ [22]. Figure 1 shows one example, in which the caging region of $P[q]$ is the area surrounded by bold solid curves and the caging region of $P[\alpha]$ is the union of both the area surrounded by bold solid curves and the area surrounded by bold dashed curves. Moreover, given the caging region of $P[\alpha]$, the caging region of $P[q]$ can be computed easily [22].

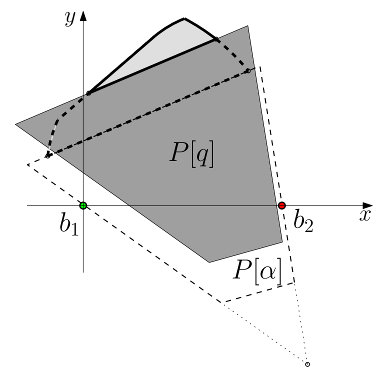

Fig. 1. $P[q]$ and $P[\alpha]$ are rigid translates. Caging region of $P[\alpha]$ is a superset of the caging region of $P[q]$.

Consider P at a placement $q = (x, y, \alpha)$. Every horizontal line intersects the polygon in at most two points. The *base-diameter* of the polygon is the maximum length intersection among all horizontal lines. The *base-diameter* of any $P[q']$ with $q' = (x', y', \alpha)$ equals that of $P[q]$. When the base-diameter of $P[q]$ with $q = (x, y, \alpha)$ is less than d, $P[\alpha]$ is not defined, as it is not possible to place a translated copy of P with orientation α such that it touches both b_1 and b_2. A *critical angle* is an angle α for which d is equal to the base-diameter of $P[\alpha]$.

Every three-finger grasp in which the polygon is at $P[\alpha]$ can be specified by $(p, \alpha) \in \mathbb{R}^2 \times [0, 2\pi)$, to which we refer as a canonical grasp. (The canonical grasp of a grasp is uniquely defined by the definition of canonical placement.) The parameter α specifies the orientation of the polygon and $p = (x, y)$ specifies the location of the third finger. We refer to the space $\mathbb{R}^2 \times [0, 2\pi)$ of all such three-finger grasps as the *canonical grasp space*. Throughout the paper the y-axis is the vertical axis both in object plane and also in canonical grasp space. The canonical grasp (p, α) is the canonical grasp of every other grasp of P at a placement q with orientation α.

[1] $\log^* n = \min\{i \geq 0 : \log^{(i)} n \leq 1\}$, where i is a nonnegative integer, and $\log^{(i)} n$ is the logarithm function applied i times in succession, starting with argument n.

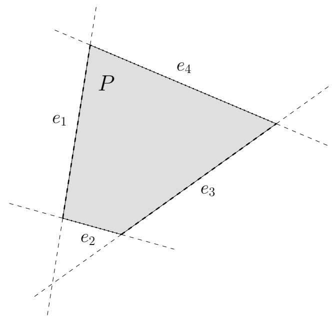

Fig. 2. The triple of edges (e_1, e_2, e_3) is non-triangular, while the triple of edges (e_3, e_4, e_1) is triangular.

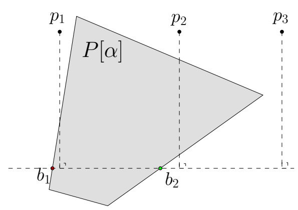

Fig. 3. (p_2, α) is an enclosing grasp while both (p_1, α) and (p_3, α) are non-enclosing grasps.

A triple of edges of P is called *triangular* if the supporting lines of the edges form a triangle that encloses P, and is called *non-triangular* otherwise. Consider Figure 2 to see examples of triangular and non-triangular triples of edges.

Consider a canonical grasp (p, α). The downward vertical ray emanating from the point $p = (x, y)$ in the plane $\theta = \alpha$ may intersect $P[\alpha]$ at an edge. This edge together with the edges of $P[\alpha]$ touched by the base fingers forms a triple of edges. The canonical grasp (p, α) is an *enclosing grasp* if the vertical ray intersects $P[\alpha]$ and if the triple of edges is triangular. Otherwise if the vertical ray does not intersect $P[\alpha]$, or if the triple of edges is non-triangular, then the canonical grasp is a *non-enclosing grasp*. See Figure 3 for examples of enclosing and non-enclosing grasps.

The following two lemmas explain two important facts about the canonical grasps of caging grasps.

Lemma 2.1: [22] The canonical grasp of every caging grasp exists and it is an enclosing grasp.

Lemma 2.2: [22] A caging grasp and its canonical grasp are reachable from each other by a sequence of translations.

If a canonical grasp is a non-enclosing grasp then it is a non-caging grasp by Lemma 2.1. The caging region of any placement q of P is a subset of the caging region of its canonical placement by Lemma 2.2. Moreover, if the canonical grasp of a three-finger grasp is non-caging then the grasp itself is non-caging too.

Let $C(\alpha)$ be the set of all placements of the third finger that together with the base fingers form a caging grasp of $P[\alpha]$. In this paper we are interested in the combinatorial complexity of $C(\alpha)$ and its computation in the worst case.

The boundary of $C(\alpha)$ consists of two x-monotone chain of curves of which the lower one is a subset of the boundary of $P[\alpha]$ [22]. Let $K(\alpha)$ be the upper part of the curves on the boundary of $C(\alpha)$. Clearly the complexity of $C(\alpha)$ is proportional to the complexity of $K(\alpha)$ plus $O(n)$ in the worst case.

Vahedi and van der Stappen [23] have proven that the placement of the third finger on $K(\alpha)$ corresponds to equilibrium grasps, which we mention here in form of a lemma.

Lemma 2.3: [23] Every placement of the third finger on $K(\alpha)$ corresponds to a two-finger equilibrium grasp or a three-finger equilibrium grasp.

Every three-finger equilibrium grasp is a canonical grasp and, thus, corresponds to a point in canonical grasp space. Two-finger equilibrium grasps, however, are not defined inside canonical grasp space because only one of the base fingers contacts the polygon. Instead, as we explain in Subsection II-C, we can represent them with their canonical grasp.

A. Visibility in Canonical Grasp Space

In this subsection we define $\mathcal{P}$ in canonical grasp space as a 3D object defined by sliding the polygon P on both base fingers (i.e. keeping the contact with both base fingers). The surface patches of $\mathcal{P}$ play an important role in the next sections in establishing a bound on the complexity of the caging region. Recall that every canonical grasp corresponds to a point in canonical grasp space. As we consider the surface patches of $\mathcal{P}$ as obstacles, we define visibility between two points in canonical grasp space (with the same x and y coordinates) as a sufficient condition that the corresponding canonical grasps have similar caging properties. Then we explain a number of properties of the surface patches of $\mathcal{P}$.

Consider an edge e of $P[\alpha]$. The edge e together with the edges e_1 and e_2 touched by the base fingers forms a triple of edges. Consider a motion of P at orientation α in which P is rotated and translated while keeping the contact with the base fingers. We refer to this motion as the sliding of P at orientation α. Clearly it is possible to slide P at orientation α in either clockwise or counterclockwise directions. As we slide P at orientation α in each direction one of the base fingers will eventually reach a vertex of the polygon, and the pair of edges touched by the base fingers change. Meanwhile, the trace of the edge e forms a surface patch, $s(e_1, e_2, e) \subset \mathbb{R}^2 \times [0, 2\pi)$, in canonical grasp space, that corresponds to the triple of edges. Clearly, $s(e_1, e_2, e)$ has a constant complexity. Let $\overline{s}(e_1, e_2, e)$ be part of $s(e_1, e_2, e)$ that are induced by all angles α for which the polygon $P[\alpha]$ is below the edge e along the y-axis. The surface patch $\overline{s}(e_1, e_2, e)$ has a constant complexity as well. If the triple (e_1, e_2, e) of edges is triangular then we call $\overline{s}(e_1, e_2, e)$ a triangular surface patch.

We define $\mathcal{P}$ in canonical grasp space as the set of patches $\overline{s}(e_1, e_2, e)$ for all edges e of the polygon P and all pairs e_1 and e_2 of edges touched by the base fingers. In other words, $\mathcal{P}$ is the set of all surface patches that are formed by considering the upper part of $P[\alpha]$ for all angles α for which d is less than the base diameter of $P[\alpha]$. The intersection of the plane $\theta = \alpha$ with $\mathcal{P}$ is the upper part of $P[\alpha]$ where α is an angle

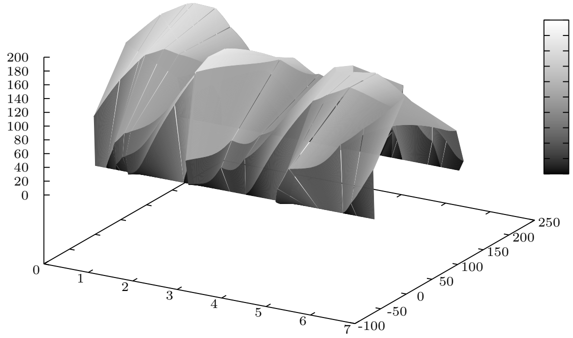

Fig. 4. The surface patches of $\mathcal{P}$.

for which d is less than the base diameter of $P[\alpha]$. Every surface patch in $\mathcal{P}$ corresponds to a set of three-finger grasps whose finger placements are on a unique triple of edges of P two of which are touched by the base fingers and the other one is touched by the third finger. Figure 4 displays $\mathcal{P}$ for the convex polygon shown in Figure 5. In Figure 5 the surface patches of $\mathcal{P}$ are projected to the space (x, θ). The horizontal gray lines display the orientations in which the pair of edges touched by the base fingers change. The surface patches of $\mathcal{P}$ form a number of connected terrains that are bounded by critical angles along the θ-axis. Since there are $O(n^3)$ triples of edges, there are $O(n^3)$ surface patches in $\mathcal{P}$.

Consider an angle α for which $P[\alpha]$ is defined. The point (p, α) in canonical grasp space corresponds to a valid canonical grasp provided that p is not inside $P[\alpha]$.

Consider two points (p, α) and (p, β) in canonical grasp space that correspond to valid canonical grasps. Here we define *visibility* between two such grasps. In this paper, the visibility is defined only along the θ-axis. In other words, we define visibility between two points in the canonical grasp space only when the line connecting the two points is parallel to the θ-axis. There are two different line segments that connect (p, α) and (p, β), which correspond to positive and negative directions along the θ-axis. We define (p, β) to be visible from (p, α) if and only if at least one of the two line segments that connect the two points intersect no surface patches of $\mathcal{P}$. According to the definition, when (p, β) and (p, α) are visible from each other, they are reachable from each other by rotation and translation; thus, they are either both caging or both non-caging. In fact, the visibility condition is a sufficient condition for the two grasps to be both caging or both non-caging but it is not a necessary condition.

B. Triangular Borders

The set $\mathcal{P}$ of surface patches consists of a number of triangular and non-triangular surface patches. The set of triangular surface patches forms a number of connected components. *Triangular borders* are the outer boundary of these connected components. In other words, triangular borders are the common boundary between the triangular surface patches and non-triangular surface patches, and also the boundary between triangular surface patches and critical angles. Every canonical grasp that corresponds to a point on the triangular borders is a non-caging grasp.

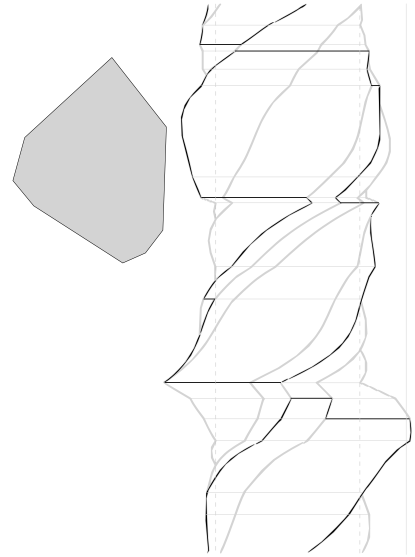

Fig. 5. Triangular borders of P projected to the space (x, θ).

In Figure 5 the surface patches of $\mathcal{P}$ are projected to the space (x, θ), in which the θ-axis is the vertical axis and the x-axis is the horizontal axis. The triangular borders are displayed in bold and the loci of vertices are displayed in gray; the distance between the two dotted vertical lines equals the distance between the base fingers. As it is displayed there are two connected components enclosed by triangular borders. (Recall that the θ-axis is circular.)

C. Two-Finger Equilibrium Grasps

In this subsection we explain about two-finger equilibrium grasps of convex polygons. Due to a general position assumption, two-finger equilibrium grasps involve the third finger and one of the base fingers.

Consider a two-finger equilibrium grasp of P at placement q with orientation α. We can show that if $P[\alpha]$ is defined, the equilibrium grasp and its canonical grasp are reachable from each other by translation, and thus they are non-caging because two-finger equilibrium grasp is non-caging. We state this fact in form of a lemma.

Lemma 2.4: Every two-finger equilibrium grasp and its canonical grasp are reachable from each other, and thus they are non-caging.

According to Lemma 2.4, since every two-finger equilibrium grasp is reachable from its canonical grasp, we can represent the two-finger equilibrium grasps in canonical grasp space by their canonical grasps.

Since for every edge there is at most one vertex with which it can form a two-finger equilibrium grasp, there are at most n pairs consisting of a vertex and an edge that induce a two-finger equilibrium grasp. Each pair of an edge and a vertex that forms a two-finger equilibrium grasp induces four curves in canonical grasp space based on the finger that is at the vertex and the base finger that is involved. We refer to these curves as the two-finger equilibrium curves each of which has a constant complexity. Similarly there are $O(n)$ pairs consisting of two vertices that define a two-finger equilibrium grasp. Two vertices that define a two-finger equilibrium grasp are necessarily antipodal.

The two-finger equilibrium curves in canonical grasp space can be classified into two groups, such that in each group the curves are circular centered around one of the base fingers in the projection to the space (x, y) [23]. The two-finger equilibrium curves of each group involve the same base finger either b_1 or b_2.

D. Escaping by Translation

In this subsection we first present some definitions and a result by Erickson et al. [4] to identify all placements of the third finger that prevent $P[\alpha]$ from escaping by pure translation; these placements form a two-dimensional region. Then we investigate the relationship between this region and the caging region.

Let the convex polygon $Q[\alpha, b_1]$ be the union of the set of all translated copies of $P[\alpha]$ touching the base finger b_1. The polygon $Q[\alpha, b_1]$ has twice number of edges of $P[\alpha]$. Every edge of $P[\alpha]$ is parallel to exactly two edges of $Q[\alpha, b_1]$. Similarly let $Q[\alpha, b_2]$ be the union of the set of all translated copies of $P[\alpha]$ touching b_2.

Let $X[\alpha]$ be the points inside the intersection of $Q[\alpha, b_1]$ and $Q[\alpha, b_2]$ that are above $P[\alpha]$. Figure 6 shows $Q[\alpha, b_1]$ and $Q[\alpha, b_2]$; the polygon is displayed in dark-gray and $X[\alpha]$ is displayed in light-gray.

Erickson et al. [4] have proven that the set $X[\alpha]$ is the set of all placements of the third finger that prevents $P[\alpha]$ from escaping by pure translation.

Lemma 2.5: [4] The polygon $P[\alpha]$ can escape by pure translation if and only if the third finger placement is outside $X[\alpha]$.

By Lemma 2.5, each edge of $P[\alpha]$ on the boundary of $X[\alpha]$ forms an enclosing triple of edges together with the edges touched by the base fingers. Let $\partial X^u[\alpha]$ be the upper-boundary of $X[\alpha]$. The set of edges of $\partial X^u[\alpha]$ consists of two continuous sets of edges: one set of edges belonging to the boundary of $Q[\alpha, b_1]$ and one set of edges belonging to the boundary of $Q[\alpha, b_2]$.

Clearly, $X[\alpha]$ is a superset of $C(\alpha)$. The following lemma proves that if a point on $\partial X^u[\alpha]$ is on the caging boundary then either the corresponding grasp corresponds to a two-finger equilibrium grasp or the corresponding grasp is on triangular borders, from which we omit the proof.

Lemma 2.6: A point on $\partial X^u[\alpha]$ that belongs to $K(\alpha)$ either corresponds to a two-finger equilibrium grasp or it is on a triangular borders.

In Figure 6, the points on $\partial X^u[\alpha]$ that can possibly be on $K(\alpha)$ are marked with small gray circles.

E. Caging and Non-caging in Canonical Grasp Space

In this subsection we prove that it is possible to compute $C(\alpha)$ for a given orientation α by using the surface patches of $\mathcal{P}$, and two types of non-caging grasps: grasps on triangular borders and canonical grasps of two-finger equilibrium grasps.

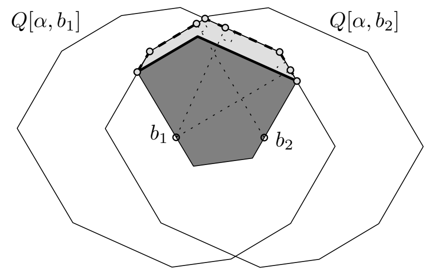

Fig. 6. Illustration of Lemma 2.6.

If $p \in C(\alpha)$, then the grasp (p, α) in canonical grasp space is only visible from other caging grasps. In other words, if (p, β) is a non-caging grasp and it is visible from (p, α) then $p \notin C(\alpha)$. In this subsection, we formulate a way to identify all points in the plane $\theta = \alpha$ that are visible from non-caging grasps. In previous sections, we have introduced two groups of non-caging grasps: grasps on triangular-borders and two-finger equilibrium grasps. In this subsection, we define vertical walls on grasps that are on triangular borders, and also on canonical grasps of two-finger equilibrium grasps (which we explain more in the next paragraphs). These walls represent a set of non-caging grasps, such that if a point on the plane $\theta = \alpha$ is visible from a point on one of these walls then that point represents a non-caging grasp. We prove the important fact that if an enclosing grasp is not visible from any of the mentioned vertical walls, then it is a caging grasp. Therefore, these walls and the surface patches of $\mathcal{P}$ provide enough information to compute $C(\alpha)$.

Consider a non-caging grasp $((x, y'), \beta)$ in canonical grasp space. Consider another grasp $((x, y), \beta)$ where $y > y'$; thus the point $((x, y), \beta)$ is vertically above $((x, y'), \beta)$ in canonical grasp space, and $((x, y), \beta)$ is a non-caging grasp as well. Consider another enclosing grasp $((x, y), \alpha)$ in canonical grasp space. If $((x, y), \alpha)$ is visible from $((x, y), \beta)$, then since $((x, y), \beta)$ is non-caging, $((x, y), \alpha)$ is non-caging as well.

The set of all grasps $((x, y), \beta)$ in canonical grasp space in which $y > y'$ and $((x, y'), \beta)$ is on triangular borders, defines a number of vertical walls, to which we refer as the *triangular-border walls*. No point in $C(\alpha)$ in the plane $\theta = \alpha$ can be visible from a point on a triangular-border wall. Therefore, the set of points in the plane $\theta = \alpha$ that are visible from no point on the triangular-border walls is a superset of $C(\alpha)$.

Recall that the canonical grasps of the two-finger equilibrium grasps form a number of curves in canonical grasp space. We define the upward vertical walls on all points on these curves, to which we refer as the *two-finger equilibrium walls*. Similar to the triangular-border walls, if (p, α) is visible from a point on a two-finger equilibrium wall then (p, α) is non-caging. The two-finger equilibrium walls intersect no surface patches of $\mathcal{P}$ by Lemma 2.4.

We define the *non-caging walls* as the union of the set of triangular-border walls and two-finger equilibrium walls. Let

$V(\alpha)$ be the set of points in the plane $\theta = \alpha$ that are visible from no point on the non-caging walls. We already know that $C(\alpha)$ is a subset of $V(\alpha)$. We prove that $C(\alpha)$ is equal to $V(\alpha)$.

First we provide a lemma which we use to prove the main result of this subsection. The following lemma states that the local minima of y-coordinates in the interior of the intersection of a triangular surface patch of $\mathcal{P}$ with a plane $\boldsymbol{x} = x$ are immobilizing grasps, from which we omit the proof. In its proof we have used a number of results proven by Vahedi and van der Stappen [22]. (Every immobilizing grasp is a caging grasp.)

Lemma 2.7: The local minima of y-coordinates in the interior of the intersection of a triangular surface patch of $\mathcal{P}$ with a plane $\boldsymbol{x} = x$ are immobilizing grasps.

The following lemma states that $C(\alpha)$ is equal to $V(\alpha)$. This result will be the foundation of our approach to establish the complexity bound. To prove the claim, we prove the equivalent lemma that every non-caging enclosing grasp $((x, y), \alpha)$ is visible from a non-caging wall. Note that if $((x, y), \alpha)$ is visible from a non-caging wall then $((x, y'), \alpha)$ with $y' > y$ is also visible from that wall. Therefore, it suffices to prove the claim for the points on $K(\alpha)$.

Lemma 2.8: For every point (x, y) on $K(\alpha)$, the point $((x, y), \alpha)$ is visible from a point on a non-caging wall.

Proof: If the point $((x, y), \alpha)$ corresponds to a two-finger equilibrium grasp then $((x, y), \alpha)$ is visible from a two-finger equilibrium wall by Lemma 2.4, and the claim follows. Assume that $((x, y), \alpha)$ does not correspond to a two-finger equilibrium grasp. As a result, the point $((x, y), \alpha)$ must correspond to a three-finger equilibrium grasp by Lemma 2.3. The number of three-finger equilibrium grasps in which the x-coordinate of the third finger is fixed and also the distance between the base fingers is d, is limited. Therefore, assume that $((x, y), \alpha)$ is a point in canonical grasp space that corresponds to a three-finger equilibrium grasp, $(x, y) \in K(\alpha)$, it is not visible from a two-finger equilibrium wall or a triangular-border wall, and has a y-coordinate that is minimal among all such three-finger equilibrium grasps.

Since $((x, y), \alpha)$ is non-caging, the canonical placement of P at orientation α can escape by first sliding along the base fingers and then by translating, according to Erickson et al. [4]. Let $((x, y), \beta)$ be the closest reachable canonical grasp at which it is possible for the canonical placement of P at orientation β to escape by translation. The set of walls visible from $((x, y), \beta)$ is the same as the walls visible from $((x, y), \alpha)$. Since the polygon can escape by translation through the grasp $((x, y), \beta)$, (x, y) is neither a point inside $C(\beta)$ nor a point inside $X[\beta]$. If (x, y) is on $K(\beta)$, then it corresponds to a two-finger equilibrium grasp or it is on a triangular border by Lemma 2.6. Therefore, assume that (x, y) is outside $C(\beta)$. Let Q be the set of all points in canonical grasp space whose corresponding canonical grasps are reachable from $((x, y), \beta)$ by sliding the polygon on the base fingers while allowing the third finger to monotonically be squeezed. Every non-caging wall visible from a point in Q is also visible from the points $((x, y), \beta)$ and $((x, y), \alpha)$. The set Q contains a number of local minima along the θ-axis with respect to the y-coordinate. Since $((x, y), \beta)$ is not visible from a triangular-border wall the local minima of Q are immobilizing grasps by Lemma 2.7. Let $((x, y_m), \alpha_m)$ be one of those immobilizing grasps and consider $(x, y') \in K(\alpha_m)$. We have $((x, y'), \alpha_m) \in Q$ and thus all non-caging walls that are visible from $((x, y'), \alpha_m)$ are also visible from $((x, y), \alpha)$. If $((x, y'), \alpha_m)$ corresponds to a two-finger equilibrium grasp the claim follows. Otherwise, the grasp $((x, y'), \alpha_m)$ corresponds to a three-finger equilibrium grasp by Lemma 2.3 for which $y' < y$. The existence of the grasp $((x, y'), \alpha_m)$ contradicts the assumption. ■

Corollary 2.9: $K(\alpha)$ is the lower boundary of the non-caging walls projected to the plane $\theta = \alpha$ not obstructed by the surface patches of $\mathcal{P}$.

III. Complexity and Computation of the Caging Region

In this section we prove that the complexity of the caging region is close to $O(n^3)$ in the worst case. We also propose an algorithm that efficiently computes the boundary of the caging region in a time that is close to $O(n^3 \log(n))$ in the worst case. The main fact we prove is that the complexity of the visible part of a surface patch of $\mathcal{P}$ not hindered by the non-caging walls is constant. This fact gives us both a way to establish an upper bound on the complexity of the caging region and obtain a solution to compute the caging region efficiently.

According to Corollary 2.9, $K(\alpha)$ is the lower boundary of the non-caging walls not obstructed by the patches of $\mathcal{P}$ projected onto the plane $\theta = \alpha$. Here, however, we formulate a slightly different way to compute $K(\alpha)$. We consider the clockwise and counterclockwise viewing direction separately. We consider only the surface patches and non-caging walls within the same connected component of triangular surface patches of $\mathcal{P}$ intersected by the plane $\theta = \alpha$. For each direction, we project the visible part of each surface patch to the plane $\theta = \alpha$ by considering only the non-caging walls as obstacles, and compute the upper boundary of the projections. Without considering the surface patches of $\mathcal{P}$, we separately project the non-caging walls to the plane $\theta = \alpha$ and compute the lower boundary of the projections. Let $V^+(\alpha)$ be the maximum of the two resulting boundaries in the clockwise viewing direction. Define $V^-(\alpha)$ similarly for the counterclockwise direction. We have the following lemma, from which we omit the proof.

Lemma 3.1: $K(\alpha)$ is the minimum of $V^+(\alpha)$ and $V^-(\alpha)$.

As we consider the non-caging walls as obstacles, we prove in Lemmas 3.5 and 3.6 that the complexity of the visible part of each surface patch is constant. Before that, we mention three results that can be easily verified.

Lemma 3.2: The two-finger equilibrium walls involving the same base finger do not intersect each other.

Observation 3.3: The triangular-border walls do not intersect each other.

Lemma 3.4: The non-caging walls do not intersect the surface patches of $\mathcal{P}$.

First we prove that the visible part of a surface patch not obstructed by triangular-border walls has constant complexity. In Lemma 3.6 we consider the two-finger equilibrium walls as well.

Lemma 3.5: The visible part of a surface patch of $\mathcal{P}$ not obstructed by triangular-border walls has constant complexity.

Proof: Since the triangular-border walls are built on the surface patches of $\mathcal{P}$ we can regard them as unbounded in both upward and downward directions. To see, consider a point which is visible behind and below a triangular-border wall. Then that point is hindered by the surface patch upon which the triangular-border wall is built.

We divide the triangular-border walls into two groups being on the left or on the right with respect to the x-axis. The walls in each group have a complete order along θ-axis according to their distance from the plane $\theta = \alpha$. Consider the left group. (The right group can be treated similarly.) Since we cannot see the side that does not face the plane $\theta = \alpha$ we consider each local maximum with respect to the x-coordinate as a wall perpendicular to the θ-axis and ignore the walls in between two consecutive local maxima. We consider the set of local maxima of the triangular borders along x-axis and then we consider the sub-sequence of local maxima in increasing order. The reason is that a local maximum is completely invisible behind another local maximum with a larger x-coordinate.

Since the complexity of a surface patch is constant it has a constant number of local minima with respect to the x-axis. Consider a local minimum that is hindered by a triangular-border wall (from the increasing sub-list). The hindering wall is the wall that is the closest to the local minimum with respect to θ-axis and is between the plane $\theta = \alpha$ and the local minimum. Then every other point of the surface patch that its x-coordinate is larger than the local maximum of the triangular-border wall and it is connected by a x-monotone curve (on the surface patch) to the hindered point (i.e. the local minimum point), is not hindered by any other triangular-border wall of the same group. Every other point of the surface patch that its x-coordinate is smaller than the local maximum of the triangular-border wall and it is connected by a x-monotone curve to the hindered point (i.e. the local minimum point), is hindered. Every point on the surface patch is connected to at least one local minimum with a x-monotone curve. ■

The following lemma is the main lemma we use to provide an upper bound on the complexity of caging region.

Lemma 3.6: The visible part of a surface patch of $\mathcal{P}$ not obstructed by non-caging walls has constant complexity.

Proof: Consider one side of the plane $\theta = \alpha$ and an arbitrary surface patch of $\mathcal{P}$. Consider the visible part of the surface patch not obstructed by the triangular-border walls. By Lemma 3.5 the complexity of this visible part is constant. Therefore, we compute the visible part of the surface patch not obstrcuted by the triangular-border walls and then remove (or ignore) the triangular-border walls. Consider the two-finger equilibrium walls that involve the base finger b_1. (We can show that there is no need to consider part of the surface patch on the right or left side of a two-finger equilibrium wall, but we do not explain it here.)

There is a complete order between the two-finger equilibrium walls along θ-axis according to their distance from the plane $\boldsymbol{\theta} = \alpha$. We traverse the two-finger equilibrium walls according to that order and we compute a sub-sequence with decreasing radii. The reason is that, a two-finger equilibrium wall with a larger radius is completely invisible behind a two-finger equilibrium wall with a smaller radius by Lemma 3.2.

Consider the local maxima of the surface patch with respect to the distance to the line $(0, 0, \theta)$. Since the complexity of a surface patch is constant the number of such local maxima is constant. If a local maximum is visible then all points that are monotonically connected to the local maximum and have have less distance with respect to the local maximum are visible too. If a local maximum is not visible, then it is hindered by a number of two-finger equilibrium walls from which consider the wall (from the decreasing sub-list) that is the furthest away from the plane $\theta = \alpha$. All points of the surface patch that are monotonically connected to the local maximum (by a path on the surface patch) and have distance less than the radius of this wall are not hindered by any other two-finger equilibrium wall of this group. All points of the surface patch that are monotonically connected (by a path on the surface patch) to the local maximum and have distance larger than the radius of this wall, are hindered.

We can similarly argue about the other group of two-finger equilibrium walls that involve the base finger b_2. ■

In the following lemma and theorem we provide an upper bound on the worst-case complexity of $K(\alpha)$.

Lemma 3.7: The complexity of $K(\alpha)$ is at most the sum of complexities of $V^+(\alpha)$ and $V^-(\alpha)$.

Proof: The set $K(\alpha)$ is the minimum of $V^+(\alpha)$ and $V^-(\alpha)$ by Lemma 3.1. Let the complexity of $V^+(\alpha)$ be of order $O(f(n))$. Consider the sequence of the breaking points of both $V^+(\alpha)$ and $V^-(\alpha)$ in increasing order with respect to their x-coordinates. Between every two consecutive breaking point, exactly one sub-curve of $V^+(\alpha)$ and exactly one sub-curve of $V^-(\alpha)$ lie within the interval. These two sub-curves intersect each other a constant number of times. Since, the total number of breaking points of both $V^+(\alpha)$ and $V^-(\alpha)$ is of order $O(f(n))$, the total number of breaking points of $K(\alpha)$ is also of order $O(f(n))$. ■

A Davenport-Schinzel sequence, $DS(m, s)$-sequence, is a sequence of m symbols in which no two symbols alternate more than s times. The lower boundary of m two-dimensional x-monotone curve segments in which no two curve segments intersect each other more than $s - 2$ times is a $DS(m, s)$-sequence. The maximum length of a $DS(m, s)$-sequence is $\lambda_s(m)$ [15].

Theorem 3.8: The complexities of both $V^+(\alpha)$ and $V^-(\alpha)$ are of order $O(\lambda_{10}(n^3))$.

Proof: The degree of the silhouette curves of the visible part of each surface patch is at most four [4]. Therefore each two curves intersect each other at most eight times. The complexity of the upper-boundary of the visible part of each surface patch projected to the plane $\theta = \alpha$ is $O(\lambda_{10}(n^3))$ by

Lemma 3.6. The complexity of the lower-boundary of the non-caging walls projected to the plane $\theta = \alpha$ is $O(\lambda_{10}(n^3))$ too. The complexity of the maximum of the two resulting chain of arcs is $O(\lambda_{10}(n^3))$ too. The proof for the last part is the same as the proof explained in Lemma 3.7. ■

We briefly explain a way to compute $K(\alpha)$ in $O(\lambda_{10}(n^3)\log n)$ time. The visible parts of all surface patches can be computed in $O(n^3 \log n)$ time. To compute the upper-boundary of the projected visible parts we use a divide-and-conquer technique. We divide the projected visible parts into two groups and compute the upper-boundary for each group separately. Then we merge the results to compute the final upper-boundary. We use the same technique to compute the lower-boundary for the projected patches of non-caing walls.

Theorem 3.9: $K(\alpha)$ can be computed in $O(\lambda_{10}(n^3)\log n)$ time.

IV. Conclusion

We have provided a worst-case bound of almost $O(n^3)$ on the combinatorial complexity of the caging region of a convex polygon with n vertices, and an algorithm with a running time close to $O(n^3 \log n)$ to compute the caging region. Both results present a major improvement over previous results. Our results have been obtained by exploiting a novel formulation of caging in terms of visibility in canonical grasp space.

The first question that comes to mind is whether the bounds reported here are tight. To gain insight into this question we will aim to construct convex polygons that have a caging region with a complexity that approaches the upper bound of $O(n^3)$. Another challenge is to extend the bounds obtained in this paper to the caging region of non-convex polygons.

Acknowledgment

M. Vahedi is supported by the Netherlands Organisation for Scientific Research (NWO). A. F. van der Stappen is partially supported by the Dutch BSIK/BRICKS-project.

References

[1] C. Davidson and A. Blake. Caging planar objects with a three-finger one-parameter gripper. In *ICRA*, pages 2722–2727. IEEE, 1998.

[2] C. Davidson and A. Blake. Error-tolerant visual planning of planar grasp. In *Sixth International Conference on Computer Vision (ICCV)*, pages 911–916, 1998.

[3] Rosen Diankov, Siddhartha Srinivasa, David Ferguson, and James Kuffner. Manipulation planning with caging grasps. In *IEEE International Conference on Humanoid Robots*, December 2008.

[4] J. Erickson, S. Thite, F. Rothganger, and J. Ponce. Capturing a convex object with three discs. *IEEE Tr. on Robotics*, 23(6):1133–1140, 2007.

[5] J. Fink, M. A. Hsieh, and V. Kumar. Multi-robot manipulation via caging in environments with obstacles. In *ICRA*, Pasedena, CA, May 2008.

[6] J. Fink, N. Michael, and V. Kumar. Composition of vector fields for multi-robot manipulation via caging. In *Proc. of Robotics: Science and Systems*, Atlanta, GA, June 2007.

[7] W. Kuperberg. Problems on polytopes and convex sets. *DIMACS Workshop on Polytopes*, pages 584–589, 1990.

[8] M. Mason. *Mechanics of Robotic Manipulation*. MIT Press, August 2001. Intelligent Robotics and Autonomous Agents Series, ISBN 0-262-13396-2.

[9] G. A. S. Pereira, V. Kumar, and M. F. M. Campos. Decentralized algorithms for multirobot manipulation via caging. *Int. J. Robotics Res.*, 2004.

[10] P. Pipattanasomporn and A. Sudsang. Two-finger caging of concave polygon. In *ICRA*, pages 2137–2142. IEEE, 2006.

[11] P. Pipattanasomporn, P. Vongmasa, and A. Sudsang. Two-finger squeezing caging of polygonal and polyhedral object. In *ICRA*, pages 205–210. IEEE, 2007.

[12] E. Rimon and A. Blake. Caging 2d bodies by 1-parameter two-fingered gripping systems. In *ICRA*, volume 2, pages 1458–1464. IEEE, 1996.

[13] E. Rimon and J. W. Burdick. Mobility of bodies in contact—part i: A 2nd–order mobility index for multiple–finger grasps. *IEEE Tr. on Robotics and Automation*, 14(5):696–717, 1998.

[14] A. Rodriguez and M. Mason. Two finger caging: squeezing and stretching. In *Eighth Workshop on the Algorithmic Foundations of Robotics (WAFR)*, 2008.

[15] M. Sharir and P. K. Agarwal. *Davenport-Schinzel sequences and their geometric applications*. Cambridge University Press, New York, NY, USA, 1996.

[16] A. Sudsang. Grasping and in-hand manipulation: Geometry and algorithms. *Algorithmica*, 26(3-4):466–493, 2000.

[17] A. Sudsang and T. Luewirawong. Capturing a concave polygon with two disc-shaped fingers. In *ICRA*, pages 1121–1126. IEEE, 2003.

[18] A. Sudsang and J. Ponce. A new approach to motion planning for disc-shaped robots manipulating a polygonal object in the plane. In *ICRA*, pages 1068–1075. IEEE, 2000.

[19] A. Sudsang, J. Ponce, M. Hyman, and D. J. Kriegman. On manipulating polygonal objects with three 2-dof robots in the plane. In *ICRA*, pages 2227–2234, 1999.

[20] A. Sudsang, J. Ponce, and N. Srinivasa. Grasping and in-hand manipulation: Experiments with a reconfigurable gripper. *Advanced Robotics*, 12(5):509–533, 1998.

[21] A. Sudsang, F. Rothganger, and J. Ponce. Motion planning for disc-shaped robots pushing a polygonal object in the plane. *IEEE Tr. on Robotics and Automation*, 18(4):550–562, 2002.

[22] M. Vahedi and A. F. van der Stappen. Geometric properties and computation of three-finger caging grasps of convex polygons. In *IEEE International Conference on Automation Science and Engineering (CASE)*, 2007.

[23] M. Vahedi and A. F. van der Stappen. Caging polygons with two and three fingers. *Int. J. Robotics Res.*, 27(11/12):1308–1324, 2008.