Probing Convex Polytopes

D. Dobkin†, *H. Edelsbrunner*‡, *C. K. Yap*§

†Department of Computer Science
Princeton University
Princeton, New Jersey

‡Department of Computer Science
University of Illinois
Urbana-Champaign, Illinois

§Courant Institute of
Mathematical Sciences
New York University
New York, New York

ABSTRACT

We investigate the complexity of determining the shape and presentation (i.e. position with orientation) of convex polytopes in multi-dimensional Euclidean space using a variety of probe models.

1. Introduction

This paper considers the problem of discovering the environment by means of simple sensory equipment. We are motivated by robots that can make 'probes' into their surroundings. The term 'probe' is intended to cover a whole range of crude sensory devices. An example of a probe might be a robot arm moving in a fixed direction until it contacts an obstacle. Such a probe might yield the spatial position of a point on the obstacle. Another example is an ultra-sound device that can detect the proximity (but not the precise location) of other objects. In contrast to such crude devices, one might think of using vision to discover one's environment: the amount of data gathered and processed in vision systems is many orders of magnitude greater than the kind of data gathered by probes. Consequently, the computational issues are rather different. While vision is a much studied topic, the area of processing crude sensory data is relatively new. It is reasonable to ask why should one bother with probes when we can accomplish much more using data-intensive sensory equipment. The answer is that data-intensive sensory equipment (such as a camera) may be uneconomical, too delicate or physically impossible to install. Even if we could gather such data, the processing of such data may be computationally too expensive. Probe devices are usually more robust, cheaper and smaller; the processing of such data is also expected to be relatively cheap. Device engineers have invented many ingenious methods to gather such crude sensory data. The theoretical understanding of processing such data is clearly important for the effective use of such devices. Furthermore, our understanding of the problem enhances our understanding of convex polytopes.

One of the first papers in this area is [CY87] which considers the problem of determining a convex polygon using 'finger' probes. Such probes can be imagined as a point moving from infinity along a straight line until the point contacts an object. The data yielded by the probe is this contact point. (If the probe misses all objects, the contact point is 'at infinity', by convention.) They proved that for a convex n-gon, $3n$ probes are sufficient and $3n - 1$ probes are necessary. This may appear a little surprising since their algorithm only assumes that the object is a convex n-gon for some unspecified n. In this paper, we consider probing convex polytopes in d-dimensional Euclidean spaces $\mathcal{E}^d$, $d \geq 2$, and look at several reasonable models of probes:

- a moving point ('finger probe')

† NSF Grant MCS-83-03926 and DCR-85-05517

‡ Amoco Foundation Faculty Development in Computer Science

§ NSF Grant DCR-84-01633 and DCR-84-01898

- a moving hyperplane ('hand probe')
- a light-source yielding a silhouette
- a plane yielding a cross-section
- a moving line in $\mathcal{E}^3$
- finger probe with $\varepsilon > 0$ uncertainty

We consider a complexity model that counts the worst case number of probes where the successive probes are adaptive, i.e., can depend of the previous probe outcomes. Note that we do not count the cost of determining successive probes. This assumption, similar to that of counting only comparisons in sorting-related problems, can be justified under various circumstances.

We use the following standard terminology: for any point set $X \subseteq \mathcal{E}^d$, its boundary and interior is denoted ∂X and $int(X)$, respectively. A (possibly non-convex) closed polyhedron P is *star-shaped* if there is a point x in the polyhedron such that the line segment connecting x with any other point of P is also contained in P. The set of all such points x is the *kernel* of P. A *polytope* is a bounded polyhedron. The boundary of a polyhedron $P \subseteq \mathcal{E}^d$ is partitioned in the usual way into i-faces for $i = 0, \ldots, d-1$ where we define an *i-face* to be any relatively open subset of $\mathcal{E}^d$ homeomorphic to $\mathcal{E}^i$. Thus the interior of P is a d-face. Two faces are *incident* if one of the faces is contained in the closure of the other. We use *vertex*, *edge*, *facet* and *cell* as synonyms for 0-face, 1-face, $(d-1)$-face and d-face. Let $f_i(P)$ denote the number of i-faces of P. The partition of $\mathcal{E}^d$ into a set of i-faces ($i = 0, \ldots, d$) such that the closure of each face is equal to a union of faces is called a *cell-complex*. In this paper, the cell complex is determined by a finite set of hyperplanes.

The organization of this paper is as follows: section 2 investigates the complexity of finger-probing a convex polytope $P \subseteq \mathcal{E}^d$. We show that $f_0(P) + f_{d-1}(P)$ finger probes are necessary and $f_0(P) + (d+2)f_{d-1}(P)$ finger probes are sufficient. We also show that probing with hyperplanes is dual to finger probing. Section 3 addresses more complex probes such as light-sources which yield silhouettes and intersecting hyperplanes that yield cross-sections. In section 4 we consider line probes in $\mathcal{E}^3$. It is important to realize that in line probes, no information about the points of contact between the line and the polytope is given; that such probes can determine polytopes is not obvious. We give an algorithm using a number of probes that is linear in the number of faces of the polytope. Section 5 examines probes that have uncertainty of $\varepsilon > 0$. The issues here are considerably more subtle. It is not even clear that a meaningful generalization of our other probe models can be posed in this context. Finally, section 6 discusses the results of this paper and indicates directions for future work.

2. Finger Probing

Let $P \subseteq \mathcal{E}^d$ be a convex polytope (so P is closed and bounded) whose interior contains the origin O. Nothing else is known about P. The goal is to infer complete information about P on the basis of probes as we now describe. A *finger probe* F is an unbounded directed line. We imagine a point moving along the directed line F. The *contact point* $C(F, P)$ of F at P is the first point on P that F intersects along its direction. The *probe path* $\pi(F, P)$ of F to P is the directed half-line contained in the line F, with direction consistent with F and terminating at $C(F, P)$. We also write $C(F)$ and $\pi(F)$ for the contact point and path, respectively, when P is understood. If F does not intersect P, we say that the probe *misses* P and the contact point is at infinity, denoted $C(F) = \infty$; the probe path $\pi(F)$ is equal to F in this case. The finger probe model was first studied in [CY87] for the case $d = 2$. We first investigate the case $d = 3$ and then generalize the result to arbitrary d.

2.1. Bounds in $\mathcal{E}^3$

Let $F_1, \ldots, F_m$ be a sequence of probes and let $c_i = C(F_i)$ and $\pi_i = \pi(F_i)$, $i = 1, \ldots, m$, denote the corresponding contact points and probe paths. We next introduce some structures to capture the information gained from these probes. For any set of points X, let $conv(X)$ denote the smallest closed convex set containing X.

$\mathcal{P}_m = \{ c_1, \ldots, c_m \}$ is the (current) set of contact points.

$\mathcal{H}_m = conv(\mathcal{P}_m \cup \{ O \})$ is the convex hull of the set of contact points together with the origin O.

$\mathcal{A}_m$ is the cell complex defined by the planes that support facets of $\mathcal{H}_m$.

Clearly the convexity of P implies

$$\mathcal{H}_m \subseteq P.$$

Next we want to define the smallest set O_m^* that is guaranteed to contain P. Each probe path π_i casts a 'shadow' which consists of all those points x that cannot be contained in P by virtue of the convexity of P and the fact that every point in $\mathcal{P}_m$ lies on the boundary of P. This shadow is precisely defined by the set S_i of those points x such that the probe path π_i would intersect the interior of

$$conv(\mathcal{P}_m \cup \{ O \} \cup \{ x \}).$$

We can define

$$O_m^* = \mathcal{E}^3 - \bigcup_{i=1}^{m} S_i$$

Thus $O_m^* = \mathcal{E}^3$ for $m = 0$. The set O_m^* can be rather complicated to understand, and for the purposes of upper bounds, we prefer to use an alternative set defined as follows. Let v be a contact point in $\mathcal{P}_m$. Let $h_1, \ldots, h_k$ ($k \geq 1$) be the open half-spaces where each h_i does not contain O but is bounded by a plane containing a facet incident to v. Define K_v to be the intersection of $h_1, \ldots, h_k$. If v is a vertex of $\mathcal{H}_m$ then K_v is a cone with vertex v; if v is in the relative interior of an edge (resp. facet) of $\mathcal{H}_m$ then K_v is a wedge (resp. a half-space). Notice that K_v is the set of points x such that v belongs to the interior of $conv(\mathcal{P}_m \cup \{ O, x \})$. It follows that K_v is contained in S_i if $v = c_i$. We now define

$$O_m = \mathcal{E}^3 - \bigcup_{v} K_v$$

where v ranges over $\mathcal{P}_m$. Clearly, for all $m \geq 0$,

$$O_m^* \subseteq O_m$$

and

$$\mathcal{H}_m \subseteq P \subseteq O_m.$$

Lemma 1. O_m is star-shaped with $\mathcal{H}_m$ as its kernel.

Proof. Define $\bar{K}_v$ as the complement of K_v, and define H_v as the kernel of $\bar{K}_v$. Thus, H_v is the intersection of the closed half-spaces that are complements of the half-spaces $h_1, \cdots, h_k$ used in defining K_v. If v is a vertex of $\mathcal{H}_m$ then H_v is the smallest cone with apex v that contains $\mathcal{H}_m$; if v belongs to an edge of $\mathcal{H}_m$ then H_v is a wedge; and if v belongs to a facet of $\mathcal{H}_m$ then H_v is a half-space. In any case, H_v contains $\mathcal{H}_m$, and furthermore, $\mathcal{H}_m$ is the intersection of all H_v, for v ranging over all vertices of $\mathcal{H}_m$. By definition of O_m as the intersection of all $\bar{K}_v$, the kernel of O_m contains the intersection of all H_v. Therefore, $\mathcal{H}_m$ is subset of the kernel of O_m. We conclude that $\mathcal{H}_m$ is exactly the kernel of O_m since every facet of $\bar{K}_v$ contains a facet of O_m, where v is a vertex of $\mathcal{H}_m$. Q.E.D.

If O belongs to the interior of $\mathcal{H}_m$, and we only know about the contact points $\mathcal{P}_m$ (i.e. we know nothing about the probe paths) then $\mathcal{H}_m$ and O_m are the strongest possible sets in the sense that $\mathcal{H}_m$ (resp. O_m) cannot be replaced by a larger (resp. smaller) set in the above lemma. Initially, $\mathcal{H}_0$ consists of just the origin and $O_0 = \mathcal{E}^3$.

It is natural to say that P is *determined* when $\mathcal{H}_m = P = O_m$. Now we propose a probing strategy to determine P.

Definition. A facet of $\mathcal{H}_m$ or of O_m is said to be *verified* if the plane defined by that facet contains at least four co-planar points of $\mathcal{P}_m$ such that one of these points is in the relative interior of convex hull formed by the other points; otherwise the facet is *unverified.* The plane containing such a facet is said to be verified or unverified according as the facet is verified or unverified. A probe aimed at the relative interior of an unverified facet is said to be *trying to verify* that facet.

Definition. A vertex of $\mathcal{H}_m$ is said to be *verified* if it is incident to at least three verified facets of $\mathcal{H}_m$; otherwise, it is *unverified.* A vertex of O_m is verified if it is a verified vertex of $\mathcal{H}_m$.

We sometimes call an unverified facet a 'conjectured' facet. An attempted verification of a facet *succeeds* if the contact point lands on the facet. This terminology is justified by the observation that a verified facet of $\mathcal{H}_m$ is necessarily contained in a facet of P. Roughly speaking, our strategy consists of sending probes where each probe tries to verify a conjectured facet. If this objective fails then we get a contact point p which may or may not be co-planar with other facets of $\mathcal{H}_m$; in any case, p is a new vertex of $\mathcal{H}_{m+1}$. If we can somehow guarantee that an unverified facet of $\mathcal{H}_m$ has at most a constant number of vertices, then the pigeonhole principle guarantees that $\mathcal{H}_m$ cannot have too many vertices so that we must eventually succeed in verifying a facet. To guarantee this constant, our probes must avoid verified facets and we need to make sure that we land on at most one unverified facet. To meet the first goal, we use the fact that a verified facet of O_m contains a facet of P and can therefore be used as an upper bound for the facets of P they contain. To meet the second goal, we use the arrangement $\mathcal{A}_m$: we can land on two or more unverified facets only if we land on an edge or a vertex of $\mathcal{A}_m$. The union of edges and vertices of $\mathcal{A}_m$ can be easily avoided, however, since it is only a one-dimensional subcomplex. The constant that we will obtain is four which implies that each verified facet will contain at most five contact points.

Note that P is not completely explored as long as we do not have a contact point at each vertex. For, it is possible that a facet of the polytope (of potentially arbitrarily small area) has been overlooked. The verification of the vertices, however, is automatic: this is because in choosing probes that avoid verified facets of O_m, we are forced to aim at vertices of O_m.

The method just outlined is a generalization of the method of Cole and Yap. It should be noted, however, that the necessity of controlling the number of contact points in a single unverified facet is a new phenomenon in three dimensions. In two dimensions, one of any three collinear points is in the relative interior of the convex hull of the other two. In three dimensions, one can have an arbitrarily large number of co-planar points without having any point in the relative interior of the convex hull of the others. Surprisingly, nothing new arises in dimension four or greater.

The next lemma establishes a bijective correspondence between verified facets of $\mathcal{H}_m$ and verified facets of O_m.

Lemma 2. Each verified facet of $\mathcal{H}_m$ is contained in a verified facet of O_m, and each verified facet of O_m contains a verified facet of $\mathcal{H}_m$.

Proof. Let f be a verified facet of $\mathcal{H}_m$. If v is a point in $\mathcal{P}_m$ in the relative interior of f then K_v is a half-space bounded by the plane through the facet f. This K_v in turn determines a unique verified facet f' of O_m containing f. Conversely, a verified facet of O_m contains four contact points with one in the relative interior of the convex hull of the other three, by definition. These four contact points necessarily lie in a verified facet of $\mathcal{H}_m$. Q.E.D.

This lemma allows us to be sloppy when we refer to 'verified facets' without saying whether they are facets of $\mathcal{H}_m$ or of O_m. For the next lemma to make sense, the reader should realize that even when all the facets of O_m are verified, there could be facets of $\mathcal{H}_m$ that are not yet verified.

Lemma 3. The following are equivalent statements: (i) All facets of O_m are verified. (ii) O_m is convex. (iii) each vertex of $\mathcal{H}_m$ is incident to some verified facet of $\mathcal{H}_m$.

Proof. (i) $\Rightarrow$ (ii). Suppose all facets of O_m are verified. If a facet f of O_m is verified then O_m is contained in the half-space containing O determined by the plane of f. Let C denote the intersection of all such (closed) half-spaces. Thus $O_m \subseteq C$. Equality between O_m and C follows since O_m has no facet that does not belong to a facet of C.

(ii) $\Rightarrow$ (iii). Observe that if $v \in \mathcal{H}_m$ is not incident to a verified facet of $\mathcal{H}_m$ then for any ball B centered at v with sufficiently small radius, $B \cap K_v$ does not intersect K_u for all $u \neq v$, $u \in \mathcal{P}_m$. This means that $B - K_v = B \cap O_m$. This shows that O_m is not convex, since $B - K_v$ is not convex.

(iii) $\Rightarrow$ (i). Observe that the proof is complete if we show that any unverified facet of O_m is contained in a facet of K_v for some v that is not incident to any verified facet of $\mathcal{H}_m$. Let f be any unverified facet of O_m. There exists a $v \in \mathcal{P}_m$ such that f is contained in some facet f' of K_v. For the sake of contradiction, assume that some facet of K_v contains a verified facet g of O_m. Let w be a point of $\mathcal{P}_m$ in g (w exists because g is verified). Then $K_v \subseteq K_w$ which implies that one facet of K_v belongs to the plane ∂K_w and that the other facets of K_v are in the interior of K_w. Hence f cannot be a facet of O_m since $O_m \subseteq \mathcal{E}^3 - K_w$. This is a contradiction. **Q.E.D.**

For convenience, we define $\mathcal{V}_m$ to be the intersection of the 'verified' half-spaces where a 'verified' half-space is one containing $\mathcal{H}_m$ and bounded by a verified plane. Initially $\mathcal{V}_m = \mathcal{V}_0$ is the entire space $\mathcal{E}^3$. Clearly $\mathcal{V}_m$ is convex and $O_m \subseteq \mathcal{V}_m$. From the preceding lemma, we see that equality holds precisely when all the facets of O_m are verified. We will use the vertices of $\mathcal{V}_m$ to choose our probe lines.

We shall maintain the following invariant (H) in our probing strategy. This guarantees that each unverified facet of P will contain at most five contact points.

(H) Each unverified facet of $\mathcal{H}_m$ is a triangle, with at most one exception which may be a quadrilateral.

This invariant is easily initialized by forming a tetrahedron $\mathcal{H}_4$ about the origin. Our probing strategy is as follows:

> If all facets of $\mathcal{H}_m$ are verified, we are done. Otherwise, choose any unverified facet f of $\mathcal{H}_m$. If there are any unverified quadrilateral, choose f to be it. Treating the plane of f to be horizontal and lying above the origin, there are two cases: if there is any vertex v of $\mathcal{V}_m$ above the plane of f, then v is unverified and we let F be a probe aimed from v to any point x in the relative interior of f. If v does not exist, pick any F aimed at an interior point x where the path of F right up to x lies entirely in $\mathcal{V}_m$. By a suitable perturbation of x, we can ensure that F satisfies the additional property of not intersecting any vertex or edge of the cell complex $\mathcal{A}_m$.

It is clear that F as specified above exists. For a later application, we made an additional observation about the preceding probing strategy: it is easy to see that we can further assume that the probe we choose is aimed at the origin.

Lemma 4. The invariant (H) is maintained by the outcome of probe F.

Proof. Note that the contact point p of F will lie between v and x; if v does not exist, then p occurs before or at x. If $p = x$ then we have verified the face f. If v exists and $p = v$ then we have verified v; v is now a contact point and it is incident at at least three verified facets of O_{m+1}. Finally, if p is equal to neither x nor v then p is a new vertex of $\mathcal{H}_{m+1}$ and we form some new facets, each incident at p. Then f is no longer a facet. Notice that if a point lies in a cell of $\mathcal{A}_m$ then it is not coplanar with any three contact points that belong to the closure of a common facet of $\mathcal{H}_m$. Furthermore, if a point belongs to a facet of $\mathcal{H}_m$ then it is coplanar with the vertices of only one facet of $\mathcal{H}_m$. Our probing strategy assures that F does not hit a verified facet, and it hits at most one plane defined by an unverified facet of $\mathcal{H}_m$. It follows that F creates at most one new quadrilateral while the only old quadrilateral, if any, disappears. **Q.E.D.**

Using this strategy, we continue until P is determined. The partial correctness of the algorithm is clear. To show termination, note that with each probe, we either increase the number of unverified vertices of $\mathcal{H}_m$ or verify a facet or verify a vertex. By the pigeonhole principle, the number of unverified vertices is at most $3(f_2(P) - k) + 1$ where k is the number of verified facets. Termination is assured since we never verify a facet or a vertex more than once.

There is a contact point at each vertex of $\mathcal{H}_m$, and each facet has at most 5 contact points incident to it (one of these contact points lies in the relative interior of the triangle or quadrilateral formed by the others). This gives an upper bound of $f_0(P) + 5f_2(P)$.

We prove a lower bound by a straightforward adversary argument. First observe that every vertex of P must be probed. Next note that the relative interior of each facet of P must be probed. Combining this with the preceding upper bound, we obtain:

Theorem 5. Let P be a convex polytope in $\mathcal{E}^3$. Let $T_F(P)$ be the worst case number of finger probes necessary to determine P. Then

$$f_0(P) + f_2(P) \le T_F(P) \le f_0(P) + 5f_2(P).$$

2.2. Higher-dimensional finger probing

Seeing that probing in 3-dimensions encounters a difficulty that has no analogue in the planar case, it is not immediately clear whether we will encounter yet new difficulties in 4-dimensions and beyond. Fortunately, nothing new arises and we can obtain a fairly straightforward generalization. We merely sketch a proof.

Theorem 6. Let P be a convex polytope in $\mathcal{E}^d$, $d \ge 3$, and $T_F(P)$ be the worst case number of finger probes necessary to determine P. Then

$$f_0(P) + f_{d-1}(P) \le T_F(P) \le f_0(P) + (d+2)f_{d-1}(P).$$

Proof. The lower bound argument of the case $d = 3$ clearly generalizes. To obtain the upper bound, we organize our probes so that probes either land on $\mathcal{H}_m$, on a facet of $\mathcal{A}_m$, in a cell of $\mathcal{A}_m$, or at a vertex of O_m. To see this, we can check that all the lemmas for the 3-dimensional algorithm hold. The corresponding invariant is that an unverified facet of $\mathcal{H}_m$ have d contact points, with at most one exception that may have $d + 1$ contact points. **Q.E.D.**

2.3. Hyperplane probes

In contrast to the above, we now define a *probe* to be a moving hyperplane H approaching from infinity in the direction of its normal. The *contact hyperplane* $C(H)$ is the location of H where it first contacts P. Interestingly, this probe model can be reduced to finger probing using a duality transformation. For a point $p \ne O$ in $\mathcal{E}^d$, let its *dual* be the hyperplane $D(p)$ where the vector p is normal to $D(p)$ and the point $\frac{p}{|p|^2}$ belongs to $D(p)$. For any hyperplane h that avoids O, h^+ denotes the closed halfspace bounded by h and containing O. Then define

$$D(P) = \bigcap_{p \in P} D(p)^+$$

as the dual image of P. It is straightforward to verify that $D(p)$ is a hyperplane that touches $D(P)$ in a vertex if and only if p belongs to a facet of P, and that $D(p)$ supports a facet of $D(P)$ if and only if p is a vertex of P. For all computations it is therefore sufficient to consider the dual images of all responses (that is, images under D inverse) and to apply the strategy for finger probing as explained above. It is important to realize that we can apply the finger probing strategy of section 2.1 because, by an earlier remark, we can assume the finger probes of that strategy are directed at the origin. This yields

Theorem 7. Let P be a convex polytope in $\mathcal{E}^d$, $d \ge 3$, and let $L_H(P)$ be the worst case number of hyperplane probes necessary to determine P. Then

$$f_0(P) + f_{d-1}(P) \le L_H(P) \le (d + 2)f_0(P) + f_{d-1}(P).$$

3. Cross section and silhouette probing

We consider two models of probing a polytope in 3-dimensions here. In the first case, a probe consists of a direction in which a polytope cross section is to be taken. All cross sections are assumed to pass through the origin so that it suffices to specify the direction via a normal vector. In the second case, we consider probes which are specified by the position of a point source of light. Again, to simplify the results, we assume the point source is at infinity. The result of such a probe is the silhouette of the polytope generated by the light source. Duality shows that these two models are really the same:

Theorem 8. Cross section probing and silhouette probing are duals.

To obtain upper bounds, we reduce finger probing to cross section probing as follows: a finger probe F together with the origin O determine a plane P which we use to define our cross section probe. If F passes through the origin then we can choose P to be any plane containing F. Clearly, the outcome from this cross section probe yields at least as much information as the outcome of F. Thus, it follows from the results in section 2 that for any convex polytope $P \subseteq \mathcal{E}^3$, if $T_C(P)$ is the worse case number of cross section probes through the origin necessary to determine P then

$$T_C(P) \le f_0(P) + 5f_2(P)$$

It is easy to show a linear lower bound. First we observe that for any P, there is a perturbation P' of P such that $T_C(P') \ge \frac{f_0(P')}{2}$. This comes from the fact that any correct algorithm must pass a cross section through each vertex of P. The bound then follows from the fact that we perturb P to P' so that every three vertices of P' define a plane that avoids the origin O. In fact this bound holds for any 'certificate' (non-deterministic algorithm) for P'. *A fortiori* no sublinear algorithm is possible. We may conclude that our above upper bound on $T_C(P)$, which is given by an apparently wasteful reduction of finger-probes to cross-section probes, is at most a constant factor from the optimal.

4. Line probing

We consider probing a polytope in 3-dimensions using line probes. Intuitively, a line probe λ consists of a line sweeping out a plane H_λ. The position of the line at time $t \in \mathbf{R}$ ($\mathbf{R}$ is the set of all real numbers) is λ_t, where all the lines λ_t are parallel. The probe contacts a given bounded closed subset $P \subseteq \mathcal{E}^3$ the first time t_0 such that $\lambda_{t_0} \cap P \ne \emptyset$. It is important to realize that we are told the line λ_{t_0} but get no information about the set $\lambda_{t_0} \cap P$. One motivation for such probes is the IBM RS-1 robot which has LED sensors on two opposing robot fingers. The invisible ray between the two finger corresponds to a sweeping line that is cut off on contact with an object. The paper [CY87] uses the same model to motivate finger probes: here we imagine an object to be a polygonal piece of cardboard somehow supported to stand with one of its edges on a table.

More precisely, a *line probe* λ is a triple

$$(d, \hat{\mathbf{n}}, \hat{\mathbf{v}})$$

where $d \in \mathbf{R}$, $\hat{\mathbf{n}}$ and $\hat{\mathbf{v}}$ are unit vectors that are orthogonal to each other. See figure 1. The probe determines the *probe plane* H_λ that is normal to $\hat{\mathbf{n}}$ at distance d from the origin O. For any $t \in \mathbf{R}$, let λ_t denote the line

$$\{ s\hat{\mathbf{u}} + (d\hat{\mathbf{n}} + t\hat{\mathbf{v}}) : s \in \mathbf{R} \}$$

where $\hat{\mathbf{u}} = \hat{\mathbf{n}} \times \hat{\mathbf{v}}$. Note that λ_t is contained in the plane H_λ and the line moves in the direction of $\hat{\mathbf{v}}$. The *result line* of the probe on a bounded closed subset P is the line λ_{t_0} where

$$t_0 = \min\{ t \in \mathbf{R} : \lambda_t \cap P \ne \emptyset \}.$$

If $\lambda_t \cap P = \emptyset$ for all t then we set $t_0 = \infty$ above and we say that the λ_∞ is undefined. The result of the probe is then said to be *infinite* or *finite* according to whether $t_0 = \infty$ or not.

We call $\hat{\mathbf{u}}$ the *probe orientation* and $\hat{\mathbf{v}}$ the *probe direction*. A probe $\lambda = (d, \hat{\mathbf{n}}, \hat{\mathbf{v}})$ is *centered* if $d = 0$. If $\lambda' = (d', \hat{\mathbf{n}}', \hat{\mathbf{v}}')$ is another probe, then λ and λ' are *opposite* if $H_\lambda = H_{\lambda'}$ and $\hat{\mathbf{v}}' = -\hat{\mathbf{v}}$. Again, λ and λ' are said to be *parallel* if $\hat{\mathbf{u}} = \pm\hat{\mathbf{u}}'$ where as usual $\hat{\mathbf{u}} = \hat{\mathbf{n}} \times \hat{\mathbf{v}}$ is the probe orientation. Note that in this case, either H_λ and $H_{\lambda'}$ are parallel, or else $\hat{\mathbf{u}}$ is parallel to the line $H_\lambda \cap H_{\lambda'}$. More importantly, the result lines (if both finite) of a pair of parallel probes are either identical or they determine a

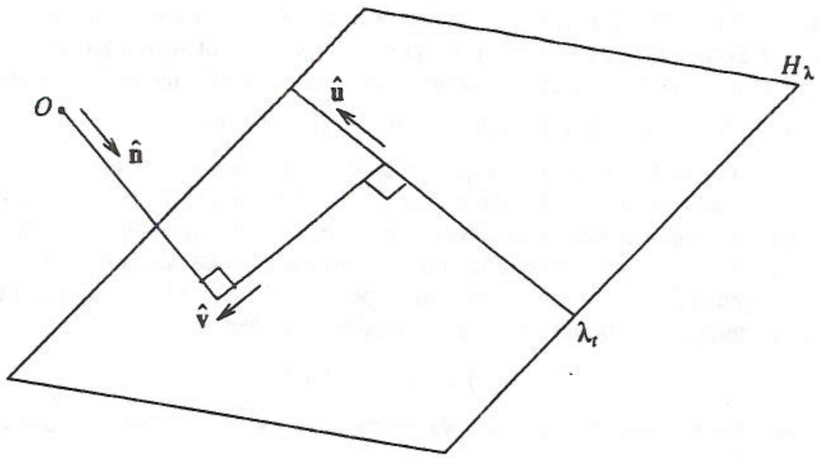

Figure 1.

plane.

Assume we are given a convex polytope P whose interior contains the origin O. In the following, we will say that a point (resp. line, plane) is *verified* at a particular moment if we can deduce from the result lines that the point (resp. line, plane) is a vertex (resp. contains an edge, contains a facet) of P. Before describing the algorithm we describe a basic subroutine.

Half-profile computation. Let H be any closed half-space such that its bounding plane ∂H intersects P. Let $\hat{w}$ be any direction parallel to the plane ∂H. Let $h(\hat{w})$ denote the plane normal to $\hat{w}$ and passing through the origin. Then a *$(H, \hat{w})$-profile* (of P) is the planar projection of $H \cap P$ along the direction of $\hat{w}$ onto $h(\hat{w})$. Let the projection of ∂H in $h(\hat{w})$ be L. Thus the $(H, \hat{w})$-profile is a polygon $R \subseteq h(\hat{w})$ with one side abutting L. Suppose R has m vertices, that is,

$$R = (u_1, u_2, \ldots, u_m).$$

In our application, we can assume that $m \geq 2$ and without loss of generality, assume that the edge $[u_1, u_m]$ abuts L. Again, in our application below, we are told two distinct points on the edge $[u_1, u_m]$. It is not hard to see that a simple modification of the algorithm of Cole and Yap, we can either (a) determine that $m = 2$ using one probe or else (b) determine that $m > 2$ and compute R using $3m - 2$ line probes. The procedure is called a *half-profile computation*. To see why the half-profile computation is essentially a finger probe problem in the plane, observe that if we restrict ourselves to line probes λ whose probe orientations are parallel to $\hat{w}$, then the behavior of the moving line λ_t is faithfully recorded by the intersection of λ_t with $h(\hat{w})$. Clearly the intersection of the probe plane H_λ with $h(\hat{w})$ gives rise to a (planar) finger probe of the $(H, \hat{w})$-profile.

Remark: Note the similarity between this subcomputation and the shadow probes of the previous section. Since this subcomputation already makes a linear number of probes, and there is a linear lower bound on the number of shadow probes, this suggests a quadratic lower bound in the number of line probes. Thus our linear upper bound on line probing is somewhat surprising.

We are ready to give the overall algorithm. At the beginning of the nth step below ($n \geq 0$), we maintain the following invariant. There is a set $\mathcal{V}_n$ of verified vertices, and $\mathcal{H}_n$ is defined to be the convex hull of $\mathcal{V}_n \cup \{ O \}$. Let $\mathcal{F}_n$ denote the subset of the facets of $\mathcal{H}_n$ each of whose plane is verified.

To initialize, we do a 'full profile' of P along any direction $\hat{w}$, in exact analogy with the half-profile computation above. That is, we restrict ourselves to line probes whose orientations are parallel to $\hat{w}$, and hence such line probes are faithfully represented by finger probes in the plane $h(\hat{w})$, and a direct application of the planar probing problem in [CY87] gives us the desired profile, which is a projection Q

of P onto $h(\hat{w})$.

We then continue as follows. Note that the polygon $Q \subseteq h(\hat{w})$ is such that its vertices are projections of vertices (possibly edges) of P and its edges are projections of edges (and possibly facets) of P. Let e be an edge of Q. Consider the plane $h'(e)$ containing e and perpendicular to $h(\hat{w})$. We know that there is at least one edge of P, and possibly one facet of P, lying in $h'(e)$. Our goal is to verify the vertices of P in $h'(e)$. This is easily reduced to "hyperplane probe" in $h'(e)$, as described in section 2.3. In other words, we restrict our attentions to line probes λ whose probe planes are $h'(e)$. By duality, hyperplane probes are transformed to the planar finger probing problem of Cole and Yap. Note that we already know two parallel hyperplane probe outcomes – namely the two lines parallel to $\hat{w}$ and passing through the endpoints of e. Using at at most $3m$ additional hyperplane probes, we can determine the $m \geq 1$ edges of P lying in $h'(e)$. Although the original formulation of Cole and Yap does not allow $m=1$, this is easily modified to suit our situation, allowing us to charge at most 3 probes to each edge of P found in this way. For reference, we call this probing of $h'(e)$ an *edge verification procedure*. The probes made in the full-profile computation of Q can be charged to the vertices of P that project onto vertices of Q; and hence we charge at most 3 probes to each vertex of P found. We now have our initial set $\mathcal{V}_0$ of verified vertices, and of verified facets $\mathcal{F}_0$. The general nth step is as follows:

> If every facet of $\mathcal{H}_n$ is verified then we halt and $P = \mathcal{H}_n$. Suppose some facet f is still unverified. Let H be the half-space where the plane ∂H supports $\mathcal{H}_n$ at f and H does not contain the interior of $\mathcal{H}_n$. Choose any direction $\hat{w}$ in ∂H such that $\hat{w}$ is not parallel to any edge of f. Construct the $(H, \hat{w})$-profile. If f is actually a facet of P, then one probe would verify this. In this case $\mathcal{H}_{n+1} = \mathcal{H}_n$. Otherwise, we get a non-trivial profile $R = (u_1, \ldots, u_m)$ $(m \geq 3)$ and we next perform an edge verification procedure for each edge $[u_i, u_{i+1}]$ $(i=1, \ldots, u_{m-1})$ as described above.

Complexity analysis: A profile computation that yields a profile of the form $R = (u_1, \ldots, u_m)$ with $m \geq 3$ costs $3m - 2$ probes. Since at least $m-2$ new vertices are revealed in this process, we charge each of the new vertices with at most 7 probes from the profile computation. The edge verification procedure for each of the edges $[u_i, u_{i+1}]$ charges at most 3 probes to each edge that is revealed. This proves a bound of $7f_0(P) + 2f_1(P) + f_2(P)$.

Theorem 9. Let the worst case number of line probes to determine a convex polytope P in $\mathcal{E}^3$ be $T_L(P)$. Then

$$f_0(P) + f_2(P) \leq T_L(P) \leq 7f_0(P) + 2f_1(P) + f_2(P).$$

Proof. We only have to show the lower bound. For each facet f, there is at least one line probe whose result line lies in the plane of f. For each vertex v, there is at least one line probe whose result line passes through v but does not lie in any of the planes defined by a facet incident to v. **Q.E.D.**

It ought to be remarked that the edge verification procedure is a highly unstable numerical procedure. It would be nice to give an alternative algorithm with better stability properties.

5. Probing in the presence of uncertainty

5.1. The model

The issue of probing in the presence of uncertainty is sufficiently intricate that we will restrict ourselves to finger probes in the plane. Let $\varepsilon > 0$ denote a fixed number that quantifies the uncertainty in our model. For any point x let $E_\varepsilon(x)$ denote the closed disk of radius ε centered at x. For any set X of points, define its ε-*expansion* $E_\varepsilon X$ and its ε-*interior* $I_\varepsilon X$ as follows:

$$E_\varepsilon X = \bigcup_{x \in X} E_\varepsilon(x)$$

$$I_\varepsilon X = \{ x : E_\varepsilon(x) \subseteq interior(X) \}$$

$$EI_\varepsilon X = E_\varepsilon X - I_\varepsilon X$$

Clearly $I_\varepsilon X \subseteq X \subseteq E_\varepsilon X$. Let P be a (closed) convex polygonal region. Then $I_\varepsilon P$ is an open set, and $E_\varepsilon P$ and $EI_\varepsilon P$ are closed sets. The set $I_\varepsilon P$, if non-empty, has boundary that forms a convex polygon with at most the same number of sides of P. The boundary of $E_\varepsilon P$ is decomposed into straight line segments and circular arcs as illustrated in figure 2.

We use the above sets to assign meaning to the term 'a probe has ε uncertainty' or 'ε-probes'. Such a probe is again denoted by a directed line. To distinguish probes with uncertainty from our original finger probes, we refer to the latter as 'standard probes' or '0-probes'; we may also write $C_0(F)$ and $\pi_0(F)$ for the contact point and probe path of standard probes. Out of the many possible interpretations of error probes, we choose one with the following semantics: we imagine an ε-probe F as specifying a collection of standard probes F' such that each F' is parallel to F, similarly directed and such that there is a vector $\bar{v}$ of length at most ε where $F' = F + \bar{v}$. Call this collection of standard probes the *ε-bundle* of probes defined by F. Each such F' defines the contact point $C_0(F', P)$, which corresponds to a point $C_0(F', P) - \bar{v}$ on F; this point $C_0(F', P) - \bar{v}$ is called a potential ε-contact point. Note that two different vectors $\bar{v}$ and $\bar{u}$ may give rise to the same standard probe F' but they could be distinguished in that they give rise to distinct ε-contact points. This is captured as follows.

Definition. Let F be an ε-probe and P be any polygon, not necessarily convex. A point p (possibly $p = \infty$) on F is called a *potential ε-contact point* of F at P if there exists a standard probe F' in the ε-bundle defined by F such that either $p = \infty = C_0(F', P)$ or $d(p, C_0(F', P)) \le \varepsilon$.

For any potential ε-contact point x of F, we have a corresponding probe path $\pi(F, x)$ defined analogously as for standard probes: $\pi(F, x)$ is a directed half-line that terminates at x. For simplicity (mainly in upper bound proofs), we prefer to ignore the information provided by the entire probe path $\pi(F, x)$ and simply consider the contact point x. Note that if ∞ is a possible ε-contact point for F then F does not intersect $I_\varepsilon P$; the converse is not true. In any case, let $C_\varepsilon(F, P) \subseteq F$ denote the set of finite potential ε-contact points. If P is understood, we write $C_\varepsilon(F)$ instead of $C_\varepsilon(F, P)$. If P is convex then $C_\varepsilon(F, P)$ is a connected interval of F; if P is non-convex then $C_\varepsilon(F, P)$ need not be connected.

Lemma 10. The set $C_\varepsilon(F, P)$ is equal to $F \cap E_\varepsilon X$ where

$$X = \{ C_0(F', P) : F' \textit{ is in the } \varepsilon\textit{-bundle of } F \}.$$

It follows easily that

(i) For all P, not necessarily convex, $\bigcup_F C_\varepsilon(F, P) \subseteq EI_\varepsilon P$ where the F ranges over all ε-probes.

(ii) If P is *star-shaped* then $\bigcup_F C_\varepsilon(F, P) = EI_\varepsilon P$.

In our model of computation, it is important to realize that the algorithm can only specify ε-probes F but some adversary (relative to P) will specify some ε-contact point in $C_\varepsilon(F)$. There is no assumption that the adversary will even consistently reply with the same contact point for two identical probes.

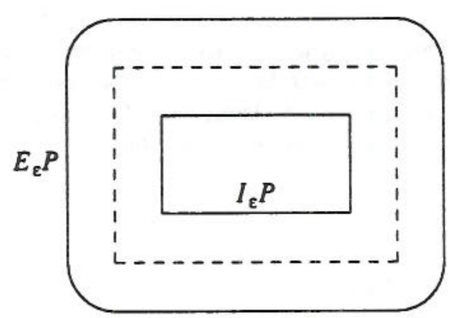

Figure 2.

However it is easy to see that the adversary can do no worse by being consistent. Thus, we may define a P*-adversary* A to be a rule which for each sequence of ε-probes $\bar{F} = (F_1, \ldots, F_k)$, $k \geq 1$, specifies a point $A(\bar{F})$ in $C_\varepsilon(F_k, P)$. We call A a *forgetful* adversary if $A(\bar{F})$ depends only on the last probe F_k. We may similarly define a *probe algorithm* to be a rule B such that for any sequence $\bar{\pi} = (\pi_1, \ldots, \pi_k)$, $k \geq 0$, of probe paths, the rule either decides that it is an 'end-game' or else specifies a probe $B(\bar{\pi})$. We usually consider probe algorithms that are determined only by the contact points on the probe paths π_i $(i = 1, \ldots, k)$.

Let B be any probe algorithm and A be any adversary. Together they determine a unique sequence of probes $F_1, F_2,.$ and a unique sequence of probe paths $\pi_1, \pi_2,.$ where for each $k \geq 1$,

$$F_k = B(\pi_1, \ldots, \pi_{k-1}), \text{ and } p_k = A(F_1, \ldots, F_k)$$

and $B(\pi_1, \ldots, \pi_k)$ is an 'end-game'. We call the sequence $(F_1, \pi_1, F_2, \ldots, F_k, \pi_k)$ a *game*.

Say a set of points X is *extremal* if

$$X \subseteq \partial(conv(X)).$$

Observe that the set of ε-contact points for P need not be extremal. Despite such difficulties, we will attempt to recover in this new setting concepts from the standard probe case. For example, we can recover the notion of collinearity by saying that three distinct points x_1, x_2, x_3 are ε*-collinear* if there is a line L that has non-empty intersection with each of $E_\varepsilon(x_i)$, $i = 1, 2, 3$.

Definition. Let $\delta > 0$. For any finite set Π of ε-probe paths, let $X(\Pi)$ be the set of finite contact points in Π. A convex polygon Q is said to δ*-fit* Π if Π is consistent with the response of a Q-adversary with δ-error. Let $H_\delta(\Pi)$ denote the set of all convex polygons that δ-fit Π. We define Π to be δ*-consistent* if $H_\delta(\Pi)$ is non-empty. If $\Pi' \subseteq \Pi$ and Π is δ-consistent then we call Π a δ*-consistent extension* of Π'.

A slight clarification may be necessary to our definition of δ-fit since Π is a set of ε-probe paths and yet we consider Q-adversaries with δ-error: this should cause no confusion since an ε-probe path is formally no different than a δ-probe path (both are specified by directed rays). A Q-adversary with δ-error is simply one that treats each ε-probe F as if it were a δ-probe, i.e., it ignores the value of ε but returns a point in $C_\delta(F,Q)$. The key definition in this section is the next idea of a 'certificate'.

Definition. Let $\delta > 0$ and Π a finite set of probe paths. We say that Π δ*-determines* a convex polygon P if for all ε-consistent extensions Π' of Π, P δ-fits Π'. If S is a finite sequence of ε-probes then we say S is a δ*-certificate* for P if the following holds: if Π is any set of probe paths defined by some P-adversary with ε-error in response to S, then Π δ-determines P.

Note that if Π δ-determines P, since Π is a consistent extension of itself, we must have $P \in H_\delta(\Pi)$. The crucial step in this development is to introduce a new parameter $\delta > 0$ that is independent of ε. There is a kind of double-standard in our definition of 'Π δ-determining P': we want P to δ-fit Π' but Π' is an ε-consistent (rather than a δ-consistent) extension of Π. A certificate is essentially a non-deterministic non-adaptive algorithm; it allows one to verify (up to δ precision) whether P has a certain shape. We first show an absolute lower limit on the precision δ of any certificate:

Theorem 11. For any P, if $\delta < 2\varepsilon$ then there does not exist any δ-certificate of P.

Proof. Suppose S is a purported δ-certificate for P. Consider the P-adversary that for any probe gives the response corresponding to the case of no errors. Let Π be the corresponding set of probe paths for S with respect to this adversary. Now choose any edge e of P and let $\bar{v}$ be the vector of length ε that points normal outward from e. Let Q be the polygon obtained by translating P a distance of ε in the direction perpendicularly outward from e, so $Q = P + \bar{v}$. We show that Π could be the response of a Q-adversary (with ε-error) to the probes in S. Let F be a probe in S and let its contact point in Π be x. Let F' be the probe obtained by translating F by the vector $\bar{v}$ (so F' is in the bundle of probes defined by F) First suppose that x is at infinity, Then F misses P implies F' misses Q, so ∞ is an appropriate response of a Q-adversary to F. If x is finite, then $x' = x + \bar{v}$ is the standard contact point of F' with Q. Since $x' \in E_\varepsilon(x)$, again x is an appropriate response of a Q-adversary to F. Finally, to show a ε-consistent extension to Π, let π be the probe path that is aimed normally at the midpoint of edge e with contact point at distance 2ε outside of P. Clearly $\Pi' = \Pi \cup \{ \pi \}$ is a ε-consistent extension of Π. But P does not δ-fit Π, contradiction. **Q.E.D.**

The intuition for the preceding lemma is that there is an ε-uncertainty associated with contact points, but there is an added ε-uncertainty arising from the desire to ensure that 'no further probes' can nullify our present guess about the shape of the polygon.

5.2. Analysis of a certificate

In contrast to the last result, we now show a positive result about certificates. In the context of standard probes, [CY87] noted that every P has a certificate of $2n$ probes, and there are no certificates with fewer probes. We need only a very mild 'largeness' requirement for P, namely:

(*) $I_{4\varepsilon}P$ is non-empty and each edge of P has length $|e| > 2\varepsilon$.

We show that with only two probes more than in the standard certificate, we can achieve a certificate of precision $\delta = O(\varepsilon)$. Our goal is to analyse the following set $S(P)$ of probes:

(i) For each corner c of P, arbitrarily pick one of the two edges e incident on c and let c' to be the point on e at distance ε away. Send a probe F_c aimed at c' such that F_c is normal to e.

(ii) For each edge e of P, send a probe F_e that is parallel to e and at distance ε+ (ε+ means any value that is infinitesimally greater than ε) from e and such that F_c misses P.

(iii) Let c be a corner with angle less than 60°. Send a probe that misses P but which intersects F_c orthogonally at a distance ε+ from c. Since there are at most two such corners, we use at most two such probes. Call these the *special* probes.

First we prove the following lemma:

Lemma 12. Let Π be any set of probe paths produced by a P-adversary in response to the probes in $S(P)$. Let $\pi = (F, p)$ be a probe path such that if $\Pi \cup \{ \pi \}$ is ε-consistent.

(a) *If F intersects $I_{(2+\sqrt{2})\varepsilon}P$ then $p \neq \infty$.*

(b) *If $p \neq \infty$ then p lies outside of $I_{(2+\sqrt{2})\varepsilon}P$.*

Proof. (a) Since $\Pi \cup \{ \pi \}$ is ε-consistent, let Q be a polygon that ε-fits $\Pi \cup \{ \pi \}$. Assume that F is horizontal and at distance $(2+\sqrt{2})\varepsilon$ below the x-axis. For the sake of contradiction, suppose $p = \infty$ and there is a point $y \in F \cap I_{(2+\sqrt{2})\varepsilon}P$. So there is a standard probe F' in the ε-bundle defined by F such that F' misses Q. Say F' lies above F. Since $E_{(2+\sqrt{2})\varepsilon}(y) \subseteq P$, there is some corner c of P that lies on or above the x-axis. Similarly, there is some corner c' of P that lies on or below the horizontal level $y = -2(2+\sqrt{2})\varepsilon$. Consider the probe path (F_c, x) in Π for some x. It is not hard to see that x is finite (since each edge of P is at least 2ε long. Furthermore, x is at a distance at most $\sqrt{2}\varepsilon$ from c. Hence x is at least 2ε above F. Since x can be the response of a Q-adversary (with ε-error) to F_c, ∂Q has some point q at a distance less than ε from x. Thus q is at least ε distance above F. A similar argument with respect to the corner c' shows that there is a point q' in ∂Q at a distance at least ε below F. This implies that the Q-adversary could not respond with $p = \infty$ for the probe F. This proves that Q does not ε-fit $\Pi \cup \{ \pi \}$, contradiction.

(b) This is similarly shown, and is omitted in this proceedings. **Q.E.D.**

We are now ready to prove the main positive result about certificates.

Theorem 13. Let P be a convex n-gon satisfying the 'largeness requirement' (*) and $\delta > 2 + \sqrt{2}\varepsilon$. Then $S(P)$ as specified above is a δ-certificate for P.

Proof. We have already observed in the above proof that each of the probes F_c gives a finite contact point. Also each F_e and also each of the two special probes yield contact points at infinity. Let Π be any set of probe paths from some P-adversary's response to $S(P)$. For each probe F that has contact point at infinity, we let H_F be the half-space bounded by F which contains P. Let Q be the convex polygon obtained by intersecting the $n + 2$ half-spaces H_F. We observe that $Q \subseteq E_{2\varepsilon}P$: it is here that we need the two special probes of $S(P)$.

Now consider any ε-consistent extension Π' of Π. We want to show that P δ-fits Π'. Take any probe path $(F, x) \in \Pi' - \Pi$. If $x = \infty$, then by the last lemma F must miss $I_\delta P$, so (F, x) is consistent with the response of a P-adversary with δ-error. So suppose x is finite. Then by the last lemma again, x is not in $I_\delta P$. It is also easy to see that x is not outside of $E_\varepsilon Q$. This implies that x is inside $E_\delta P$. Therefore $x \in EI_\delta P$. We are not yet done since $x \in EI_\delta P$ is a necessary but not sufficient condition for (F, x) to be a response of a P-adversary with δ-error: we must consider the entire probe path itself. In particular, if the probe path (i.e. the half-line of F terminating at x) intersects an edge e of P then we must be sure that $E_\delta(x) \cap e$ is not empty. This can be checked. **Q.E.D.**

By $2n + O(1)$ probes, we can improve the precision of our certificate $S(P)$ arbitrarily closer to 3ε. To do this, we first generalize the observation that there are at most two angles less than 60°:

Claim: Let $\theta = (1 - \frac{2}{k})\pi$ for some $k \geq 3$. Then every convex polygon P has at most $k-1$ angles that are $< \theta$.

[Proof of claim: If P has k angles which are $< \theta$, form the subpolygon Q with k vertices formed from these vertices of P. Each angle of Q is no larger than the corresponding angle in P. Hence we obtain the contradiction that $k\theta = (k-2)\pi$ is strictly larger than the total interior angles of Q, which is equal to $(k-2)\pi$.] Using this claim, we can now send probes to 'cut off sharp corners' analogous to our two special probes in $S(P)$. The number of 'sharp corners' is a function $k(\delta)$ of δ, and we use $O(k(\delta)) = O(1)$ special probes.

6. Concluding remarks

The study of 'crude' sensory data is relatively new in robotics. In this paper we studied the problem of determining a convex polytope using various models of probing. At first glance, it is surprising that we can even obtain finite bounds on the algorithms without any further information about these polytopes (compare (2) below). Many interesting questions remain.

(1) The section on probes with error offers many interesting prospects for further work. We have isolated some key concepts to show that some of the results on standard probes can be extended in a meaningful way. An open problem is to find efficient algorithms to determine polytopes up to δ-precision. Another intriguing question is the trade-offs between accuracy $\delta = \delta(\varepsilon)$ and the number of probes.

(2) It is natural to try to extend our results on standard probes to non-convex objects. However, *without further assumptions about these objects*, there is no meaningful extension of our results even for star-shaped objects. In other words, there is no algorithm that can determine the objects exactly in a finite number of steps. On the other hand, the problem becomes meaningful if we make additional assumptions on the star-shaped objects similar to (*) in section 5.2. Similarly, the extension to multiple convex polygons can only work with additional information.

(3) The approach to vision known as the 'model-based vision' has been extensively investigated, especially by researchers at Stanford. Abstractly, this is essentially the problem of preprocessing information about the objects or scenes to be identified. Suppose that, given a polygon P, we want to know whether P is a translation and rotation of some convex polygon taken from a given finite set S. The upper and lower bounds of Yap and Cole no longer holds. Indeed, H. Bernstein [Ber86] has observed that $2n + O(1)$ probes suffice (the necessity of this bound is easy).

(4) The probing strategies reported in this paper rely on substantial, although polynomial time, computations which process the results of the past probes and then figures out the next probe. It would be interesting to analyze the complexity of these background computations.

REFERENCES

[Ber86]. Herbert J. Bernstein,"Determining the shape of a convex n-sided polygon by using $2n+k$ tactile probes," *Information Processing Letters*, **22**, pp. 255-260 (1986).

[CY87]. Cole and C. Yap, "Shape from probing", *Journal of Algorithms*, **8**, pp. 19-38 (1987).