Journal of Computer and System Science, In Press

Randomized Query Processing in Robot Path Planning

Lydia E. Kavraki* Jean-Claude Latombe* Rajeev Motwani†
Prabhakar Raghavan‡

Abstract

The subject of this paper is the analysis of a randomized preprocessing scheme that has been used for query processing in robot path planning. The attractiveness of the scheme stems from its general applicability to virtually any path-planning problem, and its empirically observed success. In this paper we initiate a theoretical basis for explaining this empirical success. Under a simple assumption about the configuration space, we show that it is possible to perform preprocessing following which queries can be answered quickly. En route, we consider related problems on graph connectivity in the evasiveness model, and art-gallery theorems.

*Robotics Laboratory, Department of Computer Science, Stanford University, Stanford, CA 94305-2140. Partially supported by ARPA grant N00014-92-J-1809 and ONR grant N00014-94-1-0721.

†Department of Computer Science, Stanford University, Stanford, CA 94305-2140. Supported by an Alfred P. Sloan Research Fellowship, an IBM Faculty Development Award, and NSF Young Investigator Award CCR-9357849, with matching funds from IBM, Schlumberger Foundation, Shell Foundation, and Xerox.

‡IBM Research Division, T.J. Watson Research Center, Yorktown Heights, NY 10598.

1 Introduction

Planning obstacle-avoiding motion for a rigid or articulated robot from a given initial configuration to a goal configuration is an important problem in robotics [4, 19]. The set of feasible paths depends on the robot geometry, its motion capabilities and the workspace geometry. In many applications the environment is static and the robot must perform a series of complicated maneuvers to achieve a sequence of goals. Examples of such applications include maintenance of cooling pipes in nuclear plants, point-to-point welding for car assembly, and cleaning airplane fuselages. The path planning problem is receiving increased recent attention because of further applications (e.g., in movie generation, where planning can drastically reduce the need for input from human animators [18], and pharmeceutical drug design where motion planning can be used to animate the motion of a ligand molecule docking against a larged receptor molecule.

A number of recent papers in the robotics literature [10, 14, 15, 16, 27, 28] have described the success of a class of randomized preprocessing heuristics for query processing in robot path planning. The key idea is the use of random sampling in a preprocessing stage, following which queries of the form "Is configuration B reachable from configuration A?" can be answered quickly. The method is very general and can be applied to virtually any type of holonomic robot. It has proved especially effective for robots with many degrees of freedom, where traditional methods have either failed to yield algorithms or have yielded algorithms that are too slow for normal use. Figure 1 depicts several positions of a robot with 7 revolute joints to which the method has been successfully applied.

There is another motivation for such a general query processing scheme not bound to the specifics of any particular robot: it is clearly infeasible to invest effort in tailor-made algorithms for every robot in existence. While the scheme is general, it is possible to tailor it to any specific type of robot and further enhance its performance [15].

This paper initiates a theoretical basis for explaining the success of this method. In our work we focus on the geometric motion planning problem where the goal is only to determine a geometric description of a feasible path. The actual implementation of the path would require dealing with mechanical issues such as path smoothing, handling uncertainties, mechanical control, and dynamic constraints [19].

Before we describe our results in detail, we review previous work in this area and discuss the difficulties encountered in extending that work to path planning for robots with large degrees of freedom.

1.1 Background and Motivation

The configuration of a robot at any instant is described by an ordered tuple of real values, each entry of which is the value of one component of its position. For example, a unit square moving freely in the plane is captured by a triple: the x- and y-coordinates of a designated corner, together with the angle made by the line containing a designated edge with the x-axis. We therefore say that such a square has 3 degrees of freedom, and represent its position by a point in 3-space. The motion of the square forms a trajectory in this space. Given static obstacles in the plane that constrain the motion of the square, we may represent them in the space as a set of forbidden regions that must never be entered by the motion trajectory. The 3-dimensional space representing the position of the square together with these forbidden regions is known as the *configuration space* for this setting.

In general, the configuration of any robot with d degrees of freedom can be represented in a parametric d-dimensional configuration space $\mathcal{C}$. Such a configuration space can be defined for any

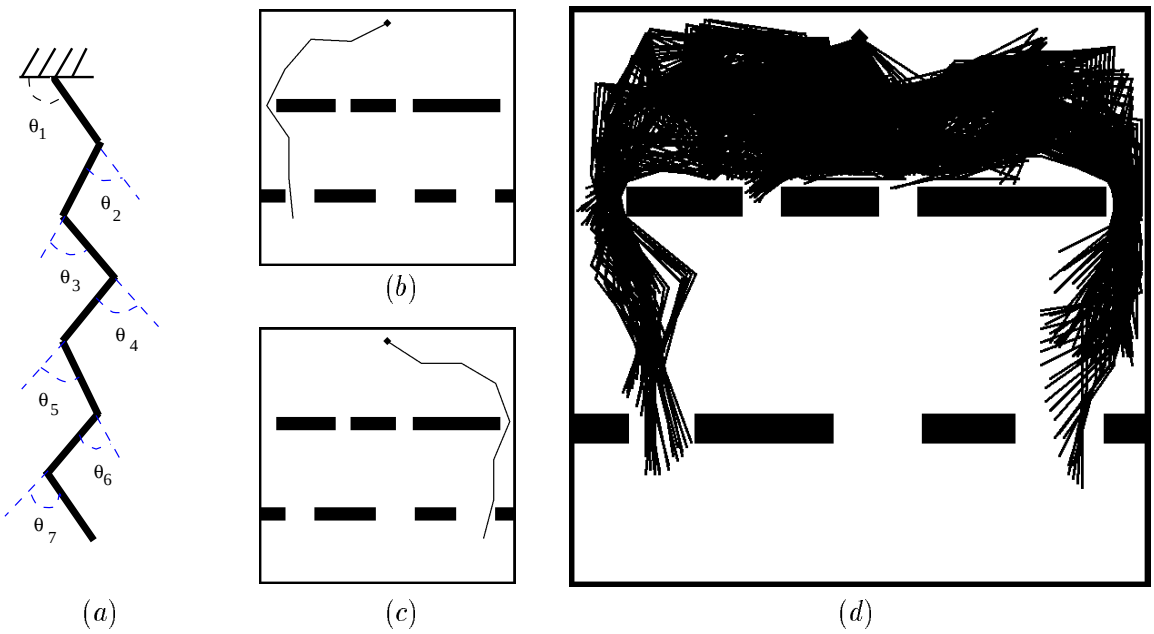

Figure 1: *Several configurations of a robot arm with a fixed base. This arm has 7 revolute joints, and must maneuver through gaps in two walls. (a) a planar arm with 7 revolute joints, (b) and (c) two different configurations of the arm, and (d) a path followed by the arm when it moves between configurations (b) and (c).*

motion planning problem and, together with a cost measure and possible constraints on the shapes of trajectories, defines the problem completely. The obstacles in the workspace induce $\mathcal{C}$-*obstacles* in the configuration space that are forbidden configurations of the robot. We refer to the subset $\mathcal{F}$ of the configuration space that is *not* forbidden as the *free space*; it may consist of more than one connected component. When a query asks whether the robot can move from one configuration to another, it is asking whether the corresponding points of the configuration space lie in the same connected component of free space.

For instance, the position of the arm in Figure 1 may be represented in a space with 7 dimensions, with each dimension corresponding to the angular position of one of the revolute joints. Figure 1 (a) depicts the 7 angles giving rise to the seven degrees of freedom. Figure 1 (b) and 1 (c) depict a pair of start and finish configurations, while Figure 1 (d) depicts a sequence of configurations found by the algorithm for going from the start configuration to the finish configuration.

The configuration space representation of a motion planning problem is due to Lozano-Perez [22, 23], although it was implicit in some earlier work [24, 32]. This representation is particularly amenable to an algebraic treatment of the problem, such as in terms of semi-algebraic sets [19]. A semi-algebraic set is a subspace of $\Re^d$ defined by quantified first-order formulas over comparisons of polynomials with constants. It can be shown that if the obstacles and obstacle boundaries are semi-algebraic sets, then the $\mathcal{C}$-obstacles are also semi-algebraic sets. This view led to the development of several planning methods all of which modeled obstacles as having low-degree algebraic surfaces, and reduced the planning problem to deciding the satisfiability of sentences in the first-order theory of reals. Specifically, Schwartz and Sharir [31] considered the setting where the d-dimensional configuration space contains a free space defined by n polynomial constraints of maximum degree

m; they presented an algorithm that is polynomial in n and m, but doubly-exponential in d. Canny [4] gave a different algorithm that is only singly-exponential in d. Thus, the practical impact of such methods has been restricted to the case of a small number of degrees of freedom, e.g., the motion planning problem for rigid robots in the plane with polygonal obstacles. Fully polynomial algorithms are known only for extremely special situations such as motion planning for a planar arm in a circular workspace without any obstacles (due to Hopcroft, Joseph, and Whitesides [7] and Kantabutra and Kosaraju [13]).

Unfortunately, it is believed that these positive results are close to the best possible since most versions of the motion planning problem are known to be PSPACE-hard, or even NEXPTIME-hard [19]. The most basic of these results is due to Reif [29] who established the PSPACE-hardness of the so-called *generalized mover's problem* — the path planning problem for a collection of polyhedral bodies linked at vertices and moving amongst a finite set of polyhedral obstacles. Path planning for various special cases are also known to be PSPACE-hard: Hopcroft, Joseph, and Whitesides [6] established PSPACE-hardness for *planar linkages*, a set of rigid 1-dimensional links connected by revolute or fixed (possibly, multiple) joints moving in the plane; Hopcroft, Schwartz, and Sharir [8] and Hopcroft and Wilfong [9] established the PSPACE-completeness for *multiple rectangles*, a set of rectangles executing coordinated axis-parallel translations in a rectangular workspace. Joseph and Plantiga [11] established the PSPACE-hardness for *planar arms*, a sequence of links connected by revolute joints moving in the plane among polygonal obstacles. Similar results have been obtained for problems such as motion planning with moving obstacles [4, 30] and compliant motion planning with uncertainty [5]. In all these cases, the basic source of complexity appears to be the dramatic increase in the difficulty of the problem with increasing dimensionality.

Practical approaches to path planning can be viewed as falling into one of the following two categories: *global methods* which involve exhaustive preprocessing of the free space and require time exponential in the dimension; and, *local methods* which are essentially localized heuristics that are fast for special cases but suffer from the possibility of failure to find paths even when one exists. One example of a global method is the technique of Schwartz and Sharir [31] that computes a cell decomposition of the free space and uses a search graph based on this decomposition. Another example is the *roadmap method* where a set of canonical paths is used to cover the components of the free space, and the planning task reduces to determining a connection to the canonical paths. Examples of this include the visibility graph approach of Nilsson [25] and the silhouette approach of Canny [4]. An example of a local method is the *potential field technique* [1, 17] where a potential function is defined at each point in the free space, based on an attraction component from the goal point and a repulsion component from the obstacle boundaries, and the planning process corresponds to a determination of the global minima of the potential function using a greedy, local search. The problem with this approach is that with increasing dimensionality, it becomes increasingly difficult to determine good potential functions which do not cause the planner to get trapped in local minima.

All these approaches suffer from severe performance degradation with increasing dimensionality, and there is a trade-off between slow but complete global planners and fast but incomplete local planners. In this context, a recent research trend in motion planning has been the use of randomized preprocessing to obtain global completeness using only local computations in query processing. For example, Barraquand and Latombe [2] present an approach based on potential functions with the twist that Brownian motion (random walk) is used to escape from local minima. While this approach gave improved performance for simple geometries in up to 30 dimensions, it still suffers from the need for application-specific potential functions and exhibits performance degradation with more complex geometry. Another randomized approach proposed by Kavraki and Latombe [14] is

a probabilistic variant of the roadmap technique. This approach has been successful in practice and is the focus of our work.

1.2 A Model for Analyzing a class of Randomized Preprocessing Schemes

We assume here that the configuration space is the cube $[0,1]^d$, where d is the number of degrees of freedom for the robot. (Our definitions and results can be extended to cases where one or more dimensions of the configuration space — say the angular position of a joint of an arm — can "wrap around", but for simplicity we assume $[0,1]^d$ here.) For the purposes of this abstract, we also assume that the space is *reflexive:* if a point p_1 in free space is reachable from p_2, then p_2 is reachable from p_1. Non-reflexive spaces arise, for instance, when there are moving obstacles so that time becomes one dimension, or if the robot has asymmetric motion capabilities such as when only forward motion is permitted.

A key ingredient of the probabilistic raodmap method is a fast *simple planner* that, given two points p_1 and p_2 in the configuration space, tries to connect them using a fast but simple strategy. For example, one simple planner that has been used for this purpose [15, 16] checks whether the line segment between p_1 and p_2 lies entirely in free space; if not, it reports failure (even though a more complicated path might exist). This is usually implemented by a walk along the line segment (suitably discretized), checking whether each of these discrete points is in free space. In addition we assume that we have access to a *complex planner* that is expensive to run, but is error-free in that it discovers a path between p_1 and p_2 whenever one exists, and reports failure when there is none. One example of such a complex planner for general configuration spaces is due to Barraquand and Latombe [2]. Such an error-free planner may be extremely slow and may not be run to completion in practice. However, if even the complex planner cannot discover a path between two connected configurations, then we may as well assume that these points are disconnected (i.e., we can view connectivity between configurations as being defined by the ability of the complex planner to find connections). Because of its expense, we seek to use this complex planner sparingly. As we will show, with high probability the preprocessing will ensure that only the simple planner is needed for answering queries. Our randomized preprocessing scheme may be summarized as follows:

1. **[Sampling]** Pick a random set of points in the free space. Call these points *milestones*.
2. **[Simple Permeation]** Try to connect all pairs of milestones using the simple planner.
3. **[Resampling]** For any milestones that are connected to relatively few others in this process, pick additional milestones "near" them at random.
4. **[Complex Permeation]** As a last resort, try using the complex planner to connect some pairs of milestones.

Step 4 is seldom used in practice, and would ideally be eliminated. In certain settings in practice this elimination may be possible with resampling and other related techniques.

The result of this preprocessing may be viewed as a graph G each of whose vertices corresponds to a milestone, with an edge signifying that its end-points are in the same component of free space. This graph is sometimes called a *probabilistic roadmap* [16].

Given a query pair of configurations q_1 and q_2 in free space, we detect whether it is possible to move from q_1 to q_2 as follows: we use the simple planner to connect q_1 and q_2 to milestones m_1 and m_2 respectively. We then use a graph search algorithm to determine whether the milestones m_1 and m_2 are in the same connected component of the roadmap G. Queries are never answered incorrectly; with some probability though, the query processing algorithm may fail to give an answer.

In our analysis, we assume that the configuration space is available as a membership oracle: given a point p in the configuration space, we can decide whether or not the point is in free space. This is reasonable in implementations [15, 19]: such a membership test corresponds to checking whether a configuration violates any of the constraints in the input, and this can be done rather efficiently. We treat the simple planner (denoted B_S) and the complex planner (B_C) as black-boxes. We assume without loss of generality that both planners are *reflexive:* i.e., if a planner succeeds in connecting p_1 to p_2, it can also connect p_2 to p_1. (A non-reflexive planner can be made reflexive by applying it to both directions — from p_1 to p_2, and from p_2 to p_1 — simultaneously.)

A word about the random sampling in Step 1 of the preprocessing: in the experimental work [14, 15, 16, 28] this is done simply by choosing a point at random from $[0,1]^d$. If the chosen point is in the free space, it is retained; else it is discarded and the process repeated. Clearly a point chosen at random in this fashion is uniformly distributed in the free space, but in order for the number of repetitions to be reasonably small we need the free space to constitute a good fraction of the configuration space. We assume this is the case based on empirical evidence (else no analysis is possible). Choosing a random sample has a minuscule cost in practice compared with the other operations, and can be repeated a very large number of times if necessary (see also Section 5).

Our main thesis is that the empirically observed success of the scheme stems from a property we call ϵ-goodness which we now define. Let $\mathcal{F}$ denote the free space. For a point $p \in \mathcal{F}$, let $S(p)$ consist of those points of $\mathcal{F}$ that can be connected to p by the simple planner B_S. For a subset $\mathcal{X}$ of the configuration space, let $\mu(\mathcal{X})$ denote its volume.

Definition 1.1 *Let ϵ be a positive real. We say that a point p in the free space $\mathcal{F}$ is ϵ-good if $\mu(S(p)) \geq \epsilon\mu(\mathcal{F})$. We say that the free space $\mathcal{F}$ is ϵ-good if for all points $p \in \mathcal{F}$ we have $\mu(S(p)) \geq \epsilon\mu(\mathcal{F})$.*

While any non-degenerate configuration space is ϵ-good for some positive ϵ, the intent in this definition is that the space be ϵ-good for a "reasonably large" value of ϵ. Many configuration spaces arising in practice do not have the ϵ-good property; for example, consider a crescent-shaped region or one where a circular obstacle is tangential to a rectangular obstacle. However, in these cases, our definition applies to the subset of free space obtained by removing a small neighborhood of the cusp or tangency points from the configuration space. (See Section 5.2 for a more rigorous treatment of this issue.)

1.3 Overview of Results

The first contribution of this paper is a model of computation appropriate for the analysis of the probabilistic roadmap scheme, taking into account the realities of the problem at hand. In Section 2 we define a concrete algorithm based on the high-level outline given above. This algorithm and its analysis do not make use of resampling (Step 3 above); we present this simplified version first because it succinctly outlines the main ideas using only the simple notion of ϵ-goodness. We argue in Section 3 that if the free space is ϵ-good then every point of the free space $\mathcal{F}$ can, with high probability, be connected to a milestone using only B_S. In Section 4 we give a bound on the number of invocations of the complex planner B_C in constructing the probabilistic roadmap; this involves a new randomized algorithm for determining connected components in a model related to the decision tree model used in the study of *evasive* graph properties [21], and may be of independent interest. We complement this with tight bounds for deterministic algorithms. These results imply bounds on the work done in preprocessing and in query processing, in terms of the number of times B_C and B_S are invoked; in particular, the complex planner is not used for answering queries. Section 5

summarizes results from experiments with the robot arm of Figure 1; these suggest that most but not all points in the corresponding free space are ϵ-good for a reasonably large value of ϵ. Interestingly, the resampling step seems to be helpful for settings such as this arm. We therefore extend (Section 5.2) the definition of ϵ-goodness and use it to explain these observations: assuming the configuration space satisfies a weaker condition we call (ϵ, t)-goodness for a small integer t, we give an explanation for the resampling step similar to the analysis in Sections 3 and 4. Finally, our work is related to classic problems in art-gallery theorems. In Section 6 we establish this connection, give some new results related to our work, and mention some resulting open problems in art-gallery theorems.

2 Algorithms and Results

For the remainder of the paper, we say that two points $p_1, p_2 \in \mathcal{F}$ are mutually *visible* when B_S can connect p_1 and p_2. We use this terminology primarily for brevity, and our usage is inspired by a commonly used simple planner [16, 15] that checks whether the straight line segment joining p_1 and p_2 is in $\mathcal{F}$ (equivalently, p_1 and p_2 are mutually visible in $\mathcal{F}$); however, our entire analysis works for any simple planner B_S.

Let $\beta \in (0, 1]$ be a positive real constant which represents the failure probability we can tolerate in the preprocessing (this will become clear in the statements of Theorems 2.1, 2.2 and 2.3). Let $s = (c/\epsilon)(\ln 1/\epsilon + \ln 4/\beta)$, where c is a fixed positive constant large enough that for any $x \in (0, 1]$, $(1-x)^{(c/x)(\ln 1/x + \ln 4/\beta)} \leq x\beta/4$; clearly, $c = O(1)$ suffices. The algorithm for preprocessing is listed in Figure 2.

1. Pick s points in $\mathcal{F}$ at random, and call these milestones.
2. Invoke B_S on every pair of milestones.
3. Pick a representative milestone from each component that results. Let V be the set of these representatives and $|V| = n$.
4. Invoke the **Randomized Permeation** algorithm (Figure 5) on these representatives.

Figure 2: *The Preprocessing Algorithm*

As we will see in Section 4, Step 4 probes the "edge-slots" of the roadmap, trying to determine the structure of the connected components without expending too many calls to B_C. Note that the algorithm in Figure 2 does not make use of resampling; we will get to this in Section 5. In practice Step 4 is a last resort; much if not all of the connectivity information should have been discovered before this step. Step 4 is the only preprocessing step in which the B_C is invoked; this will become clear in Figures 4 and 5.

The query processing algorithm is listed in Figure 3. Given the query points q_1 and q_2, we connect them to milestones m_1 and m_2 using B_S as in Figure 3. Here $\gamma \in (0, 1]$ is the allowable failure probability for a query. For each i, Step 1a can be implemented using s invocations of B_S, one for each milestone. Each trial of Step 1b can be implemented using s invocations of B_S.

For an ϵ-good free space $\mathcal{F}$ call a set of milestones M *adequate* if the volume of the subset of $\mathcal{F}$ not visible from any milestone of M is at most $(\epsilon/2)\mu(\mathcal{F})$. Intuitively, if we were to place a point source of light at each milestone, we would like a fraction at least $1 - \epsilon/2$ of $\mathcal{F}$ to be illuminated.

1. **For** $i = 1, 2$ **do**:
   (a) **If** q_i can see a milestone v, set $m_i = v$.
   (b) **Else Repeat** $\log(2/\gamma)$ times:
      i. Choose v_i uniformly at random from $S(q_i)$;
      ii. **If** a milestone is visible from v_i **then** set m_i to be that milestone.
   (c) **If** all $\log(2/\gamma)$ trials fail **then** declare FAILURE and **halt.**
2. **If** m_1 and m_2 are in the same component of G **then** output YES **else** output NO.

Figure 3: *The Query Processing Algorithm*

Note that as ϵ increases, the requirement for adequacy grows weaker but the number of milestones needed becomes smaller.

Theorem 2.1 *The preprocessing stage will generate an adequate set of milestones with probability at least* $1 - \beta$.

Theorem 2.1 only says that most of $\mathcal{F}$ is likely to be visible from some milestone in M; using this property alone, we can show that queries can be answered *quickly.* However, we need a stronger property — which we may think of as *permeation* — to guarantee that queries can be answered *correctly.* Permeation is essentially the following: for any two milestones in the same connected region of $\mathcal{F}$, we can infer this connectedness from the preprocessing algorithm. Theoretically, we cannot hope to show that the use of B_S alone will provide such permeation: if $\mathcal{F}$ consists of two spheres each of diameter $1/2$ and the spheres touch at a single point p, we have a free space that is ϵ-good for $\epsilon = 0.5$. Yet it is extremely unlikely that B_S can yield permeation in this case (if for instance B_S simply checks visibility between milestones). In such configuration spaces, the use of the complex planner B_C in Step 4 is inevitable to ensure a good overall success probability. Define a function $g()$ on an ordered k-tuple of positive integers $n_1, n_2, \ldots, n_k$ by $g(n_1, n_2, \ldots, n_k) = \sum_{i=1}^{k} i n_i$.

Theorem 2.2 *Let S be a set of n milestones lying in k connected components denoted $S_1, \ldots, S_k$ such that $|S_1| \geq |S_2| \geq \ldots \geq |S_k|$. The preprocessing stage will determine the partition correctly and the expected number of invocations of B_C is at most*

$$2g(|S_1|, |S_2|, \ldots, |S_k|).$$

which is $O(nk)$ in the worst case. With high probability, the number of invocations of B_C is within $O(\log n)$ of its expectation.

Theorem 2.3 *Suppose that the set of milestones chosen during preprocessing is adequate. Then the probability that the query processing algorithm outputs* FAILURE *is at most γ. When the query processing algorithm does not output* FAILURE*, it correctly answers the query by either producing a path or declaring that none exists.*

In fact, our analysis will imply that the expected number of executions of Step 1b in the query processing algorithm (Figure 3) is at most 2.

3 Nearly Complete Coverage

This section establishes Theorems 2.1 and 2.3. The expectation of the volume of points not visible from any of the s randomly chosen milestones in M is

$$\mathbf{E}[\mu(\{p \in \mathcal{F} \mid p \notin \cup_{m \in M} S(m)\})] = \int_{p \in \mathcal{F}} \mathbf{Pr}[p \notin \cup_{m \in M} S(m)]. \quad (1)$$

The probability that a fixed point is not visible from any of the s milestones is at most $(1-\epsilon)^s$. Thus, the above is bounded by

$$\int_{p \in \mathcal{F}} (1-\epsilon)^s = \mu(\mathcal{F})(1-\epsilon)^s \leq \mu(\mathcal{F})\epsilon\beta/4, \quad (2)$$

By the Markov inequality, it follows that

$$\mathbf{Pr}[\mu(\{p \in \mathcal{F} \mid p \notin \cup_{m \in M} S(m)\}) > \mu(\mathcal{F})\epsilon/2] \leq \beta/2.$$

Thus with probability $1 - \beta/2$ the "shadow region" not visible from any $m \in M$ has volume at most $\mu(\mathcal{F})\epsilon/2$, in which case it follows that for any $p \in \mathcal{F}$, the volume of the subset of $S(p)$ visible from some $m \in M$ is at least $\mu(S(p)) - \mu(\mathcal{F})\epsilon/2 \geq \mu(\mathcal{F})\epsilon/2$.

This establishes Theorem 2.1 and leads to Theorem 2.3: for either query point q_i, the probability that a random point chosen from $S(q_i)$ is not visible from any $m \in M$ is $(\epsilon/2)/S(q_i) < 1/2$. The probability that we fail on $\log(2/\gamma)$ trials is less than $\gamma/2$. Since we do this for the two query points, the overall failure probability is at most γ.

4 Permeation

This section establishes Theorem 2.2. En route, we connect our problem to the decision tree model used to study evasive graph properties, and prove some related results. The permeation problem is the following: given a free space $\mathcal{F}$ containing $n \leq s$ milestones, determine which milestones are reachable from each other. (Note that because of Step 2 in the Preprocessing Algorithm of Section 2, n may be much smaller than s.) Given any pair of milestones the complex planner B_C will decide whether they are connected. The graph G can be computed with $O(n^2)$ invocations of B_C by trying it on every pair of points, but we show that far fewer invocations may suffice.

We work with the following abstract version of the permeation problem. The input is a graph $G(V, E)$ with n vertices, consisting of k disjoint cliques. The goal is to determine this clique partition of G. The cost of an algorithm is measured by the number of entries it examines in the adjacency matrix of G. This is the *edge probe model* used in the study of *evasive graph properties* [21].

The vertices correspond to the points in $\mathcal{F}$, and an edge is present between two vertices if the corresponding points are connected, i.e., lie in the same component of $\mathcal{F}$. Since the complex planner is error-free, we obtain that the points in a given component of $\mathcal{F}$ form a clique in G, and that there are no edges between two distinct cliques. A probe into the adjacency matrix corresponds to an invocation of B_S, and in this abstract version we do not distinguish between deciding the presence of an edge and actually searching for a path in $\mathcal{F}$ but it is easy to see that there is no loss of generality in this.

Let $N(n, K)$ denote the non-deterministic complexity of this problem. A non-deterministic algorithm is only required to *verify* that some partition into k cliques is the right partition.

Theorem 4.1 *For* $1 \leq k \leq n$, $N(n, k) = \Theta(n + k^2)$.

1. Mark all vertices in V as being LIVE.
2. Initialize $x \leftarrow 1$.
3. **While** $x \leq n$ **do:**
   (a) $\Gamma(x) \leftarrow \emptyset$.
   (b) **For** $y = x + 1$ to n **do:**
      i. **If** vertex y is marked LIVE
      then probe the edge (x, y) in G.
      ii. **If** edge (x, y) is probed and found present
      then mark y as DEAD and add y to $\Gamma(x)$.
   (c) Output $\{x\} \cup \Gamma(x)$ as being a clique.
   (d) Mark x as being DEAD.
   (e) Set x to the smallest numbered LIVE vertex, or $n + 1$ if there are no LIVE vertices left.

Figure 4: *The Deterministic Permeation Algorithm*

Proof: The algorithm must make at least one probe on each of the n vertices. It must also have verified that each of the $\binom{k}{2}$ pairs of cliques is in fact disconnected. □

We now characterize the worst-case *deterministic* complexity of this problem, denoted $T(n, k)$. Consider the following deterministic algorithm: by probing all edge slots incident on an arbitrary vertex x, determine the neighborhood of x, say $\Gamma(x)$; let $C_x = \{x\} \cup \Gamma(x)$, and output C_x; then, recur on the vertex-induced subgraph $G[V \setminus C_x]$. The proof of correctness is obvious, so we focus on the analysis of the running time. The number of levels in the recursion is k, since one of the k cliques is removed from G prior to each recursive call. The number of probes made in the process of determining each such clique is at most n. The total number of probes is $O(nk)$. In Figure 4, we illustrate the Deterministic Permeation Algorithm, which is an iterative version of the recursive algorithm. The iterative version will prove useful when describing a randomized algorithm. By the preceding discussion, we have:

Theorem 4.2 *The Deterministic Permeation Algorithm correctly solves the permeation problem using $O(nk)$ probes.*

The following lower bound establishes that the Deterministic Permeation Algorithm is optimal.

Theorem 4.3 *For $1 \leq k \leq n$, $T(n, k) = \Omega(nk)$.*

Proof: We present an adversary argument in terms of the complementary problem: given a graph $\overline{G}$ which is a complete k-partite graph for some k, determine the k-partition of the vertices of $\overline{G}$ into independent sets. The adversary responds to each probe for an edge by some deterministic algorithm, and its strategy is to say that edges are present, as far as possible. The adversary chooses a value k initially, and ensures that the graph it constructs (adaptively) is a complete K-partite graph for some $K \in \{k - 1, k, k + 1\}$. The algorithm can be provided this information without affecting the following argument.

The adversary maintains a graph H in which the edges are those edges of $\overline{G}$ which have been probed already *and* for which the response was that the edge is present. When the adversary is

forced to concede that an edge (i, j) is absent in $\overline{G}$, it then *collapses* the two vertices i and j into a single meta-vertex whose neighborhood is the union of the neighbors of i and j. Collapsing two vertices is equivalent to conceding that they are in the same independent set of the k-partition; meta-vertices can also be collapsed into each other. The missing edges in H correspond to edge slots in G that have not been probed so far.

Any probe involving an edge (i, j), where i is contained in a meta-vertex i^* obtained by some earlier collapses, will be treated as referring to the edge (i^*, j) since all vertices in i^* have exactly the same set of neighbors. The adversary can reveal this graph H together with the meta-vertex structure to the algorithm without affecting the lower bound argument, and so we can assume that the algorithm never makes redundant queries such as probing for an edge between two vertices which belong to the same meta-vertex.

The adversary maintains the following invariants at all times.

1. The chromatic number of H is k; in particular, it maintains a partition of the (meta)-vertices into k non-empty color classes $C_1, \ldots, C_k$ such that each color class is an independent set. By the definition of H, none of the edges between the (meta-)vertices in a color class have been probed yet, and all edges that were probed and deemed present are between two distinct color classes.

2. For each meta-vertex, every vertex therein has had at least $k-1$ incident edges already probed that were deemed to be present in $\overline{G}$.

Initially, the adversary arbitrarily partitions the vertices into k non-empty color classes; since H is empty then, this ensures the first invariant. The second invariant holds trivially since there are no meta-vertices at the beginning.

Thereafter, the adversary responds as follows to a probe (i, j) by the algorithm. Note that a probe involving an edge (i, j), where i is contained in a meta-vertex i^*, will be treated as referring to the edge (i^*, j).

- If i and j belong to distinct color classes, it will say that the edge is present and will add this edge to the graph H.

- If i and j belong to the same class C_r, then it will check to see if there exists a color class C_t with $t \neq r$ such that at least one of i and j does not have neighbors in C_t. Suppose that i does not have any neighbors in C_t, then the adversary will transfer i from C_r to C_t and will then respond as in the previous case (i.e., say that the (i, j) edge is present).

- Finally, there is the case where both i and j belong to the same component C_r and each has at least one neighbor in every other color class. In this case, the adversary will concede that the edge (i, j) is indeed absent and will then collapse i and j together.

The first invariant holds since edges are only introduced between vertices in distinct color classes. The color classes remain non-empty since a vertex is transferred from a color class only when it has at least two vertices. To verify the second invariant, observe that when two vertices (i, j) are collapsed, both have at least one neighbor in the remaining $k - 1$ color classes.

The algorithm can terminate only when the number of (meta)-vertices in each color class is down to one, and there is an edge between each pair of color classes, since otherwise the algorithm cannot be certain of the k-partition of $\overline{G}$, or even whether there is a k-partition in the first place.

We claim that, upon termination, every one of the n vertices must have at least $k - 1$ edges incident on it which were probed and deemed to be present in $\overline{G}$. The second invariant implies that

1. Permute the vertices randomly so that each is labeled by an integer in $\{1, \ldots, n\}$.
2. Invoke the Deterministic Permeation Algorithm.

Figure 5: *The Randomized Permeation Algorithm*

this is true for any vertex which participated in a collapse and is a part of some meta-vertex when the algorithm terminates. A vertex which did not participate in any collapse must also have at least $k-1$ edges incident on it since it is the only vertex in its color class, and there is an edge from its color class to every other class. Thus, the total number of edges probed and deemed present in $\overline{G}$ is at least $n(k-1)/2$. Also, there must be at least $n-k$ edges which were probed and deemed absent in $\overline{G}$, since in going from n vertices to k vertices at least $n-k$ collapses need to be performed and each collapse requires a distinct absent edge. Thus, the total number of probes must be $\Omega(nk)$. □

We now give a randomized algorithm that beats the lower bound of Theorem 4.3 when the sizes of the k cliques differ significantly. This is crucial in our application to motion planning because in practice the free space $\mathcal{F}$ often consists of components of one large component and several small ones. The Randomized Permeation Algorithm (see Figure 5) labels the vertices in a random order and then invokes the Deterministic Permeation Algorithm.

Let $w_1 \geq w_2 \geq \cdots \geq w_k$ be the sizes of the cliques in an instance G arranged in a non-increasing order, where $n = \sum_{i=1}^{k} w_i$. Denote by C_i the ith largest clique in G.

Theorem 4.4 *The Randomized Permeation Algorithm correctly determines the clique structure and incurs an expected cost that is at most*

$$2g(w_1, w_2, \ldots, w_k) - n - k.$$

Furthermore, with high probability, the cost is at most

$$O(g(w_1, w_2, \ldots, w_k) \log n).$$

Remark: Observe that the worst case is when all w_i are equal to n/k, in which case the expected cost is $O(nk)$. On the other hand when there is one giant clique and $k-1$ cliques of size $O(1)$ the expected cost is $\Theta(n + k^2)$, which is essentially the non-deterministic lower bound.

Proof: The proof of correctness follows from that for the Deterministic Permeation algorithm. We start with the analysis of the expected cost.

We say that a clique C_i *beats* another clique C_j if *some* vertex of C_i occurs before *all* vertices of C_j in the random permutation chosen by the Random Permeation Algorithm. The probability that C_i beats C_j is the same as the probability that a uniformly random choice from $C_i \cup C_j$ yields a vertex of C_i, and, clearly, the latter is $w_i/(w_i + w_j)$.

We divide the edge slots of the graph into intra-clique and inter-clique edge slots. For each i, the number of intra-clique edge slots in C_i that are probed is precisely $w_i - 1$, since the only such probes are between the earliest vertex (according to the random permutation) of C_i and the remaining vertices of C_i. The total number of such probes is

$$\sum_{i=1}^{k} (w_i - 1) = n - k.$$

Fix some i and j, and suppose that C_i beats C_j. The inter-clique edge slots between these two cliques are between the earliest vertex of C_i and all vertices of C_j. This gives a total of w_j probes that are "charged" to C_j (the beaten clique). The expected total charge to clique C_j is given by

$$\sum_{i \neq j} \frac{w_i}{w_i + w_j} \times w_j.$$

To bound the expected total number of inter-clique edge slots that are probed, we sum the charges to the various cliques and obtain

$$\begin{aligned}
\sum_{j=1}^{k} \sum_{i \neq j} \frac{w_i w_j}{w_i + w_j} &= \sum_{j=1}^{k} \sum_{i<j} \frac{2 w_i w_j}{w_i + w_j} \\
&\leq \sum_{j=1}^{k} \sum_{i<j} 2 w_j \\
&= 2 \sum_{j=1}^{k} (j-1) w_j \\
&= 2g(w_1, \ldots, w_k) - 2n.
\end{aligned}$$

Adding together the bounds on the expected number of intra- and inter-clique edge-slots that are probed, we obtain the desired bound.

We now turn to the task of proving the high probability bound. Fix a clique C_j and note that the total charge to C_j is the size of C_j multiplied by the number of other cliques that beat it. Since there are $j-1$ cliques that are larger than C_j, at most $j-1$ of the cliques that beat C_j are larger than C_j. Let X_j be the random variable denoting the number of cliques *smaller* than C_j that beat C_j; let Y_j be the random variable denoting the total number of vertices from cliques smaller than C_j that are earlier than *all* vertices in C_j; and, finally, let Z_j be a random variable having the geometric distribution with parameter $p_j = w_j / \sum_{i=j}^{k} w_i$ and expectation $1/p_j$. Clearly, $X_j \leq Y_j$, and Y_j is stochastically dominated by Z_j. The probability that Z_j is larger than $2p_j^{-1} \ln n$ is bounded by

$$(1 - p_j)^{2p_j^{-1} \ln n} \leq e^{-2 \ln n} = \frac{1}{n^2}.$$

Thus with probability at least $1 - 1/n$, we have, for each j, $X_j \leq 2p_j^{-1} \ln n$. This implies that, with high probability, the total number of inter-clique edges probed is given by

$$\begin{aligned}
\sum_{j=1}^{k} (j - 1 + X_j) w_j &\leq \sum_{j=1}^{k} (j - 1 + 2p_j^{-1} \ln n) w_j \\
&\leq \sum_{j=1}^{k} ((j-1) w_j + 2 \ln n \sum_{i=j}^{k} w_i) \\
&= \sum_{j=1}^{k} j w_j - n + 2 \ln n \sum_{j=1}^{k} j w_j \\
&= (1 + 2 \ln n) g(w_1, \ldots, w_k) - n.
\end{aligned}$$

Adding in the number of intra-clique edge slots that are probed, we obtain the desired result. □

5 Experiments and Extensions

5.1 The Robot Arm Example

The robot arm of Figure 1 was tested for ϵ-goodness using 9000 random samples; on a DEC Alpha workstation, it took 9.24 seconds to create the random configurations, and 1399 seconds* to try connecting all pairs using B_S. (These figures underscore that random sampling is not a significant component of the cost.) The samples with the "most" visibility could see about 0.06 (i.e., 6%) of the remaining samples, suggesting that they are 0.06-good. As many as 3.3% of the random samples could see no other random samples, and fully 22% could see 0.001 (i.e., 0.1%) or less; in other words, only about 78% of the configuration space is 0.001-good or better. (For $\epsilon = 0.001$, we have $(1/\epsilon) \ln 1/\epsilon = 6908$, which is of the same order as our number of samples.) We conjecture that the resampling step (Step 3 from our high-level outline of Section 1) leads to better coverage of the space in situations such as Figure 1. We have observed that it helps eliminate the need for the Complex Permeation step of the outline of Section 1 in some examples. To address this we introduce a generalization of the notion of ϵ-goodness.

5.2 The Extended Definition

Let us say that a point p in free space is $(\epsilon, 1)$-good if $\mu(S(p)) \geq \epsilon\mu(\mathcal{F})$, corresponding to our original definition of ϵ-goodness for a point. Next, we say a point p in free space is (ϵ, t)-good if $\mu(\{q \in S(p) \mid q \text{ is } (\epsilon, t-1)\text{-good}\}) \geq \mu(S(p))/2$. For $t > 1$, we say that $\mathcal{F}$ is (ϵ, t)-good if $\mu(\{p \in \mathcal{F} | p \text{ is } (\epsilon, 1)\text{-good}\}) \geq \mu(\mathcal{F})/2$ and every point of $\mathcal{F}$ is (ϵ, i)-good for $i \leq t$. If $\mathcal{F}$ is (ϵ, t)-good for a small value of t, we can give a theoretical basis for the resampling step (Step 3 in the outline of Section 1). The main idea is that single links discovered by B_S in the algorithms of Section 2 are now simulated using t-link paths found by resampling and connecting using B_S. This leads to a generalized definition of an adequate set of milestones, and eventually to a version of Theorem 2.3 in which the number of invocations of B_S is larger by a factor of 2^t. This extension requires that we can still sample the visibility region of a query point. In practice, this is often accomplished by defining an appropriate "neighborhood" for any point p, from which a sample likely to be in $S(p)$ can be chosen.

6 Related Combinatorial Results

A number of combinatorial problems concerning art-gallery theorems [26] are related to our work. For instance, given a simple polygon that is ϵ-good we ask: how many guards are necessary and sufficient to cover the entire polygon? (Another way of thinking of this is to imagine point sources of light being placed in the polygon with the objective of illuminating the entire interior.)

The following would be an ideal result: given an ϵ-good configuration space S, a random sample of $poly(1/\epsilon)$ points from the free space $\mathcal{F}$ will "illuminate" the entire free space with high probability. In practice it may be reasonable to assume that the number of "holes" in the free space ω is "small" (for instance, bounded by a slowly growing function of the input size).

*In implementations used in practice, several additional techniques offer substantial savings over the timings reported here. For instance, we dynamically update a representation of the connected components after testing each pair of configurations. Thus we would not test a new pair if they are known to belong to the same connected component.

Conjecture 6.1 *A random sample of $poly(\omega + 1/\epsilon)$ points is likely to cover an ϵ-good free space with ω holes.*

At present we only have the most rudimentary results of this type; for instance, we give an upper bound on ϵ so that one guard suffices to cover an ϵ-good simply-connected region. In fact, a Helly-type theorem due to Krasnosselsky [20] immediately yields:

Theorem 6.1 *Let R be a compact, simply-connected ϵ-good region in Euclidean d-space for $\epsilon > d/(d+1)$. Then there is a point p in R such that $S(p) = R$.*

Broder, Dyer, Frieze, Raghavan and Upfal [3] have initiated progress in extending the above result: they show that if $\epsilon = 1/3 + \delta$ for a simply-connected region in the plane, the number of guards is polynomial in $1/\delta$; they also give a bound for arbitrary spaces in which the number of guards grows with the diameter of the region. Kalai and Matoušek [12] have shown that for simply connected regions in the plane, $O(\frac{1}{\epsilon}\log\frac{1}{\epsilon})$ guards suffice. Various interesting questions remain. For instance, even the existential version of Conjecture 6.1 would be useful: given an ϵ-good space $\mathcal{F}$ with ω holes, there exists a set of $poly(\omega + 1/\epsilon)$ points which covers $\mathcal{F}$.

Acknowledgement

We thank Don Knuth for helpful discussions. We are also grateful to David Hsu for pointing out an error in an earlier version of the paper.

References

[1] J. Barraquand, B. Langois and J.-C. Latombe. Potential Fields for Robot motion planning. *IEEE Transactions on Systems, Man and Cybernetics*, 22(2), March 1992, 224–241.

[2] J. Barraquand and J.-C. Latombe. Robot motion planning: A distributed representation approach. *International Journal of Robotics Research*, 10:628–649, 1991.

[3] A.Z. Broder, M.E. Dyer, A.M. Frieze, P. Raghavan and E. Upfal. Private communication, 1995.

[4] J.F. Canny. *The Complexity of Robot Motion Planning.* MIT Press, 1988.

[5] J.F. Canny and J.H. Reif. New Lower Bound Techniques for Robot Motion Planning Problems. In *Proceedings of the 28th IEEE Symposium on Foundations of Computer Science*, pp. 49-60, Los Angeles, CA, 1987.

[6] J.E. Hopcroft, D.A. Joseph, and S.H. Whitesides. Movement Problems for 2-Dimensional Linkages. *SIAM Journal of Computing*, 13(3):610–629, 1984.

[7] J.E. Hopcroft, D.A. Joseph, and S.H. Whitesides. On the Movement of Robot Arms in 2-Dimensional Bounded Regions. *SIAM Journal of Computing*, 14(2):315–333, 1985.

[8] J.E. Hopcroft, J.T. Schwartz, and M. Sharir. On the Complexity of Motion Planning for Multiple Independent Objects: PSPACE-hardness of the 'Warehouseman's Problem'. *International Journal of Robotics Research*, 3(4):76–88, 1984.

[9] J.E. Hopcroft and G.T. Wilfong. Motions of Objects in Contact. *SIAM Journal of Computing*, 15(3):768–785, 1984.

[10] Th. Horsch, F. Schwarz, and H. Tolle. Motion planning for many degrees of freedom — random reflections at c-space obstacles. In *Proceedings of the IEEE International Conference on Robotics and Automation*, pp. 3318–3323, San Diego, CA, 1994.

[11] D.A. Joseph and W.H. Plantiga. On the Complexity of Reachability and Motion Planning Questions. In *Proceedings of the First ACM Symposium on Computational Geometry*, pp. 62–66, Baltimore, MD, 1985.

[12] G Kalai and J. Matoušek. Guarding galleries in which every point can see a large area. Unpublished manuscript, 1995.

[13] V. Kantabutra and S.R. Kosaraju. New Algorithms for Multilink Robot Arms. *Journal of Computer and System Science*, 32(1):136–153.

[14] L. Kavraki and J.-C. Latombe. Randomized preprocessing of configuration space for fast path planning. In *Proceedings of the IEEE International Conference on Robotics and Automation*, pp. 2138–2145, San Diego, CA, 1994.

[15] L. Kavraki. *Random Networks in Configuration Space for Fast Path Planning.* PhD Thesis, Stanford University, 1995. Technical Report STAN-CS-TR-95-1535. Department of Computer Science, Stanford University, 1995.

[16] L. Kavraki, P. Švestka, J.-C. Latombe, and M. Overmars. Probabilistic roadmaps for path planning in high dimensional configuration spaces. Technical Report STAN-CS-TR-94-1519, Department of Computer Science, Stanford University, 1994.

[17] O. Khatib. Real-Time Obstacle Avoidance for Manipulators and Mobile Robots. *International Journal of Robotics Research*, 5(1):90–98, 1986.

[18] Y. Koga. On computing multi-arm manipulation trajectories. Technical Report STAN-CS-TR-94-1530, Stanford University, 1994.

[19] J.-C. Latombe. *Robot Motion Planning.* Kluwer Academic Publishers, 1991.

[20] S.R. Lay. *Convex Sets and their Applications.* John Wiley, 1982.

[21] L. Lovász and N. Young. *Lecture notes on evasiveness of graph properties.* Technical Report CS-TR-317-91, Computer Science Department, Princeton University, 1991.

[22] T. Lozano-Pérez. Automatic Planning of Manipulator Transfer Movements. *IEEE Transactions on Systems, Man and Cybernetics*, SMC-11(10):681–698, 1981.

[23] T. Lozano-Pérez. Spatial planning: a configuration space approach. *IEEE Transactions on Computers*, 32:108–120, 1983.

[24] T. Lozano-Pérez and M.A. Wesley. An Algorithm for Planning Collision-Free Paths Among Polyhedral Obstacles. *Communications of the ACM*, 22(10):560–570, 1979.

[25] N.J. Nilsson. A Mobile Automaton: An Application of Artificial Intelligence Techniques. In *Proceedings of the 1st International Joint Conference on Artificial Intelligence*, pp. 509–520, Washington, DC, 1969.

[26] J. O'Rourke. *Art Gallery Theorems and Algorithms.* Oxford University Press, 1987.

[27] M. Overmars and P. Švestka. A probabilistic learning approach to motion planning. In *Proceedings of the Workshop on Algorithmic Foundations of Robotics*, 1994.

[28] M. Overmars. A Random Approach to Motion Planning. Technical Report RUU-CS-92-32, Department of Computer Science, Utrecht University, 1992.

[29] J.H. Reif. Complexity of the Mover's Problem and Generalizations. In *Proceedings of the 20th IEEE Symposium on Foundations of Computer Science*, pp. 421–427, 1979.

[30] J.H. Reif and M. Sharir. Motion Planning in the Presence of Moving Obstacles. In *Proceedings of the 25th IEEE Symposium on Foundations of Computer Science*, pp. 144–154, 1985.

[31] J.T. Schwartz and M. Sharir. On the 'Piano Movers' Problem: II. General Techniques for Computing Topological Properties of Real Algebraic Manifolds. *Advances in Applied Mathematics*, 4:298–351, 1983.

[32] S. Udupa. *Collision Detection and Avoidance in Computer Controlled Manipulators.* PhD thesis, Department of Electrical Engineering, California Institute of Technology, 1977.