

Fast Construction of Near Optimal Probing Strategies

Eric Paulos, *University of California, Berkeley, CA, USA*

John Canny, *University of California, Berkeley, CA, USA*

We address the problem of defining and constructing optimal probing strategies for precisely localizing polygonal parts whose pose is approximately known. We demonstrate that this problem is dual to the well studied grasping problem of computing optimal finger placements as defined by Mishra et al. [17] and others [10, 16]. In addition we develop a quality measure for any given probing strategy based on a simple geometric construction in the displacement space of the polygon. Furthermore, we can determine a minimal set of probes that is guaranteed to be near optimal for constraining the position of the polygon. The size of the resulting set of probes is within $O(1)$ of the optimal number of probes and can be computed in $O(n \log^2 n)$ time. The result is a probing strategy useful in practice for refining part poses.

1 Introduction

In industrial manufacturing and automated assembly accuracy is money. Attaining and maintaining high precision can increase the cost of fixturing and feeding several fold [18]. The meaning of high versus low precision depends on the application, but for typical mechanical assembly, low precision tooling might provide accuracies in the tens of mils, while high precision would be around one mil or less (One mil = 10^{-3} inch = 25.4 microns). In this paper we study the *pose refinement problem*. In pose refinement, sensing is used as an inexpensive route to high precision part pose, assuming the pose is already known at low precision. Most research to date in computer vision and RISC sensing addresses the *pose acquisition problem*, where pose is determined with no knowledge of initial pose. The

result of pose refinement is a high-precision estimate, but it differs from the problem of high-precision pose acquisition. Since initial pose is approximately known in pose refinement, it can be used to make judicious choices about sensor placement. The same accuracy can be achieved with fewer or less expensive sensors for pose refinement as compared to pose acquisition, which must deal with all possible poses.

Pose refinement is a problem suggested to us by our visits to several state-of-the-art manufacturing companies, especially PTI (Productivity Technologies Inc.) and HP Labs. In typical industrial workcells, it was pose refinement rather than pose estimation that was the dominant sensing task. There are two reasons for this:

1. Feeder economics. Vibratory feeders are an inexpensive way to provide many part types in known (albeit low precision) pose. Small parts can also be fed on tape, which is more expensive (a couple of cents per part) but still costs far less than a high-precision pallet. So the initial and ongoing costs of achieving low-precision pose without sensing are small.
2. Multi-step manufacturing. In typical manufacturing, there are not one but several sequential stages, including assembly stages, testing and packaging. A single step might mate two parts whose poses are known at high precision. But the assembly step itself introduces a small amount of uncertainty, and it is expensive to transport the partial assembly at high precision to the next assembly stage. A more economical solution is to use pose refinement at the next stage. So while there might




Figure 1: A typical simple reflective sensor used for probing

be one pose acquisition step per part to get approximate initial pose, there will be several pose refinement steps for that part that start with the low precision output from the previous step and feeder, and increase the precision as needed for the next step.

So the arguments for pose refinement are (i) that it replaces the most expensive (high precision) fixturing and feeding steps (ii) it replaces the most frequent fixturing and feeding steps in multi-step assembly.

In this paper we employ simple light-beam sensors that act as line probes. A point light source and receiver define a line in space that is broken and unbroken by an object as it moves relative to the beam (see Figure 1). The positions of the object when the beam breaks give position readings, and three or more of these determine pose. Those readings are subject to error, and the pose estimate accuracy is limited by those errors and the sensor placements. This paper gives algorithms for choosing the probe placements to achieve near-optimal accuracy with a fixed number of probes, or to find a near-minimal number of probes to achieve a specified accuracy.

Problem Statement:

Given: A polygonal part, its geometry, and an initial pose estimate of $O_i = (O_{ix}, O_{iy}, O_{i\theta})$ within $\sigma = |O - O_i|_m$ of the exact actual pose of the object, O .

Solve: Find the optimal set of point probes defined as the minimal number of probes and




Figure 2: Example of a fixed array scanning beam sensor

their placement necessary to reduce the uncertainty in the position of an object to better than some acceptable level. The probes are defined by a set of fixed points and vectors denoting the direction of travel for each probe. A probe returns boolean value with a positional error of at most ϵ , where $\epsilon \ll \sigma$, based upon the presence or absence of an object's edge at a particular point along the path of the probe.

With enough time, we could simply perform a large number of probes as shown in Figure 3. However, in real industrial robotic assembly workcell design, throughput is a heavily weighted criterion. Therefore, our aim is to produce the best possible probing strategies that conform to the imposed constraints. The probing strategies we obtain can be used by any line probe of the object's 2d projection. There are natural generalizations to higher dimensions, although they are not as efficient. A typical probe, shown in Figure 1, consists of a simple reflective light beam sensor that can easily detect the presence or absence of an object. Our algorithm also allows us to construct specialized optimal probing strategies such as those for a scanning array of probes as shown in Figure 2.

We achieve these near-optimal probes with a small number of actual probes by maximizing the utility of each sensor probe we place on the object. This in turn makes the problem tractable for a real robot in a high throughput automation system. These strategies are within a constant factor of the optimal probing strategy and can be solved in $O(n \log^2 n)$ time whereas the exact optimal solution is in NP-hard [7].




Figure 3: A typical initial probe placement along the edges of an object

2 Previous and Related Work

Since our probing strategy is dual to the grasping problem and uses recent results in set covering algorithms, we report advances and related works in all of these areas.

2.1 Work in Probing and Grasping

The importance of probing in terms of localizing and identifying objects with probes has been explored by several individuals. Cole and Yap [6] and Bernstein [1] developed algorithms for choosing probes to obtain the geometry of an unknown two-dimensional convex object. A generalization of this strategy for higher dimensions is presented by Dobkin *et al.* [8] while a non-convex version was developed by Boissonnat and Yvinec [2]. Also, Lindenbaum and Bruckstein [13] describe similar probing strategies for a geometric probe composed of two line probes rotating about a common axis point.

Development of efficient algorithms for scanning objects with probes for the purpose of identification and localization have been explored by Wallack *et al.* [22, 23]. Likewise, point probing strategies have been developed for insertion operations by Paulos and Canny [19].

Jia and Erdmann [12] demonstrate an elegant technique for choosing placements of simple binary sensors to discriminate objects in the plane. In fact they also employ recent work on hitting sets and set coverings in solving their problem. Our work differs mainly in

the type of problem that is solved. Jia and Erdmann choose fixed probes to discriminate individual object poses from a large set of possible poses. The problem we tackle in this paper is how to best choose moving probes to refine the pose of a known object.

We also explore the development of optimal grasping techniques since we will show our probing problem to be dual to this problem. The need for good grasp planning algorithms for arbitrary shapes has always been important for robotics and industrial automation. The problem of optimal finger placement has been addressed by Mishra *et al.* [17] who define easily computable quality metrics for grasps. Markenscoff and Papadimitriou [14] chose to optimize the grasp with respect to minimizing the forces needed to balance the object’s weight through friction. Ponce and Faverjon [20] fix the number of fingers and solve a system of linear constraints in the positions of the fingers to optimally position them along the polygonal edges. A similar technique for three-dimensional polyhedral objects was developed by Ponce *et al.* [9]. Goldberg [11] also details a method for choosing grasps with a parallel jaw gripper when the initial pose of the object is unknown. Other optimizing grasps techniques based on simple geometric constructions have been developed by Brost [4] and later Mirtich, Canny, and Ferrari [10, 16].

2.2 Work in Set Coverings

We will show that finding the minimal set of probes is equivalent to solving the convex set covering problem. This problem is discussed by Clarkson [5] who describes a $O(cn \log^{O(1)} n)$ time randomized algorithm for finding covering sets of cardinality within $O(\log c)$ of the optimal set covering c .

More recent results by Brönniman and Goodrich [3] on the dual problem of finding minimal hitting sets improves on these bounds. They demonstrate an $O(n \log^2 n)$ algorithm that finds a hitting set of size $O(1)$ from the optimal set size. They employ work by Matoušek [15] using ϵ -nets.

3 Defining Optimality

When probing an object we would like to choose point probes that allow the minimum variation of the object pose. Our point probes inherently contain some known error so it is not enough to take k independent measurements to constrain k degrees of freedom. The placement of the probes effects the worst case object displacement. Therefore, we are interested in the relationship between the object displacements and the corresponding probe displacements. Our goal is to find a set of probe placements that minimizes the potential worst case object displacement.

3.1 Object Pose Definition

We define O as the actual pose of the object in two-dimensions as

$$O = (O_x, O_y, O_\theta)$$

From our problem statement, we note that we are interested not in locating an object, but in refining the position of a known object whose pose is known to some reasonable degree of accuracy. Our approach relies on this initial coarse accuracy pose information. We define our assumed initial pose as O_i and quantify a bound on the worse case displacement of the assumed pose from the actual pose as

$$\|O_i - O\|_m \leq \sigma$$

where $\|\cdot\|_m$ is defined to be the m -norm. At this point, we run into the usual problem of defining a metric on a space with distance and angular coordinates. There may be application-specific ways to weight the angular component, but a good default is to weight the angular component by the object's radius (i.e. the largest distance from any point in the object to its coordinate origin). With this choice, the metric bounds the maximum distance between any two corresponding points on the object at O and O_i . A typical value for σ would be tens of mils. Finally, while we are considering an m -norm for generality, the 2-norm would seem to be the most natural choice.

Recall that we will be using line probes to refine the position of the object. Therefore, for a given set of

probe measurements there will also be a set of valid poses for the object consistent with those sensor readings. We denote this object pose as \bar{O} and define it to be an object pose chosen by an adversary consistent with some sensor readings given the object is at O . We define the difference between the actual object position and the adversary's choice as o .

$$o = O - \bar{O}$$

Recall that we are attempting to refine the position of the object so that σ will always be at least an order of magnitude larger than $\|o\|$. In terms of linear displacement the initial pose uncertainty, σ , is typically on the order of tens of mils while the uncertainty under probing is on the order of a mil or less.

$$\sigma \gg \|o\|$$

To quickly summarize we have O as the actual pose of the object, O_i as our initial pose estimate, and \bar{O} as a pose that an adversary can choose that does not violate our sensor readings. That is, we cannot determine from the sensors if the object is at \bar{O} or not. We will clarify later exactly how \bar{O} is defined.

3.2 Probe Placement

We construct probes along the perimeter of the object and denote them as

$$p = (p_x, p_y, l_x, l_y)$$

where (p_x, p_y) is the point where the line probe touches the object when the object is at O_i , and (l_x, l_y) is the direction of the line probe. We must guarantee that no matter where the object actually is, this line probe always contacts the same edge. Assuming that the object radius was used to weight the angular component of the pose metric, this can be accomplished in a simple way: construct a strip about the line probe l whose boundaries are parallel to the line, and at distance σ from it; This strip represents the possible relative positions of the line probe for various actual object poses O . If the edge we are probing crosses the entire strip, it will always be probed correctly. If the edge crosses




Figure 4: Relation of placement of probes along an edge and the size of the allowable displacement regions

only part of the strip, then there is a possible O such that this line probe misses the edge completely. From now on, we assume that all line probes are chosen so that their σ strips touch only the edge of interest.

The initial probe placement consists of placing a pair of probes on each edge. Each probe is placed as near as possible to an endpoint of the edge, but subject to the strip constraint above. Our algorithm will choose a subset of these initial probes as the near optimal probe set. Our choice is based on the fact that we receive the most accurate pose information by probing near the vertices of an object. Observe that probes near the vertices give rise to large sensor displacements as a result of small rotational perturbations, while position information is the same anywhere on the edge. Figure 4 demonstrates how moving a set of probes with a given error out towards the vertices of an edge shrinks the size of allowable displacements for that edge.

This set of probes is guaranteed to contain the optimal probe placement. Any edge-interior probes would give only redundant information in the worst case, and our probe choice is based on a worst-case analysis. A typical initial probe placement example is shown in Figure 3. The remaining problem is to determine a subset of these probes that still provide a substantial gain in object pose accuracy.

3.3 Probing Function

We place a coordinate system at the center of mass of the object. In addition we define the rotational displacement of the object to be about this c.o.m. axis. In




Figure 5: Original probe p and probe p' resulting from an object displacement

Figure 5 we depict the construction of the corresponding probe displacement for a given object displacement. This will define the probing function. In this figure n is a unit normal to the edge being probed, p is the initial probe location and p' is its location after the displacement O from the origin.

Recall that σ is very small allowing us to take small angle approximations and write

$$p'_k \approx p_k + (O_x, O_y) + p_k^\perp O_\theta$$

where $p^\perp = (x, y)^\perp = (-y, x)$ and k denotes the k th probe. It follows that the change in probe position is

$$\begin{aligned} \Delta p_k &= p'_k - p_k \\ &= (O_x, O_y) + p_k^\perp O_\theta \end{aligned}$$

The probe actually only gives us useful position information normal to the edge being probed. We could freely displace the object along the edge without changing that probe reading. Therefore, the change in probe position along the edge normal n_k can be written as $n_k \cdot \Delta p_k$. Observe that even if we approach an edge at an angle, when we detect the edge we can only claim that some point of the edge must intersect the detected point. This is equivalent to the information we receive if we approach normal to the edge. Therefore, the two probe approach techniques are equivalent and thus the choice of edge approach is independent and left as a final implementation detail. It *does* affect the σ -strip described earlier, and the amount of clearance from the edge endpoint needed to ensure the correct edge is detected.

We are now ready to define the probing function $P : \mathbb{R}^3 \rightarrow \mathbb{R}^k$ to be a real valued function which maps object positions into ideal probe outputs of the form (P_1, P_2, \dots, P_k) . We define each element to be

$$P_k(O) = \Delta p_k \cdot n_k \quad (1)$$

Our probes will have a sensor error ϵ , typically a mil or so. We define the measured probes as $\bar{P} \in \mathbb{R}^k$. Given a sensor error of ϵ , we observe that the measured probe values \bar{P} must be consistent with the ideal probes given object pose O .

$$\|\bar{P} - P(O)\|_\infty \leq \epsilon \quad (2)$$

Similarly, any possible object position \bar{O} that the adversary chooses must have all measured probes within ϵ of the given measurements.

$$\|P(\bar{O}) - \bar{P}\|_\infty \leq \epsilon \quad (3)$$

Using the triangle inequality on these last two expressions, we find that the O and \bar{O} satisfy

$$\|P(O) - P(\bar{O})\|_\infty \leq 2\epsilon \quad (4)$$

and observe that for any \bar{O} satisfying this inequality that there is a \bar{P} satisfying Equations 2 and 3. Thus the combined bound is tight. We will call the set of object displacements \bar{O} that satisfy this inequality \mathcal{K} .

Recall that the actual position of the object is defined as O and the possible interpreted object position for some sensor reading is \bar{O} where \bar{O} is any \bar{O} satisfying Equation 4. We want to constrain the distance between the interpreted object position and the actual object position to be as small as possible. This in turn minimizes the worst case distance between the actual and measured poses, which is the ultimate goal of pose refinement. We represent the former quantity as

$$\|O - \bar{O}\|_m \quad (5)$$

We employ an adversarial argument and note that if an adversary is allowed to move the object to some valid \bar{O} consistent with the sensor readings it will always choose the \bar{O} such that Equation 5 is maximized. We express this as

$$\max_{\bar{O} \in \mathcal{K}} \|O - \bar{O}\|_m \quad (6)$$

However, we are allowed to choose the set of probes P . Furthermore, we desire a set of probes that will output drastically different values for different nearby object poses, thus allowing us to identify different poses easily. Essentially we would like to eliminate the possibilities of obtaining identical or near-identical sensor readings for an object in two different poses. We can write this as

$$\max_P \|P(O) - P(\bar{O})\|_\infty \quad (7)$$

or

$$\min_P \frac{1}{\|P(O) - P(\bar{O})\|_\infty} \quad (8)$$

Since both equations 6 and 8 scale linearly with $O - \bar{O}$, it is natural to combine them as a ratio which is then independent of the magnitude of $O - \bar{O}$:

$$\min_P \left(\max_{\bar{O} \in \mathcal{K}} \frac{\|O - \bar{O}\|_m}{\|P(O) - P(\bar{O})\|_\infty} \right) \quad (9)$$

Since $P(O)$ is linear in O we can simplify Equation 9 and arrive at our final optimality criterion and probe quality measurement Q .

$$Q(P) = \min_P \left(\max_{\bar{O} \in \mathcal{K}} \frac{\|O - \bar{O}\|_m}{\|P(O) - P(\bar{O})\|_\infty} \right) \quad (10)$$

4 Displacement Space

Working in displacement space, we observe that there is a simple geometric construction of the optimality criterion as given in Equation 10. Displacement space, denoted \mathbb{D} and $\mathbb{D} \in \mathbb{R}^3$, is the space of all displacements in (x, y, θ) of the object O to be probed. Each probe sensor that we introduce imposes constraints on the allowable set of displacements of the object without violating the probe value.

Equation 4 from the previous section defines a pair of halfspaces in displacement space \mathbb{D} for each probe p_k .

$$\begin{aligned} \|P_k(O) - P_k(\bar{O})\|_\infty &\leq 2\epsilon \\ \|P_k(O - \bar{O})\|_\infty &\leq 2\epsilon \\ \|P_k(o)\|_\infty &\leq 2\epsilon \\ \|n_x o_x + n_y o_y + (p^\perp \cdot n) o_\theta\|_\infty &\leq 2\epsilon \end{aligned}$$

Fast Construction of Near Optimal Probing Strategies

These two halfspace can be written as

$$n_x o_x + o_y o_y + (r^\perp \cdot n) o_\theta - 2\epsilon \leq 0 \quad (11)$$

$$n_x o_x + o_y o_y + (r^\perp \cdot n) o_\theta + 2\epsilon \geq 0 \quad (12)$$

The intersection of all $2k$ halfspaces constructed from k probes by definition represents a convex polytope in \mathbb{D} . We name this polytope S with the definition

$$S = \bigcap_{h \in \mathcal{H}(P)} h$$

where $\mathcal{H}(P)$ is the family of $2k$ halfspaces defined by the set of k probes P .

In displacement space this polytope S will have furthest outlying point which will occur in the non-degenerate case at a vertex of S . This furthest outlying point represents the largest object displacement from the assumed pose that still satisfies the given probe measurements. More formally we define this point as

$$\Gamma(S) = \sup_{q \in S} \|q\|_m$$

The distance to $\Gamma(S)$ is exactly the optimality criteria as defined in Equation 10. That follows because points in S fix the denominator of Equation 10, and $\Gamma(S)$ is chosen to maximize its numerator. We assume that P is fixed, so there is only optimization by the adversary over q . Recall that an adversary can choose the actual sensor readings \bar{P} such that the object displacement $\Gamma(S)$ is a valid interpretation of \bar{P} . Hence, this is the largest displacement of the object undetectable by the given probing strategy.

For illustrative purposes we work through a simple example without rotation. In Figure 6 we show three probes on a triangle, an admittedly simple case, but enough to demonstrate our method. Notice that each probe in real space gives rise to a pair of parallel halfspaces in displacement space, \mathbb{D} . If we remove probe p_2 , the area of polygon S in displacement space increases which represents the additional translational freedom that the object can undergo and still remain consistent with the remaining two sensor readings. Therefore, the added sensor p_2 is a useful addition since it decreases the area of S and reduces the distance to $\Gamma(S)$.

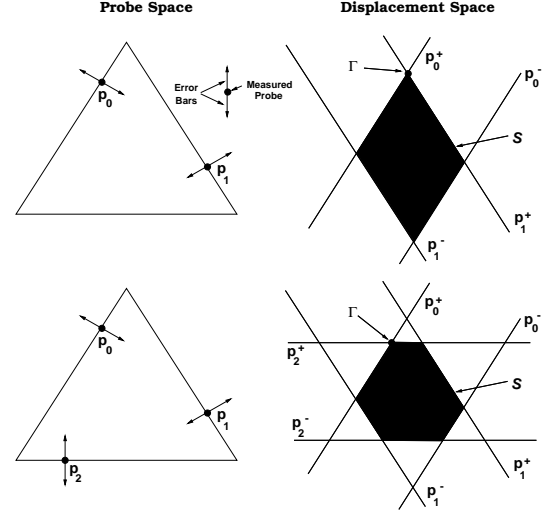


Figure 6: Simple example of probe space and displacement space without considering rotation

We are interested in probing strategies P' that have approximately the same quality metric Q . Remember that every probe we remove from P removes a pair of halfspace in \mathbb{D} . This in turn changes the shape of S but not always the point $\Gamma(S)$ which defines the optimality criteria. Therefore, we would like to find other optimal probes with fewer probes. In particular, we would like to find

$$\min_{P' \subset P} |P'| : Q(P') \approx Q(P)$$

where $|\cdot|$ is simply the cardinality of the set P' .

We define S' to be the polytope defined by the intersection of the halfspaces defined by the probes P' . Observe that

$$S' \supseteq S$$

which implies that when we remove a probe, hence two halfspaces, we expect the furthest outlier to remain where it is or increase in distance from the origin giving

$$|\Gamma(S')| \geq |\Gamma(S)|$$

Rather than removing half-planes in an *ad hoc* manner such that $Q(S')$ remains essentially unchanged, we will dualize and solve for a minimal convex set covering for the corresponding points in the dual. These minimal set of points will be exactly dual to the minimal

set of halfspaces in displacement space by definition of the minimal convex set cover problem. These resulting halfspaces in \mathbb{D} correspond to minimal set of probes, as desired. We discuss this dualization in the next section.

Observe that the production of any such probing strategy is independent of the error, 2ϵ . This is true because we are interested in optimizing the ratio shown in Equation 10. One can also note that topology of the polytope S of the solution space is independent of ϵ which only serves as a scaling factor. That is, when we double ϵ we get the same polytope at twice the size (four times the volume). Therefore, without loss of generality we set ϵ to one for the duration of the paper.

5 Displacement Space Dual

A strong relationship to grasping is shown in this section. We show that finding the optimal k probe placements is equivalent to finding the optimal push-pull grasp for a set of k fingers. We recall that a push-pull grasp is defined as a grasp that employs fingers capable of exerting a pushing or pulling force at the contact.

We define \mathbb{D}^D to be the dual of \mathbb{D} . We define the dual exactly in Table 1. In this mapping we show how points in \mathbb{D} map to planes in \mathbb{D}^D and similarly for planes in \mathbb{D} to points in \mathbb{D}^D . We note that by definition the dual of \mathbb{D}^D is \mathbb{D} , hence the duality operation is symmetric.

\mathbb{D}		\mathbb{D}^D
$p : (p_x, p_y, p_\theta)$	\leftrightarrow	$p^D : p_x x + p_y y + p_\theta \theta = 1$
$f : ax + by + c\theta = 1$	\leftrightarrow	$f^D : (a, b, c)$
$S = \text{polytope}$	\leftrightarrow	$S^D = \{f^D : f \cap \text{Int}(S) = \emptyset\}$

Table 1: Duality Mappings

Observe that a polytope S defined as the intersection of a set of halfspaces h_k becomes the polytope S^D . We define $\text{Bound}(h_k)$ to be the plane on the boundary of the halfspace h_k . The polytope S^D can also be expressed as the convex hull of the union of dual points

$\text{Bound}(h_k)^D$.

$$S^D = \text{Conv}(\cup \text{Bound}(h_k)^D)$$

Let $r \in S$. The distance of r from the origin in \mathbb{D} is simply

$$|r| = \sqrt{r_x^2 + r_y^2 + r_\theta^2}$$

The dual plane r^D in \mathbb{D}^D by definition is represented as

$$r_x x + r_y y + r_\theta \theta = 1$$

This distance of the closest point on this plane to the origin in \mathbb{D}^D is given by

$$|r^D| = \frac{1}{\sqrt{r_x^2 + r_y^2 + r_\theta^2}}$$

Setting $\alpha = \sqrt{r_x^2 + r_y^2 + r_\theta^2}$ we get that the distance of this point r from the origin in \mathbb{D} is α and the minimal distance of the dual plane r^D from the origin in \mathbb{D}^D is $\frac{1}{\alpha}$. Therefore,

$$|r| = \frac{1}{|r^D|}$$

Let f_c be the closest plane to the origin of \mathbb{D}^D not intersecting $\text{Int}(S^D)$. The distance to f_c is the same as the distance to the closest point u_c in the boundary of S^D (which is contained in f_c). And its easy to see that $\Gamma(S)^D = f_c$ where $\Gamma(S)$ was defined earlier as the furthest outlying point in S .

The closest point to the origin in the boundary of a polytope lies on the largest inscribed sphere centered at the origin. Observe that $\Gamma(S)$ lies on the smallest circumscribing sphere of S in \mathbb{D} . Therefore, finding the smallest circumscribing sphere Σ for a polytope is equivalent to finding the largest inscribed sphere Σ^D of the dual polytope S^D . This follows from the relationship

$$|r_\Sigma| = \frac{1}{|r_{\Sigma^D}|}$$

Recall that the planes $\text{Bound}(h)$ through the halfspaces in \mathbb{D} dualize to the points

$$\text{Bound}(h_k)^D = (n_x, n_y, \mathcal{D}_k^\perp)$$

These points are equivalent to the wrenches due to unit pull finger forces acting at p . Recall that in the probing problem we obtain a pair of halfspaces for each probe. Hence, the optimal probing problem is equivalent to the optimal push-pull grasping problem. We use the optimal grasping criteria as defined by Mishra *et al.* [17] and others [10, 16] which is the set of finger placements such that the one-norm of the finger forces can resist the largest externally applied wrench on the object. We also define the optimal probe placement as the minimal number of probes and their placement necessary to reduce the uncertainty in the position of an object to better than some acceptable level. Using these metric definitions, we obtain the following result.

Theorem 1 *Finding the optimal placement of k probes is equivalent to finding the optimal push-pull frictionless grasp for a set of k fingers.*

6 Hitting Sets and Set Covers

Recall that the quality of that probing strategy is given directly by the radius of the maximally inscribed sphere in S^D . We would like to remove some vertices of S^D such that the radius of the maximally inscribed sphere does not decrease by much.

This problem can be posed as a convex set cover problem which states that given L and U a set of points in \mathbb{R}^d , find $C \subseteq U$ with $L \subseteq \text{Conv}(U)$. Here we have L as the sphere of desired radius. This problem has been studied by several individuals. Recently, Clarkson [5] describes a randomized algorithm for computing the three-dimensional convex point set cover from an initial set of n points to within $O(\log c)$ of the optimal cover of c points. His algorithm has a running time of $O(cn \log^{O(1)} n)$.

Recent work by Brönniman and Goodrich [3] improve on both the running time and approximation to the optimal convex set covering. Their deterministic algorithm solves the equivalent problem of finding a minimal hitting set, where a hitting set is a subset $H \subseteq X$ such that H has a non-empty intersection with every set R in a collection of subsets of X . Their algorithm employs work by Matoušek [15] on ϵ -nets to obtain a

hitting set in $O(n \log^2 n)$ time that is within $O(1)$ of the optimal size hitting set. This set corresponds exactly to the optimal probe placement which we define as a set of c probes that reduce the uncertainty in the position of an object to at least some necessary level for the operation to be performed.

In our optimal construction we obtain pairs of halfspaces, hence pairs of points in the dual. However, in the Brönniman and Goodrich algorithm they are treated as two completely unrelated elements. This will result in near-optimal set sizes that are in the worst case twice as large as we could achieve by grouping the pairs. Alternatively, we can group them to obtain the near-optimal hitting set at a slight running time cost. This performance slowdown is a result of an increase in the VC-dimension [21] as a result of our pairing.

VC-dimension, named for Vapnik and Červonenkis, is defined for a range space (X, R) with $P \subseteq X$ as the cardinality of the largest set P that is shattered by R . A set P is shattered by R if $\Pi_R(P)$ is the power-set of P .

To obtain an optimal probing strategy for an array of scanning sensors as shown in Figure 2, we identify the co-linear points in the displacement space and assign them labels such that the hitting set algorithm will include all or none of a set of co-linear points in the probe optimization selection. This also results in an increase in the VC-dimension which affects the running time but still finds a hitting set within $O(1)$ of the optimal one.

Our algorithm successfully handles other variations similar to the co-linear constraint for the scanning sensor without major modification. This makes it well adapted to situations where optimal probing strategies under special constraints are needed and not intuitive to observe.

The Lemma below summarizes much of the results of this paper.

Lemma 1 *A near optimal set of c point probes can be found for any polygonal object in $O(n \log^2 n)$ time. Furthermore, the size of the set c will be within $O(1)$ of the size of the optimal set of c point probes.*

7 Conclusion

We have demonstrated a fast method by which optimal probe placements can be obtained for any known polygonal object. More importantly, we can guarantee that the number of probes resulting from our technique is within a constant of the actual optimal number of probes necessary. These probing strategies refine the position of an object whose pose is approximately known. Furthermore, we claim that this is a real problem often encountered in industrial manufacturing. Our paper also shows that the problem of optimal probe placement is dual to the well studied push-pull grasping problem of positioning frictionless fingers on an object.

Acknowledgements

This research was supported in part by National Science Foundation Grant #IRI-9114446. The authors would like to thank Brian Mirtich and Ken Goldberg for several insights and suggestions towards the development of this paper. We would also like to thank the anonymous reviewers who provided numerous useful comments and suggestions.

References

- [1] Herbert Bernstein. Determining the shape of a convex n -sided polygon by using $2nk+k$ tactile probes. In *Information Processing Letters*, pages 225–260, 1986.
- [2] J. Boissonnat and M. Yvinec. Probing a scene of nonconvex polyhedra. In *Algorithmica*, pages 321–342, 1992.
- [3] H. Brönnimann and M.T. Goodrich. Almost optimal set covers in finite vc-dimension. In *Proc. 10th ACM Symp. on Computational Geometry (SCG)*, pages 293–302, 1994.
- [4] Randy C. Brost. Automatic grasp planning in the presence of uncertainty. *International Journal of Robotics Research*, 7(1):3–17, 1988.
- [5] Kenneth L. Clarkson. Algorithms for polytope covering and approximation. In *Proc. 3rd Workshop Algorithms Data Struct.*, volume 709, pages 246–252, 1993.
- [6] Richard Cole and Chee Yap. Shape from probing. *Journal of Algorithms*, 8:19–38, 1987.
- [7] G. Das and D. Joseph. The complexity of minimum nested polyhedra. In *Canadian Conference on Computational Geometry*, 1990.
- [8] D. Dobkin, H. Edelsbrunner, and C. Yap. Probing convex polytopes. In *18th Annual ACM Symposium on Theory of Computing*, pages 424–432, 1986.
- [9] B. Faverjon and J. Ponce. On computing two-finger force closure grasps of curved 2d objects. In *IEEE International Conference on Robotics and Automation*, pages 424–429, 1991.
- [10] C. Ferrari and J.F. Canny. Planning optimal grasps. In *IEEE International Conference on Robotics and Automation*, pages 2290–2295, 1992.
- [11] Ken Goldberg. Orienting polygonal parts without sensors. *Algorithmica. Special Issue on Computational Robotics*, Volume 10(3):201–225, 1993.
- [12] Yan-Bin Jia and Michael Erdmann. The complexity of sensing by point sampling. *Algorithmic Foundations of Robotics*, pages 283–298, 1995.
- [13] M. Lindenbaum and A. Bruckstein. Reconstructing a convex polygon from binary perspective projections. In *Pattern Recognition*, pages 1343–1350, 1990.
- [14] Xanthippi Markenscoff and Christos H. Papadimitriou. Optimum grip of a polygon. *International Journal of Robotics Research*, 8(2):17–29, 1989.
- [15] J. Matoušek. Construction of ϵ -nets. In *Discrete Computational Geometry*, pages 5:427–448, 1990.
- [16] Brian Mirtich and John Canny. Easily computable optimum grasps in 2d and 3d. In *IEEE International Conference on Robotics and Automation*, 1994.

Fast Construction of Near Optimal Probing Strategies

- [17] B. Mishra, J. T. Schwartz, and M. Sharir. On the existence and synthesis of multifingered positive grips. *Algorithmica*, 2:541–558, 1987.
- [18] James Nevins and Daniel Whitney. Computer controlled assembly. *Scientific America*, 238(2):62–74, February 1978.
- [19] Eric Paulos and John Canny. Accurate insertion strategies using simple optical sensors. In *IEEE International Conference on Robotics and Automation*, pages 1656–1662, May 1994.
- [20] J. Ponce and B. Faverjon. On computing three-finger force closure grasps of polyhedral objects. In *International Conference on Advanced Robotics*, pages 1018–1023, 1991.
- [21] V. N. Vapnik and A. Ya. Červonenkis. On the uniform convergence of relative frequencies of events to their probabilities. In *Theory Probab. Appl.*, pages 16:264–280, 1971.
- [22] Aaron Wallack and John Canny. Planning for modular and hybrid fixtures. In *IEEE International Conference on Robotics and Automation*, May 1994.
- [23] Aaron Wallack, John Canny, and Dinesh Manocha. Object localization using crossbeam sensing. In *IEEE International Conference on Robotics and Automation*, pages 692–699, 1993.