SIAM J. COMPUT.
Vol. 17, No. 5, October 1988

002

PROBING CONVEX POLYGONS WITH X-RAYS*

HERBERT EDELSBRUNNER†‡ AND STEVEN S. SKIENA†

Abstract. An X-ray probe through a polygon measures the length of intersection between a line and the polygon. This paper considers the properties of various classes of X-ray probes, and shows how they interact to give finite strategies for completely describing convex n-gons. It is shown that $(3n/2)+6$ probes are sufficient to verify a specified n-gon, while for determining convex polygons $(3n-1)/2$ X-ray probes are necessary and $5n+O(1)$ sufficient, with $3n+O(1)$ sufficient given that a lower bound on the size of the smallest edge of P is known.

Key words. theory of robotics, computational geometry, probing, X-rays, convexity, complexity

AMS(MOS) subject classifications. 52A10, 68C25

1. Introduction. Identifying and understanding objects from sensory data is a fundamental problem in robotics and computer vision. Such sensors can include imaging devices as well as simpler detectors. Clearly imaging devices provide a tremendous amount of information; however for reasons of economy and robustness, simple sensors are often used. There has been increasing interest [1]–[3] in analyzing the number of measurements which different types of sensors need to determine convex polytopes. This paper introduces another sensor model, the *X-ray* probe, and gives strategies for using it.

Probing polytopes can be viewed as a discrete case of sampling problems encountered in signal processing. The *Nyquist rate* [4] specifies the amount of sampling needed to reconstruct a continuous waveform. Since our objects of interest, convex polygons, have much more structure than continuous waveforms, it is clear that tighter bounds can be obtained. Continuous waveforms generated by probes analogous to our model are used in such medical instrumentation systems as tomography. The techniques used in reconstruction of these waveforms are unrelated to ours, (see [5]–[7] for surveys).

The most studied geometric probe is the *finger probe* [1]–[3], which is a directed line l and returns the first intersection of l with polygon P. Intuitively, it is like moving a finger towards an object and recording where it hits. An *X-ray probe* $X(P, l)$ is defined to be the length of intersection between polygon P and the line l. We assume that we know a point O within P, which identifies the general location of P in the plane. Without such a point to provide a general idea of where P is, it is not clear how to find P in a finite number of probes. For convenience we shall assume O is within the interior of P.

A collection of X-ray probes through an object provides us with a great deal of information about it but not directly with the coordinates of a point on the surface. Obtaining such absolute information is the difficulty in working with X-ray probes. Figure 1 demonstrates some of these difficulties; the collection of probes provides very little constraint on the location, shape, or number of sides of P. Another polygon P' bears little resemblance to P, but gives identical results for the collection of probes. In fact, the complete set of probes $X(P, l(O, \theta))$ through O over all angles θ describes two different polygons, P and P reflected through O, denoted as $-P$. Thus it is a

* Received by the editors December 15, 1986; accepted for publication (in revised form) October 29, 1987.
† Department of Computer Science, University of Illinois at Urbana-Champaign, Urbana, Illinois 61801.
‡ The research of this author was supported by the Amoco Foundation Facility for the Development of Computer Science 1-6-44862.

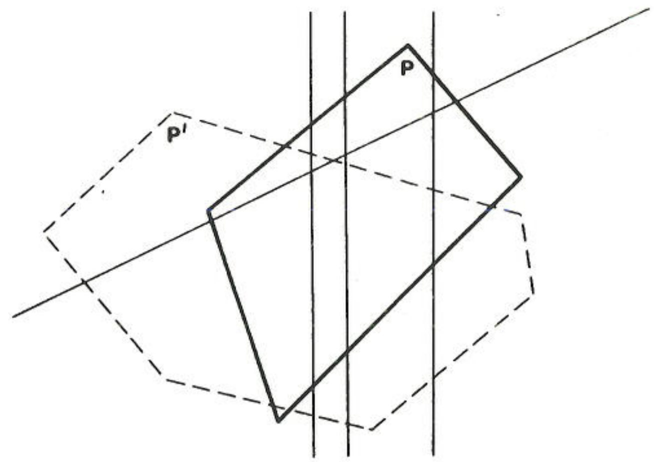

FIG. 1. *Different polygons satisfying the same collection of probes.*

fundamentally different type of device from the finger probe, although it is not clear which can be considered more powerful.

This paper will give upper and lower bounds on the number of X-ray probes necessary to determine convex n-gons, as well as for the problem of verifying a conjectured convex polygon. These bounds are based on the powers of different classes of probes, each of which has capabilities and limitations to be examined.

2. A lower bound for probing. We can prove a nontrivial lower bound for X-ray probing by a comparison to finger probing.

LEMMA 1. *An X-ray probe $X(P, l)$ can be simulated by two finger probes.*

Proof. Send finger probes down each end of the line, and compute and return the Euclidean distance between these two points. □

THEOREM 2. *At least* $(3n-1)/2$ *X-ray probes are necessary to determine a convex n-gon.*

Proof. Cole and Yap [1] prove a lower bound of $3n-1$ finger probes for determining convex polygons. The result follows from Lemma 1. □

This bound is probably loose, but the concept of a lower bound on the number of probes is complicated by difficulty in determining exactly what information can be obtained in a constant number of X-ray probes. Since their power comes from collections of probes, a tighter lower bound may be difficult to obtain.

3. Upper bounds for probing. To obtain absolute information about P from X-ray probes, it is necessary to think in terms of groups of probes which work together to determine something about P. This section considers four different classes of probes, what powers and limitations they possess and how they interact to lead to probing strategies.

3.1. Origin probes. The first class of probes are *origin probes*, a set of X-ray probes all aimed through a common point O within the object. An X-ray probe which hits a convex polygon and avoids its vertices will go through exactly two edges of the object. The largest number of such edge pairs is n.

LEMMA 3. *For any convex n-gon P containing O, there are at most n distinct pairs of edges e_1, e_2 of P such that there exists a line which intersects e_1, e_2 and O.*

Proof. Consider a line l through O which intersects a pair of edges. If we rotate this line clockwise, only when we sweep past a vertex do we intersect a new pair of edges. After rotating past n vertices and at most π radians, we will return to the original pair. No other distinct pairs can be found, so there exist at most n opposing edges. Note that there are exactly n edge pairs unless O is collinear with two vertices of the n-gon. □

We can define a mapping of $X(P, l(O, \theta))$ to two points p_1 and p_2 on $l(O, \theta)$ at a distance $X(P, l(O, \theta))$ from O, where $l(O, \theta)$ is the line through O that encloses an angle of θ radians with the x-axis. By the following result, these points lie on hyperbolas defined by the edges probed through.

LEMMA 4. *Let $l_1: y = m_1x + b_1$ and $l_2: y = m_2x + b_2$ be two distinct lines and map each angle θ to the two points on line $l(O, \theta)$ at distance d from O, where d is the distance between $l_1 \cap l(O, \theta)$ and $l_2 \cap l(O, \theta)$. Then these points define the hyperbolas*

$$y^2 - xy(m_1 + m_2) \pm y(b_1 - b_2) - x^2(m_1m_2) \pm x(m_1b_2 - m_2b_1) = 0.$$

Proof. Consider the situation in Fig. 2. The line $l(O, \theta)$ contains points (x, y) such that $y = \tan(\theta)x$. With $t = \tan(\theta)$, we have

$$d = \sqrt{(1 + t^2)(b_2/(t - m_2) - b_1/(t - m_1))^2}.$$

The x-coordinates of the corresponding points can be found by subtracting the x-coordinates x_1 and x_2 of the intersections of $l(O, \theta)$ with l_1 and with l_2, respectively:

$$x = \pm(x_2 - x_1) = \pm(b_2/(t - m_2) - b_1/(t - m_1)).$$

This can be solved for t and used with $y = tx$ to get y as a function of x. To obtain a simpler formula, we set $A = m_1 + m_2$, $B = m_1m_2$, $C = b_1 - b_2$, and $D = m_1b_2 - m_2b_1$. Then we have

$$y = \frac{Ax \mp C + \sqrt{(Ax \mp C)^2 - (4Bx^2 \pm 4Dx)}}{2}$$

and

$$y = \frac{Ax \mp C - \sqrt{(Ax \mp C)^2 - (4Bx^2 \pm 4Dx)}}{2}.$$

After simplification, this reduces to the assertion. □

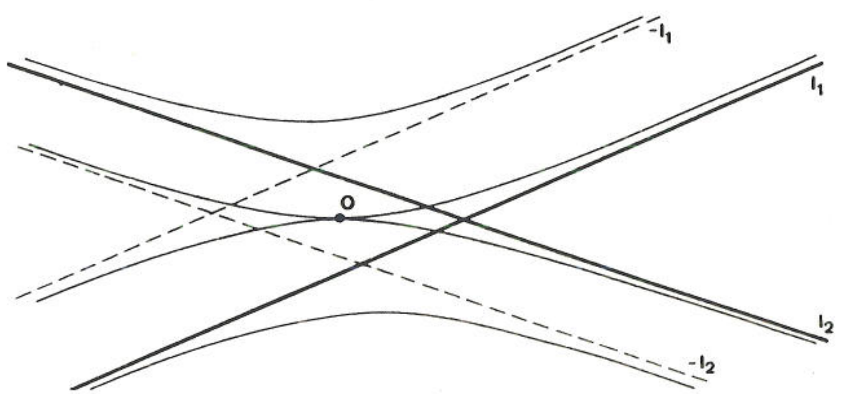

FIG. 2. *Hyperbolas associated with two straight lines.*

We will use Lemma 4 to determine the equations of the lines that contain edge pairs. If we have some number of origin probes through a common edge pair, then we can determine the hyperbolas through the associated points. From the equations of these hyperbolas, we then deduce the equations of the lines.

It is obvious from Fig. 2 that the hyperbolas which satisfy Lemma 4 have asymptotes l_1 and $-l_2$ and l_2 and $-l_1$. If the angle θ is such that $l(O, \theta)$ intersects l_1 and l_2 on opposite sides of O, then the two hyperbolas that avoid O are relevant. This will be the most frequent case in our discussion below. If $l(O, \theta)$ intersects l_1 and l_2 on the same side of O, then the two hyperbolas through O are relevant. This occurs when O is not between the two edges of the pair considered. In both cases, the two relevant hyperbolas are central reflections of each other.

The hyperbolas that show up in Lemma 4 are defined by four constants $A = m_1 + m_2$, $B = m_1m_2$, $C = b_1 - b_2$, and $D = m_1b_2 - m_2b_1$. It follows that, in general, four probes through a pair of edges is enough to determine the hyperbolas, and from the hyperbolas the equations of the lines that contain the two edges. Given A, B, C, D, we have

$$m_1 = \frac{A \pm \sqrt{A^2 - 4B}}{2}, \qquad m_2 = \frac{A \mp \sqrt{A^2 - 4B}}{2},$$

$$b_1 = \frac{C}{2} \pm \frac{2D + AC}{2\sqrt{A^2 - 4B}}, \quad \text{and} \quad b_2 = \frac{-C}{2} \pm \frac{2D + AC}{2\sqrt{A^2 - 4B}}.$$

From these equations two limitations on our ability to reconstruct the edges are apparent. First, there is the ambiguity between P and $-P$. More serious is that b_1 and b_2 are undefined when the square root vanishes, that is, when $m_1 = m_2$. Thus any probing strategy using origin probes must take special action on parallel edges. Note, however, that there is no ambiguity between P and $-P$ for parallel edges, since reflection through O is equivalent to rotation through π radians.

We can now recognize the structure formed by mapping origin probes to points. Each opposing pair of edges gives rise to their own pair of hyperbolas. Two adjacent hyperbolas meet on the line through O and each vertex. Thus the probes define the extremes of a "spider web" $S(P)$ (see Fig. 3) around the object. This permits us to interpret an origin probe for P to be a finger probe on $S(P)$. Both P and $-P$ generate the same spider web, $S(P) = S(-P)$. By Lemma 3, $S(P)$ consists of at most $2n$ pieces of hyperbolas.

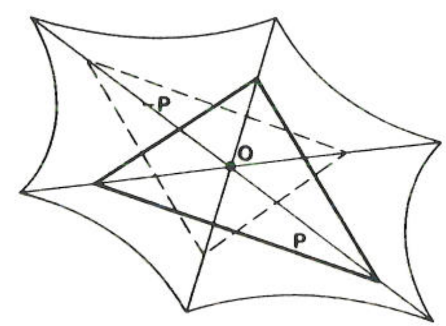

FIG. 3. *The "spider web" $S(P)$ around a convex polygon P.*

However, to make use of these properties of origin probes for a probing strategy we need some means to verify that four or possibly some higher but constant number of origin probes pass through the same pair of edges. With finger probes, it was possible to confirm an edge with three probes, two to define it and one in between to verify it. Unfortunately, a constant number of extra confirmation probes lying on the same hyperbola is not sufficient to verify an edge pair.

LEMMA 5. *There is no constant k such that k probes lying on a hyperbola defined above implies that the probes pass through the same pair of edges of P.*

Proof. The proof is by construction. Figure 4 indicates such a polygon. A convex chain of edges can be used to insure that as many points as desired of $S(P)$ can be made to lie on a hyperbola of the kind defined in Lemma 4. Let one edge e lie on the line $y = x + 1$, and let h be the line of points (x, y) that satisfy $y = -2$. We construct a curve c of points p with the following property: if l is a line through p and O, then l intersects edge e in point q and line h in point r such that the distance from p to q is the same as the distance from O to r.

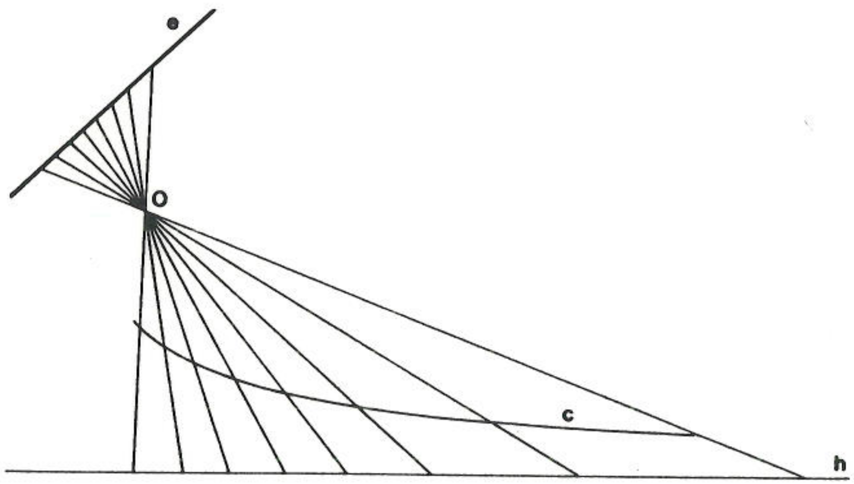

FIG. 4. *Counterexample to the notion that k origin probes mapped onto the same hyperbola must intersect the same edge pair.*

It is not hard to see that c is a piece of a hyperbola whose asymptotes are line h and the line that contains e. If we pick k vertices of P on curve c, then all origin probes through these vertices are mapped to points on line h. Thus, if a collection of probes were made through O and these vertices, they would all be taken as lying on a parallel pair of edges.

By construction, $S(P)$ has k vertices on line h and the hyperbolic pieces connecting any two consecutive vertices lie all above line h. Thus, a hyperbola of the kind considered in Lemma 4 can be found that intersects $S(P)$ at least $2k$ times if it lies close enough and above h, where the vertices of $S(P)$ on h lie. □

Verifying edge pairs is the motivation for parallel probes, discussed below.

3.2. Parallel probes. *Parallel probes* are a set of X-ray probes aimed with a common angle θ. A complete collection of parallel probes for a given angle produce a histogram $H(P, \theta)$ (see Fig. 5) of the thickness of the obstacle. This is the situation in a medical

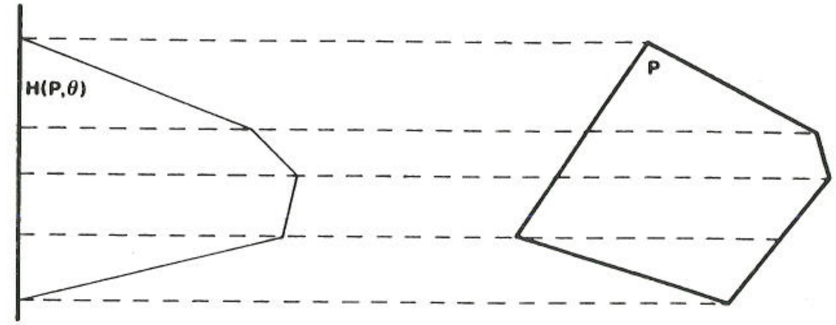

FIG. 5. *The "histogram" $H(P, \theta)$, $\theta = 0$, of a convex polygon P.*

X-ray photograph. Formally, $H(P, \theta)$ is obtained as follows. Let b, the baseline, be normal to the probing direction θ and call one of the half-planes defined by b its *positive side.* Let $l(b, p)$ be the line normal to b such that $p = b \cap l(b, p)$; thus, $l(b, p)$ has angle θ. We map the probe $X(P, l(b, p))$ into the point on $l(b, p)$ at distance $X(P, l(b, p))$ from b that lies in the positive side of b. $H(P, \theta)$ is the polygon bounded by b and the images of all probes with angle θ. For a convex n-gon this histogram will be bounded by up to n line segments, including the one on b. $H(P, \theta)$ is convex for all convex P over all angles θ. Note that an X-ray probe with angle θ can be interpreted as a finger probe on $H(P, \theta)$. Each vertex of $H(P, \theta)$ determines a line on which must lie a vertex of P. Thus they provide a capability for verifying edge pairs.

LEMMA 6. *Three parallel probes are sufficient to verify an edge pair.*

Proof. Parallel probes measure the distance between two line segments. Thus three parallel probes are mapped to three collinear points on the boundary of $H(P, \theta)$ if they intersect the same two segments. No three parallel probes hitting different edge pairs can be mapped to collinear points without violating convexity. □

There are two apparent weaknesses of parallel probes. First, a finite number of them cannot be guaranteed to locate certain vertices, specifically the extreme vertices that correspond to the vertices of $H(P, \theta)$ that lie on the baseline b. Without any bounding information, repeated probes may intersect the same pair of edges. Second, once a vertex is finally located, it is impossible without more information to distinguish whether there are one or two vertices on the line. Convexity restricts the number of collinear vertices to two. This phenomenon relates to the fact that $H(P, \theta)$ neither determines P nor the number of its vertices.

It is interesting to consider using an infinite number of parallel probes, as approximated in medical X-ray photographs where all probes perpendicular to the photographic plate are recorded at once. The first weakness vanishes although the second remains. This problem was first posed for convex sets by Hammer in 1963 [8] and has generated a substantial literature [9]–[12]. The power of such a probe which returns $H(P, \theta)$ for a specified θ is evident in that only a constant number of *photograph probes* are needed to determine any convex polygon.

THEOREM 7. *Three X-ray photograph probes are sufficient to determine a convex polygon P.*

Proof. The strategy is as follows. Each photograph probe can be defined by its baseline. For the first two probes, use baselines b_1 and b_2 that are not parallel to each other. Since according to the observation above, each of these probes defines up to n lines on which all vertices of P must lie, two intersecting probes will define up to n^2 points at the intersection of these lines. The vertices of P must be a subset of these points. The third baseline b_3 is selected such that the line through each of these points perpendicular to b_3 is unique. The n points on this histogram uniquely identify each of the vertices of P. □

Because an X-ray photograph probe captures in one probe a representation of the entire polygon, they may provide a way to extend probing results to nonconvex polygons by eliminating the need for an infinite set of probes to hunt for possible concavities.

3.3. Determining a boundary point. Since parallel and origin probes have complementary properties of verification and identification, by combining them we can obtain our first piece of absolute information about the polygon. The idea behind this strategy is to bound a section of the polygon where we know there must be at least one edge pair and to repeatedly send parallel probes to this section until three images on the boundary of the corresponding histogram are collinear. Once we have an edge pair, origin probes can be used to determine the lines containing the two edges.

First, a section of the polygon must be bounded. Sending one horizontal probe through O gives a thickness λ of P along that x-axis. At least one of two vertical probes at distance $\lambda/2$ to the left and to the right of O must intersect P and with a vertical probe l_O through O define a section which is a vertical strip such that each vertical probe in this strip intersects P. First, we discuss certain subtleties of the determined section. If both vertical probes l_1 and l_2 at distance $\lambda/2$ from O intersect P, as in Fig. 6, we cannot be sure for which side of O the x-axis is within P, and thus we cannot be certain about any new interior points. In this case, we choose our section to be either to the left or to the right of O.

To find a first edge pair, we send vertical probes in the identified section until three points on the boundary of the associated histogram $H(P, \pi/2)$ are collinear. All points lie on at most $n-1$ edges of $H(P, \pi/2)$ and all but one edge contains at most two of the points. It follows that we succeed after at most $2n-3$ additional vertical probes, since we already have two vertical probes which bounded the probed section

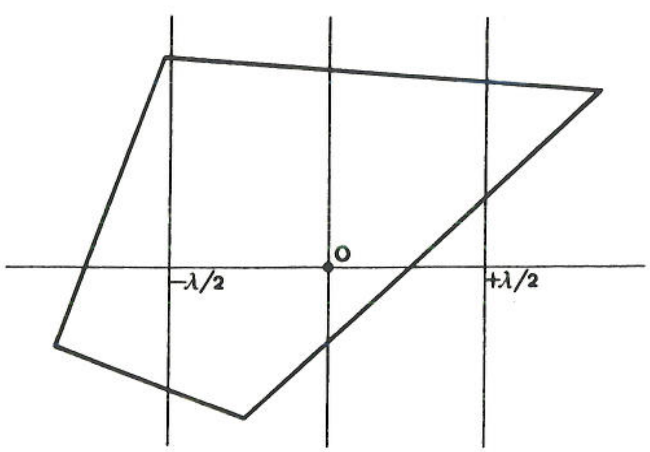

FIG. 6. *Identifying a section to parallel probe.*

of P. Once we have vertical probes through an edge pair, four origin probes through a common point, called the origin, can be used to determine the lines that contain the two edges. By picking the origin for this collection of probes to be where the center parallel probe intersects the x-axis, we need only make three new probes. One final probe can distinguish between the edges of P and those of $-P$.

In order to origin probe the edge pair between the section defined by two vertical probes, we must have an upper bound U on the distance the edges within this section are from the x-axis. Without this knowledge, we cannot design origin probes which we can be certain will intersect the boundary of P within the section, and thus whether the origin probes all hit the same pair of edges. For the situation where exactly one of the two initial vertical probes l_1, l_2 intersects P, we know that the x-axis intersects P throughout the resulting section. In this case we can bound U because it is clear that the distance of a point from the x-axis is no larger than the height of the associated histogram at the vertical line through the point.

This argument fails in the case of Fig. 6, since we cannot be certain on which side P contains the x-axis. We will use a convexity argument to put a bound on U. Let us arbitrarily select the section to the right of O. If this section contains the x-axis, we know a bound on U. If not, we know that the other section contains the x-axis within P for a distance $\lambda/2$. The slope of the upper edge of P that intersects the y-axis is greatest if it is the only upper edge of the left section, $l_1 = l((-\lambda/2, 0), \pi/2)$ intersects P entirely below the axis and $l_O = l(O, \pi/2)$ intersects P entirely above the axis. By convexity, no edge in the other section can increase faster than this line, which passes through the points $(0, X(P, l_O))$ and $(-\lambda/2, -X(P, l_1))$. Reflecting this situation along the x-axis bounds the lower edges, and together provides the information we need to origin probe.

A further complication occurs when the edge pair is parallel, meaning b_1 and b_2 are undefined. If another potential edge pair exists in this section, that is, it required more than five parallel probes within the section to locate the parallel edge pair, this nonparallel edge pair can be uncovered by a total of $2n$ parallel probes, since now two edge pairs can have three probes each.

If the first edge pair is parallel and another edge pair is not known, we must now repeat the sectioning process parallel to the original pair. Clearly, a line l_p through O parallel to the first edge pair intersects P between the edge pair. By a process of binary search, we can enlarge this strip of P known to lie between the edge pair as much as desired. If δ_1 is the distance between the two edges, a probe parallel to l_p $\delta_1/2$ below l_p widens this strip by $\delta_1/2$. If this probe intersects P, the strip is between l_p and the last probe. Otherwise, the strip is on the other side of l_p. Similarly, we can widen this known strip to $3\delta_1/4$ with a probe parallel to l_p $\delta_1/4$ on either side of the known strip. We can continue to widen this strip by this method, although for our purposes two of these probes will suffice. This strip will serve to define a section for the next set of probes. Note that there is no reason to actually probe l_p and that at this point we do not know how long the edges of the first parallel pair are.

Since these probes are parallel with our previously encountered edge pair, they will intersect at least two edges different from the parallel edge pair. We will make five of these, one at each side of the boundary of our $3\delta_1/4$ region, two more within this region $\delta_1/2+\varepsilon$ apart for $0<\varepsilon<\delta_1/4$, and one between these final two probes. Note that it may be possible to reuse the binary search probes, but only if they intersected P. If the center three of these probes are not all of the same magnitude, they do not all intersect a parallel pair of edges. Thus with up to $2n-3$ more parallel probes we can locate a nonparallel edge pair, which can be origin probed to determine the edge pair. Unfortunately, as in Fig. 7, the center three probes can instead intersect a parallel edge pair. This parallel edge pair must be greater than $\delta_1/2$ in length.

If δ_2 is the distance between the second pair of parallel edges, we can define a strip $3\delta_2/4$ wide within P through the binary search strategy using probes parallel to the second edge pair. Two parallel probes within this strip $\delta_2/2+\varepsilon$ apart for $0<\varepsilon<\delta_2/4$ can confirm that the first edge pair is greater than $\delta_2/2$ in length. If this is not the case, we can find a nonparallel edge pair between the offending probe and the appropriate end of the $3\delta_2/4$ region. Otherwise, the intersection of the two strips defines a rhombus Q which must lie within the interior of P. We note that the remainder of P must lie in strips less than $\delta_1/2$ and less than $\delta_2/2$ wide around Q. Extending these boundaries for each of the two edge pairs surrounds Q by a skewed grid of eight regions, which together contain all of the edges of P. None of these regions can contain parallel edges, since P is convex. Further, no three neighboring regions around Q including only one corner region contain any parallel edges.

Referring to Fig. 7, it is clear that a probe X from the upper left corner of the top-central region to the bottom right corner of the right-central region cannot intersect a parallel pair of edges. Because of the size and position of the central region, this probe must intersect Q, meaning it intersects a nonparallel edge pair. Along with a probe parallel to X that intersects the upper right-hand corner of Q this defines a section which can only contain nonparallel edge pairs, and thus can be parallel probed until three are collinear on the histogram. Using the earlier counting argument with $n-2$ edges, since the other two parallel edges cannot be within the section, shows finding an edge pair can require up to $2n-5$ additional probes.

Finally, one confirmation probe of the nonparallel edge pair will distinguish the edges on P from the pair on $-P$. Note that there is no ambiguity beween P and $-P$ for the parallel edge pairs.

LEMMA 8. *With restriction to origin and parallel probes,* $2n+23$ *X-ray probes are sufficient to identify the lines that contain the first pair of edges and a point on one of the two edges.*

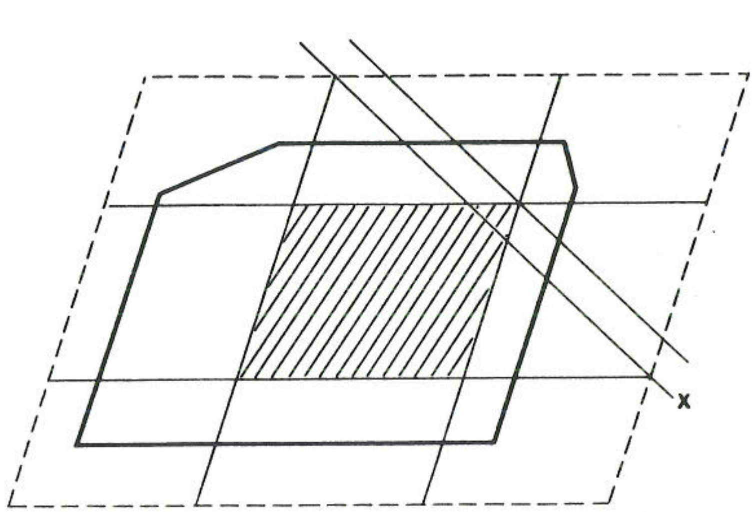

FIG. 7. *Handling parallel edge pairs.*

Proof. The above strategy will identify the lines that contain a pair of edges and a point on one of the two edges. The final accounting of probes is as follows. Four probes were used initially to define a section to probe, at least two of which can serve as parallel probes. Three more parallel probes can identify an edge pair, with the center three probes incident on a parallel pair of edges. Three origin probes found the slopes of these lines. Two parallel probes will widen the strip to $3\delta_1/4$. Up to five probes between the two edges will identify that this edge pair is also parallel, and three more are necessary to origin probe it. Two more probes widen the new strip, two more enlarge Q, and two diagonal probes select a nonparallel section. $2n-5$ parallel probes in this section will locate a nonparallel edge pair, the first two of which were the diagonal probes. Three more origin probes find the equations of these lines. Finally, there is the confirmation probe. Thus two edges can be determined in a total of

$$4+3+3+2+5+3+2+2+2+(2n-5-2)+3+1=2n+23$$

probes. Any point on these two edges within the appropriate section is on the boundary of P. □

A complete probing strategy for P could perhaps be constructed along these lines by repeating the process for each pair of edges. However, since $O(n)$ edge pairs can be parallel this would lead to a quadratic number of probes. A simpler strategy can be developed once we know a point on the boundary of P.

3.4. Boundary probes. The power of the finger probe is that it returns a point on the boundary of the polygon. To get a similar effect, we define the *boundary probe*, which relies on the observation that sending an X-ray line probe through a known point on the boundary of a convex polygon identifies another point on the boundary of the polygon. This means that once we have identified the coordinates of any point p on the boundary of the polygon, any X-ray probe through p determines another boundary point. If we also are given a boundary point we can formulate our first probing algorithm.

THEOREM 9. *With restriction to origin, boundary, and parallel probes,* $5n+19$ *probes suffice to determine a convex n-gon P.*

Proof. By Lemma 8, $2n+23$ probes suffice to find a boundary point and to *semi-verify* two edges, in the terminology of Cole and Yap [1]. Cole and Yap give a strategy requiring $3n$ finger probes to determine polygons which can be adapted by using boundary probes in place of finger probes.

Starting from one of the semi-verified edges, we will walk along the polygon, conjecturing vertices based on the intersection of the semi-verified edge and the line defined by the next two known points. Each of the n vertices will eventually be probed, and each of the $n-2$ other edges will have at most two interior points probed, for a total of $5n+19$. □

Being clever about reinterpreting the parallel probes may reduce the total by $O(n)$ more probes since once the edge one of them passed through is determined, a probe on an unverified edge can be recorded. No doubt, the additive constant of Lemma 8 can be lowered by more careful arguments.

Note that Cole and Yap's optimal $3n-1$ strategy cannot be adapted to X-ray probes since they probe along a semi-verified edge to obtain a vertex, which will not work with X-ray probes unless the location of the other vertex is known.

3.5. Close probes. If the measurements we have been using were performed on real sensing devices, there will be some uncertainty with respect to accuracy. Thus for us to completely determine an n-gon we must know that no edge has length less than this uncertainty, or else we could never find the edge. Knowing a lower bound on the length of all the edges ε gives us extra information about the polygon. We can exploit this with a collection of *close probes*, where each probe depends on intersecting a point on the boundary within some fraction of ε of another close probe.

Close probes present a problem because they suggest strategies that are somehow "unfair" as they require additional information. Certainly in any physical implementation they would be extremely nonrobust. The virtue of close probes is that they enable us to find a boundary point in a constant number of X-ray probes, as opposed to the linear probing strategy described above.

LEMMA 10. *Two lines that contain a pair of edges of a convex polygon P and a point on one of the two edges can be determined in* 37 *X-ray probes including close probes.*

Proof. Our strategy is similar to that used in Lemma 8, but modified to take advantage of close probes. It should be noted that there is no fraction α such that parallel probes within $\alpha\varepsilon$ of each other are guaranteed to intersect the same edge pair. The reason is that the angle between two edges can be arbitrarily close to π, so even a long edge can slip between two seemingly close probes. Thus the largest angle between edges will have to be bounded to take advantage of close probes.

We will replace the linear strategy of searching a bounded section of P for an edge pair by the following constant one. Within the bounded section of P, send two more parallel probes, giving four probes intersecting P labeled from left to right a, b, c, and d. By the argument in the proof of Lemma 8, a and b determine the steepest increasing slope possible between b and c, and probes c and d determine the steepest decreasing slope. We define θ_v as the greater of the two angles formed by these steepest increasing and decreasing slopes with the horizontal, so θ_v represents the angle nearest to vertical which can occur within the section without violating convexity. An edge pair will be found within seven close probes spaced $\varepsilon \cos(\theta_v)/8$ apart between b and c. Seven are required because up to two vertices, one each from the upper and lower vertex chains of P, may slip between the close probes.

Thus we can use the strategy of Lemma 8, substituting the two parallel and seven close probes for the linear edge pair search. Using the counting argument of Lemma 8 with this change, we can determine the first edge pair in 37 probes. □

THEOREM 11. $3n+33$ *X-ray probes* (*including close probes*) *are sufficient to determine a convex polygon P given a point O within P.*

Proof. Lemma 10 enables us to find an edge pair in 37 probes instead of the $2n+23$ of Lemma 8. Substituting the new strategy for the old improves Theorem 9 by $2n-14$ probes, for a total of $3n+33$. □

It is certainly possible that these constants can be reduced. Other strategies involving close probes are no doubt possible.

It would be nice to find an efficient X-ray probe simulation of a finger probe. It is possible by modifying the above strategy as follows. Make one of the parallel close probes along the desired probing line and then, if it hits on a parallel edge, perform additional boundary probes through the located point to finish the description of the edge pair, and calculate what the finger probe returned. However, since this constant will be over twenty it is unlikely the simulation can prove useful in any context.

4. Bounds for verification. A lower bound on the number of probes required to determine an object can be based on a comparison to the *verification* problem. Suppose we are given the representation of a polygon P. How many probes will be necessary to test whether P correctly describes a particular object? It is obvious that any lower bound to verification represents a lower bound to the determination problem, since it presupposes knowledge of the polygon.

For verification, clearly each vertex and edge must be confirmed. Otherwise, P could have a triangle on any unconfirmed edge or be truncated before any unconfirmed vertex. Since an X-ray probe passes through two edges or vertices, and there are at least $2n$ points of interest, the trivial lower bound is for n probes. It can be easily shown that three probes do not suffice to verify a triangle, since no matter how the three probes are taken the object would be indistinguishable with one of four or more sides. We conjecture that the actual bound for verification is $(3n/2)+k$ for some constant k. This is based on the observation that although $n/2$ probes are sufficient to verify the edges given the vertices or verify the vertices given the edges, it appears at least n probes are necessary to verify either the vertices or the edges independently. This conjectured lower bound is sufficient.

THEOREM 12. $(3n/2)+6$ *X-ray probes are sufficient to verify a convex n-gon P.*

Proof. Given the polygon to verify, three parallel probes are sufficient to verify the existence of a nonparallel edge pair and three origin probes are enough to define the hyperbola of it. One final probe to verify that P is not reflected through O identifies a boundary point.

From this boundary point, n boundary probes can verify the vertices. The remaining $n-2$ edges can be verified with $(n-2)/2$ probes, each bisecting a different pair of edges. Since P is the convex hull of its vertices, none of these probes can have length other than expected without violating convexity, unless there exists another vertex. □

5. Open problems and extensions. We have given strategies for probing with X-rays. In particular, we have shown that complete information about a convex n-gon can be obtained with a linear number of carefully planned X-ray probes. Still, the power of X-ray probes is not well understood. For example, no algorithm is known that decides whether or not a given collection of X-ray probes (and answers) determines the probed object.

An obvious extension of our results would be to three or more dimensions. Since the boundary probe generalizes to a finger probe in three or more dimensions, and both parallel and origin probes can identify edge pairs from slices of polytopes, a higher-dimensional probing strategy can be based on ideas in Dobkin, Edelsbrunner, and Yap [3].

A different type of probe to consider would measure the area or volume of intersection with a half-plane or half-space instead of a line. Such an "Archemedian" probe in two dimensions would have as its derivative an X-ray probe. In three dimensions, its derivative is a cross-sectional area probe. We have proved linear upper and lower bounds for determination with half-plane probes [13]. An interesting question is whether better probing strategies will result by having access to two different devices, such as finger and X-ray probing.

It is also worth considering the computational complexity of determining the location of the probes. Finally, a further study of X-ray photograph probes may generalize some of these results to simple polygons and provide some insight into standard image-reconstruction methods.

6. Acknowledgments. We thank Vasant Rao for motivating the problem and for his discussion of results from the field of signal processing. We also thank Raimund Seidel for suggesting the construction used in the proof of Lemma 5.

Note added in proof. A $2n$ lower bound for determination can be shown by a topological argument. See [14] for details.

REFERENCES

[1] R. COLE AND C. K. YAP, *Shape from probing*, J. Algorithms, 8 (1987), pp. 19–38.
[2] H. J. BERNSTEIN, *Determining the shape of a convex n-sided polygon by using 2n+k tactile probes*, Inform. Process. Lett., 22 (1986), pp. 255–260.
[3] D. P. DOBKIN, H. EDELSBRUNNER, AND C. K. YAP, *Probing convex polytopes*, Proc. 18th Annual ACM Symposium on Theory Computing, 1986, pp. 424–432.
[4] A. OPPENHEIM AND R. SCHAFER, *Digital Signal Processing*, Prentice-Hall, Englewood Cliffs, NJ, 1975.
[5] G. T. HERMAN, *Image Reconstruction from Projections: the Fundamentals of Computerized Tomography*, Academic Press, New York, 1980.
[6] L. A. SHEPP AND J. B. KRUSKAL, *Computerized tomography: the new medical X-ray technology*, Amer. Math. Monthly, 85 (1978), pp. 420–439.
[7] K. T. SMITH, D. C. SOLOMON, AND S. L. WAGNER, *Practical and mathematical aspects of the problem of reconstructing objects from radiographs*, Bull. Amer. Math. Soc., 33 (1977), pp. 1227–1270.
[8] P. C. HAMMER, *Problem* 2, in Proc. Symposium on Pure Math Vol. VII: Convexity, American Mathematical Society, Providence, RI, 1963, pp. 498–499.
[9] R. J. GARDNER AND P. MCMULLEN, *On Hammer's X-ray problem*, J. London Math. Soc., 21 (1980), pp. 171–175.
[10] K. J. FALCONER, *X-ray problems for point sources*, Proc. London Math. Soc., 46 (1983), pp. 241–262.
[11] ———, *Hammer's X-ray problem and the stable manifold theorem*, J. London Math. Soc., 28 (1983), pp. 149–160.
[12] R. J. GARDNER, *Symmetrals and X-rays of planar convex bodies*, Arch. Math., 41 (1983), pp. 183–189.
[13] S. S. SKIENA, *Probing convex polygons with half-planes*, Dept. of Computer Science, Report UIUCDCS-R-1380, University of Illinois at Urbana-Champaign, Urbana, IL, 1987.
[14] S. S. SKIENA, *Geometric Probing*, Ph.D. dissertation, Department of Computer Science, University of Illinois at Urbana-Champaign, Urbana, IL, 1988.

SIAM J. COMPUT.
Vol. 17, No. 5, October 1988

003

DEFERRED DATA STRUCTURING*

RICHARD M. KARP†, RAJEEV MOTWANI†, AND PRABHAKAR RAGHAVAN‡

Abstract. We consider the problem of answering a series of on-line queries on a static data set. The conventional approach to such problems involves a preprocessing phase which constructs a data structure with good search behavior. The data structure representing the data set then remains fixed throughout the processing of the queries. Our approach involves dynamic or query-driven structuring of the data set; our algorithm processes the data set only when doing so is required for answering a query. A data structure constructed progressively in this fashion is called a *deferred data structure*.

We develop the notion of deferred data structures by solving the problem of answering membership queries on an ordered set. We obtain a randomized algorithm which achieves asymptotically optimal performance with high probability. We then present optimal deferred data structures for the following problems in the plane: testing convex-hull membership, half-plane intersection queries and fixed-constraint multi-objective linear programming. We also apply the deferred data structuring technique to multi-dimensional dominance query problems.

Key words. data structure, preprocessing, query processing, lower bound, randomized algorithm, computational geometry, convex hull, linear programming, half-space intersection, dominance counting

AMS(MOS) subject classifications. 68P05, 68P10, 68P20, 68Q20, 68U05

1. Introduction. We consider several search problems where we are given a set of n elements, which we call the *data set*. We are required to answer a sequence of queries about the data set.

The conventional approach to search problems consists of preprocessing the data set in time $p(n)$, thus building up a search structure that enables queries to be answered efficiently. Subsequently, each query can be answered in time $q(n)$. The time needed for answering r queries is thus $p(n)+r\cdot q(n)$. Very often, a single query can be answered without preprocessing in time $o(p(n))$. The preprocessing approach is thus uneconomical unless the number of queries r is sufficiently large.

We present here an alternative to preprocessing, in which the search structure is built up "on-the-fly" as queries are answered. Throughout this paper we assume that an adversary generates a stream of queries which can cease at any point. Each query must be answered *on-line*, before the next one is received. If the adversary generates sufficiently many queries, we will show that we build up the complete search structure in time $O(p(n))$ so that further queries can be answered in time $q(n)$. If on the other hand the adversary generates few queries, we will show that the total work we expend in the process of answering them (which includes building the search structure partially) is asymptotically smaller than $p(n)+r\cdot q(n)$. We thus perform at least as well as the preprocessing approach, and in fact better when r is small. We do so with no a priori knowledge of r. We call our approach *deferred data structuring* since we build up the search structure gradually as queries arrive, rather than all at once. In some cases we

* Received by the editors December 22, 1986, accepted for publication (in revised form) November 10, 1987.

† Computer Science Division, University of California, Berkeley, California 94720. The work of the first two authors was supported in part by the National Science Foundation under grant DCR-8411954. The results in § 4 first appeared in R. Motwani and P. Raghavan, *Deferred data structures: query-driven preprocessing for geometric search problems*, in Proc. 2nd Annual ACM Symposium on Computational Geometry, Yorktown Heights, NY, June 1986, pp. 303–312.

‡ IBM Thomas J. Watson Research Center, Yorktown Heights, New York 10598. The work of this author was supported in part by an IBM Doctoral Fellowship while he was a graduate student at the Computer Science Division, University of California, Berkeley, California 94720.