

Sparse Roadmap Spanners

Andrew Dobson, Athanasios Krontiris and Kostas E. Bekris

Abstract Asymptotically optimal planners, such as PRM^* , guarantee that solutions approach optimal as iterations increase. Roadmaps with this property, however, may grow too large. If optimality is relaxed, asymptotically near-optimal solutions produce sparser graphs by not including all edges. The idea stems from graph spanner algorithms, which produce sparse subgraphs that guarantee near-optimal paths. Existing asymptotically optimal and near-optimal planners, however, include all sampled configurations as roadmap nodes. Consequently, only infinite graphs have the desired properties. This work proposes an approach that provides the following asymptotic properties: (a) completeness, (b) near-optimality and (c) the probability of adding nodes to the spanner converges to zero as iterations increase. Thus, the method shows that finite-size data structures can have near-optimality properties. The method brings together ideas from various planners but deviates from existing integrations of PRM^* with graph spanners. Simulations for rigid bodies show that the method indeed provides small roadmaps and results in faster query resolution. The rate of node addition is shown to decrease over time and the quality of solutions satisfies the theoretical bounds. Smoothing provides a more favorable comparison against alternatives with regards to path length.

1 Introduction

Motivation: Sampling-based roadmaps preprocess a configuration space (\mathcal{C} -space) to answer multiple path planning queries online [11]. They are practical solutions for challenging, relatively high-dimensional instances. Traditionally, the focus has been on probabilistically complete methods that return quickly collision-free paths regardless of their quality. Recent progress has led to asymptotically optimal planners [9], which, however, need to eventually include all configurations as graph nodes to provide optimality. Many applications, however, require small roadmaps that still provide good quality paths. Resource constrained robots can better communicate, store and query smaller, sparse roadmaps that are computed offline. Interactive games have a small computational budget for path planning and need to quickly load and query a roadmap online [4]. In dynamic environments, small roadmaps with multiple good quality paths are advantageous when roadmap edges must be invalidated on the fly given moving obstacles [8]. This paper describes a method that simultaneously satisfies many of the above aims. It shows that by relaxing optimality to near-optimality, it is possible to return a finite-size roadmap that asymptotically converges to near-optimal solutions.

Sampling-based Roadmaps: The first popular method for building a roadmap using sampling was the Probabilistic Roadmap Method (PRM) [11]. The algorithm samples a point in \mathcal{C}_{free} , the collision-free \mathcal{C} , and adds it as a node. It then tries to

Computer Science and Engin., Univ. of Nevada, Reno, NV, e-mail: bekris@cse.unr.edu

connect it to the k -closest neighbors (k -PRM) or those within a δ -ball (δ -PRM). If the path is in C_{free} , an edge is added. Several variations exist and various adaptations to the path planning problem have been solved with PRM [2, 23, 21, 25, 24]. This success motivated work on the study of the probabilistic completeness guarantees PRM provides [10, 5, 6, 12]. Some variations focus on small roadmaps that provide coverage, connectivity or good path quality [25]. For instance, visibility-based roadmaps reject nodes if not needed for coverage or connectivity [23]. The Useful Cycles approach tests edges for their usefulness in terms of path quality [18] and combined with the Reachability Roadmap Method returns high clearance paths in 2D and 3D C-spaces [4]. Smoothing can improve solutions and algorithms exist that produce roadmaps with paths in all homotopic classes and thus deformable to optimal ones [7, 22]. Nevertheless, they build relatively dense roadmaps and smoothing can be expensive for online query resolution. Hybridization graphs combine multiple solutions into a higher quality one [20].

Roadmaps require a steering method that exactly connects two states, which is not available for many dynamical systems. Tree-based planners, such as RRT [13] and Expansive Spaces [6], can solve problems with dynamics and already return sparse graphs. Nevertheless, they do not provide the preprocessing properties of roadmaps for multiple queries and for this reason they are not the focus of this work. Tree-based planners can benefit from roadmaps that return distances between configurations given C-space obstacles [14]. RRT has been shown to converge to a suboptimal solution almost certainly [17, 9].

Asymptotic Optimality: Recent work has provided the conditions under which PRM is asymptotically optimal [9]. Asymptotic optimality implies that as more time is invested in constructing the roadmap, the quality of solutions improves and converges to optimal at infinity. The analysis indicates that the number of neighbors is the important variable that controls asymptotic optimality [9], as indicated in Fig. 1. A simple PRM that connects samples to neighbors within a δ -ball is asymptotically optimal, but results in a dense roadmap. The roadmap's density can be reduced by considering the k -nearest neighbors, but this version is not asymptotically optimal. The PRM* and k -PRM* algorithms rectify this by selecting the minimum number of neighbors required for asymptotic optimality, which is a logarithmic function of the number of nodes. Nevertheless, the number of neighbors still grows with each iteration and the resulting roadmap size is high, as all samples are added as nodes and it is not clear when to stop sampling.

Asymptotic Near-Optimality: A way to return sparser, good-quality roadmaps is to relax the optimality guarantees by utilizing graph spanners [19]. Spanners are subgraphs, where the shortest path between two nodes on the subgraph is no longer than t times the shortest path on the original graph, and t is the *stretch factor* of the spanner. Applying an efficient spanner [3] on the output of k -PRM* resulted in a Sequential Roadmap Spanner (SRS), which reduces the expected number of

algorithm	edges	optimal?
δ -PRM	$O(n^2)$	asympt.
k -PRM	$O(kn)$	no
PRM*	$O(n \log n)$	asympt.
k -PRM*	$O(n \log n)$	asympt.
SRS	$O(an^{1+\frac{1}{a}})$	asympt. near
IRS	$O(n \log n)$	asympt. near
IRS2 ($m < n$)	$O(m \log m)$	no

Fig. 1 PRM variations and asymptotic optimality. n, m : # nodes in the roadmap.

edges and provides asymptotic near-optimality, i.e., as more time is spent on constructing the roadmap, the quality of solutions converges to a value at most t times the optimal. An incremental integration of spanners with k -PRM* (IRS) has been experimentally shown to provide even better results [15]. Moreover, the path quality degradation in roadmap spanners is quite smaller in practice than the theoretical guarantees. These approaches, however, still include every sample as a roadmap node. Recent work proposed an extension of IRS, called IRS2, which has a similar objective to the current paper [16]. IRS2 applies the spanner criterion on nodes, i.e., a node is useful if it connects at least two other nodes with a new path that is at least t times shorter than the best alternative. The method introduces some of the notions that are important in the current work. Nevertheless, it is not possible to argue about asymptotic near-optimality by applying the spanner criterion on nodes. It is also not possible to show that the algorithm will converge to a finite graph.

Contribution: This paper presents the SParse Roadmap Spanner (SPARS) algorithm, which: (a) is probabilistically complete, (b) can connect any two query points with a path of length:

$$t \cdot c^* + 4 \cdot \Delta, \quad (1)$$

where t and Δ are input to the algorithm, c^* is the cost of the optimum path between the query points in \mathcal{C}_{free} of clearance cl , if one exists, and (c) converges to a finite-size roadmap, i.e., as time increases, the probability of adding new nodes and edges converges to 0.

SPARS can be seen as a solution to the problem of finding a compact representation for answering shortest-path queries in continuous spaces, which has been proposed as an important challenge for the motion planning community [1]. To the best of the author’s knowledge, this is the first work that argues for the existence and provides an algorithm for computing finite data structures with some form of near-optimality guarantee for continuous path planning. While SPARS provides only asymptotic properties, i.e., it requires an infinite amount of time to return the finite data structure, the method allows for a natural stopping criterion inspired by Visibility PRM [23]. The criterion relates to a probabilistic measure of how close the roadmap is to a solution that provides the desired properties. This stopping criterion is not obvious for approaches that include all the samples in the roadmap.

In order to compute the spanner, the method builds an asymptotically optimal denser graph. \mathcal{C} -space samples are included in the spanner if they are useful for coverage or connectivity purposes or if they improve path quality relative to paths on the dense graph. The algorithm departs from existing roadmap spanners. In particular, it is based on the identification of boundaries of “visibility” regions of the spanner nodes through the dense graph nodes. It eventually guarantees that all shortest paths between boundaries of “visibility” regions on the dense graph can be represented by paths on the spanner that are at most t times longer, which leads to Eq. 1. The parameters t and Δ can control the sparsity of the resulting roadmap.

Simulations with rigid bodies ($SE(2)$ and $SE(3)$) indicate that the method provides very small roadmaps and that the rate of nodes added to the spanner decreases over time, resulting in very efficient online query resolution times. Moreover, the path solution quality is significantly better than theoretical bounds and close to the output of asymptotically optimal planners.

2 Setup and Nomenclature

The C-space abstraction casts a robot’s position and orientation as a point q in a d -dimensional space. The collision-free part of \mathcal{C} is denoted as \mathcal{C}_{free} . This paper focuses on the C-spaces of planar and rigid body configurations ($SE(2)$, $SE(3)$), but SPARS is applicable if appropriate metric and sampling functions exist.

Definition 1 (The Path Planning Problem). Given the set of free configurations $\mathcal{C}_{free} \subset \mathcal{C}$, initial and goal points $q_{init}, q_{goal} \in \mathcal{C}_{free}$, find a continuous path $\pi \in \Pi = \{\rho | \rho : [0, 1] \rightarrow \mathcal{C}_{free}\}$, $\pi(0) = q_{init}$ and $\pi(1) = q_{goal}$.

Definition 2 (Robust Feasible Paths). A path planning instance $(\mathcal{C}_{free}, q_{init}, q_{goal})$ is robustly feasible if there is a cl -robust path that solves it, for a positive clearance $cl > 0$. A path $\pi \in \Pi$ is cl -robust, if π lies entirely in the cl -interior of \mathcal{C}_{free} .

This work aims towards a planner that provides a compact, finite-size data structure for answering shortest-path queries in continuous spaces. Towards this objective, a planner must be able to identify which C-space samples are not needed as roadmap nodes. The standard, graph-theoretic formulation of spanners does not allow for removing nodes. An implicit, exhaustive graph $G(V, E)$ over \mathcal{C}_{free} can be defined by taking all the elements of \mathcal{C}_{free} as nodes and collision free paths between them as edges. Then, a “roadmap spanner” is a subgraph $S(V_S \subset V, E_S \subset E)$ of this implicit, exhaustive graph of the continuous space with the following properties:

- All nodes in G are connected with a path in \mathcal{C}_{free} to a node on S (coverage).
- S has as many connected components as G (connectivity).
- All shortest paths on S are no longer than t times the corresponding shortest paths in G (spanner property).

The properties allow for (a) arbitrary query points to connect to the roadmap, (b) paths to exist between any query points through S that can be connected in \mathcal{C}_{free} , and (c) asymptotic near-optimality for query points that lie on S . If the spanner, however, does not contain all \mathcal{C}_{free} points, there is no guarantee regarding path cost. To account for connecting to and departing from the roadmap, the path quality guarantees need to be further relaxed.

Definition 3 (Asympt. Near-Optimality with Additive Cost). An algorithm is asymptotically near-optimal with additive cost if, for a path planning problem $(\mathcal{C}_{free}, q_{init}, q_{goal})$ and cost function $c : \Pi \rightarrow \mathbb{R}^{\geq 0}$ with a cl -robust optimal path of finite cost c^* , the probability it will find a path with cost $c \leq t \cdot c^* + \epsilon$, for a stretch factor $t \geq 1$ and additive error $\epsilon \geq 0$, converges to 1 as the iterations approach infinity.

This paper shows that finite-size data structures can provide the above property and presents an algorithm that asymptotically converges to such a data structure. To describe the algorithm the following definitions will be useful.

Definition 4 (Local Planner). Given $q, q' \in \mathcal{C}$, a *local planner* computes a *local path* $L(q, q')$ connecting both configurations q and q' in the absence of obstacles. A straight-line between q and q' in \mathcal{C} is often sufficient.

Definition 5 (Distance function). The space \mathcal{C} is endowed with a *distance function* $d(q, q')$ that returns distances between configurations in the absence of obstacles.

Definition 6 (Shortest Paths). A cl -robust *shortest path* in \mathcal{C}_{free} is denoted as $\pi_{\mathcal{C}}$.

The shortest path between two configurations computed through the roadmap spanner \mathcal{G}_S is denoted as π_S .

3 Algorithm

SPARS builds in parallel two graphs: (a) a “sparse” graph $S(V_S, E_S)$ with the desirable properties, returned upon termination, and (b) a “dense” graph $D(V_D, E_D)$, which instead of being the exhaustive graph G , corresponds to the output of δ -PRM and asymptotically stores optimal paths. The notion of sparsity in this work deviates from its standard use in graph theory and is used to denote that graph S contains a finite number of nodes, while D asymptotically includes all configurations as nodes.

SPARS operates by sampling configurations in C_{free} and adding them to D . The algorithm adaptively selects a subset of these samples as nodes in the roadmap spanner S (i.e., $V_S \subset V_D$). By construction, there will be no edge in E_S that will have length more than $2 \cdot \Delta$, where Δ is the visibility range of a spanner node and is provided as an input parameter to the algorithm. There will also be no edge in E_D with length more than δ , the maximum radius for neighborhoods in δ -PRM. Typically $\delta \ll \Delta$.

Definition 7 (Representatives). A spanner node $v \in V_S$ will be the *representative* $rep(q)$ of sample $q \in V_D$, if they can be connected with an obstacle-free path of length less than Δ and v is the closest to q node on the spanner with this property, i.e., $rep(q) = \{v \in V_S \text{ so that } L(v, q) \in C_{free}, d(v, q) < \Delta \text{ and } \operatorname{argmin}_{v \in V_S} d(v, q)\}$. The *visibility region* of a spanner node v given the set of spanner nodes V_S is defined as the set of collision-free configurations that have v as their representative. Figure 2 (left) offers an illustration of the visibility region notion.

Definition 8 (Interfaces and support). The *interface* between two spanner nodes v and v' , $i(v, v')$, is the shared boundary of their visibility regions as in Figure 2 (right). A sample $q \in V_D$ *supports* the interface $i(v, v')$ of its representative $v = rep(q)$, if there is a sample $q' \in V_D$ so that $L(q, q') \in E_D$ and $rep(q') = v'$, where $v' \neq v$.

Definition 9 (Midpoint). The *midpoint* between two spanner nodes v and v' along the local path $L(v, v')$ will be denoted as $m(v, v')$:

$$m(v, v') \in L(v, v') \text{ and } d(v, m(v, v')) = d(m(v, v'), v')$$

If the local path $L(v, v')$ is obstacle-free, then the midpoint $m(v, v')$ lies on the interface $i(v, v')$ between the two spanner nodes.




Fig. 2 (left) Visibility region of the representative v_i : the configurations that can be connected to v_i and have it as the closest node. (right) Two neighboring spanner nodes define an interface: the shared boundary of their visibility regions. q supports the interface $i(v, v')$, where v is the representative of q , if there is a configuration q' within the δ hyper-sphere of q that has a different representative ($rep(q') = v' \neq v$) and $L(q, q') \in C_{free}$.

Algorithm 1: SPARS($C_{free}, M, t, \Delta, \delta$)

```

1  $failures \leftarrow 0; V_D \leftarrow \emptyset; E_D \leftarrow \emptyset; V_S \leftarrow \emptyset; E_S \leftarrow \emptyset;$ 
2 while  $failures < M$  do
3    $q \leftarrow \text{Sample}(C_{free});$ 
4    $\text{Add\_Node\_}\delta\text{-PRM}(D, q, \delta);$ 
5    $E_q = \emptyset;$  // initialize edges from  $q$  to spanner nodes to empty
6    $v = \emptyset;$  //  $v$  will be  $q$ 's representative
7   for  $v' \in V_S$  do // Find guards that can be connected to  $q$ 
8     if  $d(v', q) < \Delta$  and  $L(v', q) \in C_{free}$  then
9        $E_q \leftarrow E_q \cup \{L(v, q)\};$  // add local path
10      if  $v == \emptyset$  or  $d(v', q) < d(v, q)$  then
11         $v = v';$ 
12  if  $E_q == \emptyset$  then // when  $q$  not visible from existing guards
13     $\text{Add\_Guard\_Spanner}(V_S, E_S, q);$ 
14  else if  $q$  connects two spanner nodes  $v_1, v_2$  that were previously disconnected then
15     $\text{Add\_Bridge\_Spanner}(V_S, E_S, v_1, v_2, q);$ 
16  else
17     $V' = \emptyset;$  //  $V'$ : nodes with interface supported by  $q$ 
18    for all  $q' \in V_D$  so that  $L(q, q') \in E_D$  do
19      if  $\text{rep}(q') \neq \text{rep}(q)$  then  $V' \leftarrow V' \cup \text{rep}(q')$ 
20    if  $V' \neq \emptyset$  then
21      for  $v' \in V'$  do
22        if  $L(v, v') \notin E_S$  then
23           $\text{Add\_Pair\_Spanner}(V_S, E_S, v, v', q, q')$ 
24      if  $q \notin V_S$  then
25        for  $v' \in V'$  do
26          for all  $v'' \in V_S, v'' \neq v'$  so that  $L(v, v'') \in E_S, L(v', v'') \notin E_S$  do
27             $\Pi_S \leftarrow \{\pi_S(m(v, v'), m(v, v''))\};$  // midpoint paths
28          for all  $x \in V_S, x \neq v, v', v'': L(x, v), L(x, v'') \in E_S, L(x, v') \notin E_S$  do
29             $\Pi_S \leftarrow \Pi_S \cup \{\pi_S(m(v, v'), m(v, x))\};$ 
30           $\pi_S = \text{argmax}_{\pi \in \Pi_S} |\pi|;$  // max. length path
31          while  $q \notin V_S$  do
32            for all  $q''$  that support  $i(v, v'')$  do
33               $\pi_D \leftarrow \text{Shortest\_Path}(D, q, q'');$ 
34              if  $t \cdot |\pi_D| < |\pi_S|$  then
35                 $\text{Add\_Path\_Spanner}(V_S, E_S, V_D, E_D, \pi_D, v, v'');$ 
36      if  $q \notin V_S$  then  $failures++;$ 
37      else  $failures = 0;$ 
38 return  $(V_S, E_S);$ 

```

Algorithm 1 outlines the operation of SPARS. The input includes (i) the maximum number M of failures to add a node in S, (ii) t , the spanner's stretch factor, and (iii) the spanner visibility range Δ and (iv) the sample visibility range δ . The algorithm initializes the two graphs and sets the number of failed attempts to zero (line 1). While the failed attempts are fewer than M (line 2), the algorithm samples obstacle-free configurations q (line 3). Every q is added to the dense graph as a node and connected with an edge to existing samples $q' \in V_D$ so that $d(q, q') < \delta$ and $L(q, q') \in C_{free}$ (line 4). Then, SPARS identifies the connections $E_q \in C_{free}$ between the sample q and spanner nodes of distance less than Δ . Among these nodes, the representative $v = \text{rep}(q)$ is found (lines 5-11).




Fig. 3 (left) q added as a guard. (middle) q added as a bridge. (right) q, q' support $i(v, v')$, which does not have an edge crossing it. Samples q and q' are candidates for addition.

Figure 3 describes the first three cases under which SPARS adds a sample q to V_S . If none of the existing spanner nodes can be connected to q , then the sample has to be added for coverage purposes (lines 12-13). The nodes added for C-space coverage will be called “guards”. If q can connect two disconnected components, then it is also a candidate for addition (lines 14-15). The function `Add_Bridge_Spanner` first checks whether v_1 and v_2 can be connected directly with an edge. If they can, this edge is added, otherwise q is added and connected to v_1 and v_2 . If a direct edge is added, its maximum length is $2 \cdot \Delta$. Note that the previous method IRS2 employs this two ways of inserting samples into the spanner and also accepts nodes if they introduce a “useful cycle”, i.e., if they have two neighbors v_1 and v_2 so that $t \cdot (d(v_1, q) + d(q, v_2)) < \pi_S(v_1, v_2)$.

Beyond this point, the algorithm adds samples in V_S only if they support an interface. The algorithm detects if q has neighbors in D with a different representative, in which case, q lies on an interface (lines 17-19). If the representatives v and v' of q and q' do not share an edge, then q and q' are candidates for addition (lines 21-23). The objective is to guarantee that every two spanner nodes that share an interface will also share an edge in E_S . This is handled by function `Add_Pair_Spanner` (Algorithm 2), which again first attempts to connect v and v' directly. If this is possible, the edge is added and no new node or interface is introduced in S . If $L(v, v')$ is not collision-free, then `Add_Pair_Spanner` will attempt to connect v and v' through $mid = m(q, q')$. If this is not possible, then both q and q' are added to S , and connected appropriately to their neighbors. SPARS proceeds to check whether q reveals paths that are much shorter than spanner paths (lines 24-35).

Algorithm 2: `Add_Pair_Spanner` (S, D, v, v', q, q')

```

1   $mid \leftarrow m(q, q')$ ;
2  if  $L(v, v') \in \mathcal{C}_{free}$  then
3     $E_S \leftarrow E_S \cup \{L(v, v')\}$ ;
4  else if  $L(v, mid) \in \mathcal{C}_{free}$  and  $L(v', mid) \in \mathcal{C}_{free}$  then
5     $V_S \leftarrow V_S \cup \{mid\}$ ;
6     $E_S \leftarrow E_S \cup \{L(v, mid), L(v', mid)\}$ ;
7    Reduce_Interfaces( $S, D, mid$ );
8  else
9     $V_S \leftarrow V_S \cup \{q, q'\}$ ;
10    $E_S \leftarrow E_S \cup \{L(v, q), L(q, q'), L(q', v')\}$ ;
11   Reduce_Interfaces( $S, D, q$ );
12   Reduce_Interfaces( $S, D, q'$ );

```

For all interfaces $i(v, v')$ that q supports (line 25), the algorithm considers neighbors v'' of v , so that v'' is not linked to v' (line 26). Figure 4 gives an example. For such a case, the algorithm computes the spanner path $\pi_s(m(v, v'), m(v, v''))$ (line 27) and similar spanner paths between v and neighbors x of v that do not share an edge with v' (lines 28-29). Among all of them, the algorithm stores the longest path (line 30). The algorithm aims to guarantee that all these spanner paths are less than t times longer than the shortest paths between samples along $i(v, v')$ and $i(v, v'')$ (lines 31-35).




Fig. 4 q, q' support $i(v, v')$, while q'' supports $i(v, v'')$. If path $\pi_D(q, q'')$ is shorter than t times the length of spanner paths, then samples along $\pi_D(q, q'')$ are candidates for addition.

Algorithm 3: Add_Path_Spanner ($V_S, E_S, V_D, E_D, \pi, v, v''$)

```

1 prior  $\leftarrow v$ ;
2 for  $q \in \pi$  do
3    $V_S \leftarrow V_S \cup \{q\}$ ;
4    $E_S \leftarrow E_S \cup \{\mathbf{L}(\textit{prior}, q)\}$ ;
5   prior  $\leftarrow q$ ;
6  $E_S \leftarrow E_S \cup \{\mathbf{L}(\textit{prior}, v'')\}$ ;
7 for  $q \in \pi$  do
8   ReduceInterfaces( $S, D, q$ );

```

The algorithm computes the shortest path between q , which supports $i(v, v')$, and all samples that support $i(v, v'')$ and have v as their representative (lines 31-35). If the shortest path $\pi_0(q, q'')$ in D is shorter than t times the length of the longest corresponding spanner path then the algorithm considers the samples along $\pi_0(q, q'')$ for addition (lines 34-35). This is accomplished by function `Add_Path_Spanner` (Algorithm 3). The final step in `SPARS` is to make sure that if q has not been added to the spanner, that the parameter *failures* is increased (lines 36-37). An efficient implementation of `Add_Path_Spanner` would first try to directly connect v with v'' , and it would also try to smooth the final path added to S .

Algorithm 4: Reduce_Interfaces (S, D, v)

```

1 for  $v' \in V_S$  s.t.  $d(v, v') < 2 \cdot \Delta$  do
2   if  $L(v, v') \notin E_S$  then
3     for  $(q, q') \in V_D$  do
4       if  $d(q, m(v, v')) < \delta$  and  $d(q', m(v, v')) < \delta$  then
5         if  $rep(q) \neq rep(q')$  then
6           if  $L(v, rep(q')) \in C_{free}$  then
7              $E_S \leftarrow E_S \cup \{L(v, rep(q'))\};$ 

```

The method `Reduce_Interfaces` is called by all functions that add nodes to V_S . It aims to directly connect vertices which share an interface but not an edge using the local planner. Doing so prevents the algorithm from needing to add samples q and q' that support this interface.

4 Analysis

The existence of a finite set of nodes that achieve even coverage depends on properties of \mathcal{C}_{free} . Such a finite set may not always exist. This work assumes that the \mathcal{C} -space is expansive, similar to related literature [5, 6].

Definition 10 (Lookout). The lookout of a set Q is the subset of configurations q in Q , such that the measure of what is visible from q outside of Q is greater than some fraction of the measure of the part of the free space outside of Q . More formally:

$$\text{LOOKOUT}_\beta(Q) = \{q \in Q \mid \mu(V(Q) \setminus Q) \geq \beta \mu(\mathcal{C}'_{free} \setminus Q)\}$$
 where \mathcal{C}'_{free} is a \mathcal{C}_{free} connected component and $V(Q)$ is the visibility of set Q .

Definition 11 (ϵ, α, β -expansiveness). A space is expansive if it satisfies:

- $\forall q \in \mathcal{C}'_{free}$, where $\mu(V(q)) \geq \epsilon \mu(\mathcal{C}_{free})$
- $\forall Q \subseteq \mathcal{C}'_{free}$, where $\mu(\text{LOOKOUT}_\beta(Q)) \geq \alpha \mu(Q)$

To provide coverage, it has to be that the resulting spanner allows every collision-free configuration to be connected with a spanner node through a collision-free path.

Theorem 1 (Coverage). *Upon termination of SPARS and with probability $1 - \frac{1}{M}$, the following is true, $\forall q \in \mathcal{C}_{free} : \exists v \in V_S$ so that $L(q, v) \in \mathcal{C}_{free}$.*

The detailed argument to support the above statement can be found in the work that presented the Visibility PRM [23] as SPARS is following a similar approach regarding \mathcal{C} -space coverage and the stopping criterion. It also relates to the property of ϵ -goodness that expansive spaces provide [6]. At a high level, each new guard inserted to S increases the coverage of \mathcal{C}_{free} and the probability of generating configurations in non-covered regions decreases over iterations. The algorithm is then guaranteed to terminate for any finite input value M . When it stops, a probabilistic estimation of the percentage of free space not covered by spanner nodes is $\frac{1}{M}$, given uniform sampling. This means that future attempts to connect configurations to spanner nodes will succeed with probability $(1 - \frac{1}{M})$. Consequently, as M goes to infinity, the resulting graph covers the entire space. This is a conservative estimate, since the algorithm assumes an artificial visibility range limit Δ for spanner nodes. The work on Visibility PRMs has shown that even for relatively complicated problems in $SE(3)$, a relatively small number of guards is needed to probabilistically cover the space. Connectivity properties of the resulting sparse roadmap spanner can be argued in a similar manner.

Theorem 2 (Connectivity). *Upon termination of SPARS and with probability 1 as M goes to infinity, for every pair of nodes $v, v' \in V_S$ that are connected with a collision-free path in \mathcal{C}_{free} , there is a path $\pi_S(v, v')$ that connects them on S .*

The above property arises from lines 14-15 of SPARS. The algorithm adds an edge or a node every time it detects there is a way to connect two disconnected components. The probability of sampling a configuration that will connect such disconnected components of the graph depends on the environment. Narrow passages will make this connection more challenging. The probability goes to 1, however, as the value of M goes to infinity, since every configuration will be eventually sampled. The combination of the last two theorems provides probabilistic completeness, at least when the algorithm samples configurations q in a uniform way.

Regarding path quality, using δ -PRM, the approach is able to construct a dense graph that provide asymptotic optimality [9]. Thus, the following lemma holds.

Lemma 1 (Asymp. Optimality of D). *Upon termination of SPARS and with probability 1 as $M \rightarrow \infty$, for every cl -robust optimal path $\pi_C(q, q')$, there is a collision-free path on the dense graph $\pi_D(q, q')$, so that: $|\pi_D(q, q')| \rightarrow |\pi_C(q, q')|$.*

Thus, if the length of spanner paths is bounded relative to dense paths, they will be asymptotically bounded relative to optimal C-space paths.

Lemma 2 (Coverage of Optimal Paths by S). *Consider the shortest dense path $\pi_D(q_0, q_m)$. As M goes to infinity, the probability of having a sequence of guard nodes $U = (v_1, \dots, v_n)$ with the following properties goes to 1:*

- Each q along $\pi_D(q_0, q_m)$ belongs to the visibility region of a spanner node.
- v_1 is the representative of q_0 , $v_1 = rep(q_0)$ and similarly $v_n = rep(q_m)$.
- All the paths $L(v_i, v_{i+1})$ belong to the set E_S of the sparse roadmap spanner.

Proof. Fig. 5 (left) illustrates the first property, which is a direct outcome of the coverage theorem. The second point holds without loss of generality. The following discussion focuses on the validity of the last point and relates to Fig. 5 (middle). Consider q_j along $\pi_D(q_0, q_m)$ that supports $i(v_i, v_{i+1})$. With probability 1 as $M \rightarrow \infty$, samples have been generated in the $\frac{cl}{2}$ neighborhood of q_j , which is guaranteed to be collision-free. Consider two samples in the $\frac{cl}{2}$ -ball centered at q_j : (a) q_j^i with spanner node v_i as its representative and (b) q_j^{i+1} with v_{i+1} as its representative.

To guarantee the algorithm will produce such samples, it has to be that the set of such configurations has non-zero measure. The region from which the sample q_j^i must be generated from is the intersection of (a) the volume of visibility paths from v_i to the optimum path π_D and (b) the $\frac{cl}{2}$ -ball centered at q_j , which has positive measure given expansiveness [6]. An expansive space can be partitioned so that the $\frac{cl}{2}$ -ball is one partition of the space. Expansiveness states that every subset of both partitions must have visibility of the other with positive measure. It is already true that the local path $L(v_i, q_j)$ is collision-free and contributes towards the lookout of the partition containing v_i . From the above observations, a sample q_j^i will be generated and similarly for q_j^{i+1} .

The local path $L(q_j^i, q_j^{i+1})$ must be collision-free. Furthermore, for $\delta > 0$ there are q_j^i and q_j^{i+1} for which $d(q_j^i, q_j^{i+1}) < \delta$. Then, the algorithm would have added the edge $L(q_j^i, q_j^{i+1})$ in E_D . There are two cases at this point. Either v_i and v_{i+1} are connected as desired, or the representatives are not connected. Assume v_i and v_{i+1} are not connected. Then all the requirements of the lines 17-23 of the algorithm are satisfied: (i) $q_j^i, q_j^{i+1} \in V_D$ and $L(q_j^i, q_j^{i+1}) \in E_D$, (ii) $rep(q_j^i) = v_i \neq v_{i+1} = rep(q_j^{i+1})$ (iii) $L(v_i, v_{i+1}) \notin E_S$. Thus, `AddPair.Spanner` would have been called. The function would either introduce the edge $L(v_i, v_{i+1})$ or it would add the samples q_j^i, q_j^{i+1} to S together with the edges $L(v_i, q_j^i)$, $L(q_j^i, q_j^{i+1})$, $L(q_j^{i+1}, v_{i+1})$. This means that $(v_1, \dots, v_i, v_{i+1}, \dots, v_n)$ can be replaced by $(v_1, \dots, v_i, q_j^i, q_j^{i+1}, v_{i+1}, \dots, v_n)$, which still covers the dense path $\pi_D(q_0, q_m)$. In this case, all the nodes between v_i and v_{i+1} are pairwise connected. Thus, if an edge doesn't exist, a new sequence can always be created where all the guards are pairwise connected. \square




Fig. 5 (left) For every $\pi_D(q_0, q_m)$, there is a sequence of V_S nodes $\{v_1, \dots, v_n\}$ so that every q along the path is “visible” by at least one v_i . For each (v_i, v_{i+1}) there is an edge $L(v_i, v_{i+1})$ in E_S . (middle) If there is no edge $L(v_i, v_{i+1})$, the samples q_i^i and q_{i+1}^{i+1} would have been added to S . (right) The spanner path $\pi_S(q_0, q_5)$ can be split into: $\{M_{q_0}, M_1, M_2, M_3, M_{q_5}\}$. Each M_i is a path of the form $\pi_S(m(v_i, v_{i+1}), m(v_{i+1}, v_{i+2}))$. The first and last segments connect q_0 and q_5 to the spanner path.

A solution path $\pi_S(q_0, q_m)$ computed through the spanner will correspond to the concatenation of $L(q_0, v_1)$ and $L(v_n, q_m)$ with local paths $L(v_i, v_j)$ along the sequence $U = (v_1, \dots, v_n)$. Note that not all consecutive guard edges of the type $L(v_i, v_{i+1})$ need to be part of the shortest spanner path, as there can be shortcuts along U . Then, the path on the spanner can be decomposed into segments $M_i = \pi_S(m(v_{i-1}, v_i), m(v_i, v_{i+1}))$, which start at a midpoint between two spanner nodes and stop at another midpoint. There are also the following special segments: $M_{q_0} = (L(q_0, v_1), L(v_1, m(v_1, v_2)))$ and $M_{q_m} = (L(m(v_{n-1}, v_n), v_n), L(v_n, q_m))$. These segments connect the start or the final configuration with spanner midpoints. Figure 5 (right) provides an illustration. Overall: $\pi_S(q_0, q_m) \equiv (M_{q_0}, M_1, \dots, M_k, M_{q_m})$.

Lemma 3 (Connection Cost). *The length of M_{q_0} and M_{q_m} of a path between configurations q_0, q_m computed through the spanner is bounded by: 4Δ .*

Proof. Consider the segment M_{q_0} . The length of $L(q_0, v_1)$ cannot be longer than Δ because v_1 is the representative of q_0 . Similarly, the distance between v_1 and $m(v_1, v_2)$ cannot be more than Δ , since $m(v_1, v_2)$ is the midpoint between two spanner nodes that share an edge. Overall $|M_{q_0}| < \frac{4\Delta}{2}$. The same bound can be shown for the segment M_{q_m} . Consequently: $|M_{q_0}| + |M_{q_m}| < 4 \cdot \Delta$. \square

Lemma 4 (Spanner Property). *Segments M_i have length bounded by $t \cdot |\pi_D(q_{i-1}, q_i)|$, where q_i lies at the intersection of $\pi_D(q_0, q_m)$ with the interface $i(v_i, v_{i+1})$.*

Proof. The arguments presented here correspond to Fig. 6. Given lemma 2, the edges $L(v_{i-1}, v_i)$ and $L(v_i, v_{i+1})$ must belong to the set E_S , as the algorithm ensures that all interfaces will eventually have edges connecting the vertices inducing them. Assume the true optimal path goes through the visibility region of vertex v_i on sequence U . There are two cases to be considered: (a) there is no edge $L(v_{i-1}, v_{i+1})$ in E_S (Fig. 6 (left)) or (b) there is such an edge (Fig. 6 (right)).

To ensure that the spanner property is upheld in the first case, it is sufficient that the path on S from $m(v_{i-1}, v_i)$ to $m(v_i, v_{i+1})$ is shorter than t times the length of the shortest dense paths between interfaces $i(v_{i-1}, v_i)$ and $i(v_i, v_{i+1})$. Lines 24-35 of SPARS, however, ensure that this property is true. When the newly sampled configuration q , is on an interface $i(v, v')$, the algorithm will search for paths from v' to all v'' where $L(v', v'') \notin E_S$. The equivalent of v in Fig. 6 (left) is v_i , v' would be v_{i+1} and v'' would be v_{i-1} or vice versa.




Fig. 6 (left) The optimal path between q_{i-1} and q_i is represented by the segment M_i on the spanner. The algorithm checks whether the length of M_i is less than $t \cdot |\pi_D(q_{i-1}, q_i)|$. (right) Nodes u_{i-1} and u_{i+1} are connected with an edge in this case. Thus, the spanner path will not contain the node u_i . Nevertheless, the algorithm checks whether the spanner path $[m(u_{i-2}, u_{i-1}), u_{i-1}], [u_{i-1}, m(u_{i-1}, u_{i+1})]$ is shorter than t times the shortest dense path between the interfaces $i(u_{i-2}, u_{i-1})$ and $i(u_{i-1}, u_{i+1})$. The same is true for $i(u_{i-1}, u_{i+1})$ and $i(u_{i+1}, u_{i+2})$.

In the second case, the algorithm will not check the interface $i(v_{i-1}, v_i)$ and $i(v_i, v_{i+1})$, because the nodes v_{i-1} and v_{i+1} share an edge. Nevertheless, the algorithm considers additional spanner paths when it compares against the length of dense paths. In particular, the algorithm will check whether the spanner path from $m(v_{i-2}, v_{i-1})$ to $m(v_{i-1}, v_{i+1})$ is shorter than t times the length of the shortest dense paths between the interfaces $i(v_{i-2}, v_{i-1})$ and $i(v_{i-1}, v_i)$. Notice the difference in the second vertex of the second midpoint and the second interface. In this case, node v_{i-1} in Fig. 6 (right) is equivalent to node v in lines 24-35 of SPARS, v_{i-2} is equivalent to node v' , v_i is equivalent to v'' and v_{i+1} is equivalent to node x . The algorithm will also provide the same guarantee for the spanner path $m(v_{i-1}, v_{i+1})$ to $m(v_{i+1}, v_{i+2})$ and overall: $|\pi_S(m(v_{i-2}, v_{i-1}), m(v_{i+1}, v_{i+2}))| < t \cdot (|\pi_D(q_{i-2}, q_{i-1})| + |\pi_D(q_i, q_{i+1})|) < t \cdot (|\pi_D(q_{i-2}, q_{i-1})| + |\pi_D(q_{i-1}, q_i)| + |\pi_D(q_i, q_{i+1})|)$ and the spanner property is still satisfied. If the spanner property is violated, SPARS will call function `Add_Path_Spanner`, which will either directly connect the corresponding vertices or add the dense path into the spanner graph. \square

Theorem 3 (Asympt. Near-Optimality with Additive Cost). *Upon termination of SPARS, as $M \rightarrow \infty$ and $\forall q_0, q_m \in \mathcal{C}_{free}$: $|\pi_S(q_0, q_m)| < t \cdot |\pi_D(q_0, q_m)| + 4 \cdot \Delta$.*

Proof. Lemma 2 provides the existence of a sequence U of spanner nodes that can be used to solve path planning queries for any q_0, q_m (i.e., it is possible to connect q_0 and q_m to the sequence U , and all the nodes in the sequence are pairwise connected). The combination of lemmas 3 and 4 indicates that the solution path $\pi_S(q_0, q_m)$ computed through this sequence has length that is bounded relative to the length of the dense path $|\pi_S(q_0, q_m)| < t \cdot |\pi_D(q_0, q_m)| + 4 \cdot \Delta$, where the additive cost arises from the cost of connecting the query points to the spanner nodes. Lemma 1 indicates that the paths computed through the dense graph converge asymptotically to the optimal ones $\pi_S(q_0, q_n)$. The combination of all these lemmas proves this theorem. \square

An important issue is whether SPARS adds an infinite number of nodes to S in order to provide the above path quality properties.

Theorem 4 (Finiteness). *As the number of iterations goes to infinity, the probability of SPARS adding nodes to S goes to zero.*

Proof. There are four ways with which SPARS may add nodes in the spanner S and is necessary to show that the probability of adding each type of node goes to zero.

Nodes for coverage are added when a sample q lies outside the visibility range of all existing nodes in V_S . Theorem 1 already argues that the probability of adding guards diminishes to zero as the number of iterations increases.

Nodes for connectivity are added when a sample q connects two disconnected components. Eventually, enough nodes for ensuring coverage will be added in S . Given the number of nodes n at that point, the number of samples that can connect disconnected components is finite.

Nodes for ensuring interfaces have edges are added when a sample q has a neighbor q' in D and their representatives v and v' do not share an edge. Assuming coverage and connectivity has been ensured, it is possible there are interfaces without an edge. In the general case, when a node is added to S , multiple new interfaces are defined. But the algorithm employs the function `Reduce_Interfaces`, which attempts to connect a new node to all nodes with which it shares an interface. There are three cases for two neighboring nodes: (i) The interface lies outside C_{free} and there is no need to add an edge. (ii) $L(v, v')$ lies within C_{free} and the edge can be immediately added. (iii) $L(v, v')$ is not within C_{free} , but there exists a portion of the interface in C_{free} . In this case, new nodes need to be added, which, however, will also add new interfaces, which need to be connected by an edge. Nevertheless, the newly generated interfaces will eventually have to fall into case (i) or case (ii) if the C -space is nicely behaved. Otherwise, it has to be that every time a node is added, there will exist an obstacle which partially covers the local path to neighbors with which the node shares an interface. Overall, all interfaces will be eventually connected with an edge, so the probability of adding nodes for this reason goes to 0.

Nodes for ensuring path quality are added when a sample q supports an interface and reveals a path in D that is significantly shorter than the corresponding path in S . SPARS always adds the minimal amount of nodes possible while ensuring that the path quality properties of the algorithm hold. When adding a dense path to S , SPARS will first try to smooth this path. In some cases, the path cannot be smoothed, which will result in the addition of nodes to S . The addition of new nodes could potentially always create new visibility regions where a dense path could be found that is sufficiently shorter than the spanner paths. Eventually, however, the spanner nodes which the algorithm attempts to connect will be within cl of each other and it would be guaranteed that $L(v, v')$ is collision-free and the nodes can be connected directly. Therefore, SPARS will eventually stop adding nodes in this fashion.

5 Evaluation

5.1 Implementation Details and Setup

The algorithm was implemented in a simulation software written in C++. To improve space efficiency, information was kept to a minimum on the dense graph's nodes and edges. Dense nodes take up approximately 40 bytes of memory in $SE(2)$ examples and 72 bytes in $SE(3)$ examples. Sparse nodes use a larger amount of memory, as they retain a lists of all the nodes in the dense graph which they repre-




Fig. 7 Environments for testing (left to right): $SE(2)$: 2D Maze, Teeth, $SE(3)$: 3D Hole and Many Holes. The 2D Maze tests performance for constrained spaces, while the Teeth environment tests performance in large, open spaces. The output roadmap is visualized for the $SE(2)$ case.

sent, introducing a cost of 4 bytes per dense node. All edges use approximately 12 bytes of memory. The implementation performed several optimizations compared to the algorithm described. In order to speed up convergence, connections are attempted to neighbors of newly added sparse nodes in the spanner. The neighbors considered are all within $2 * \Delta$ distance. It often happens that the algorithm will not be able to quickly identify that two spanner nodes share an interface, as not enough samples will be present in the dense graph D to correctly identify the interfaces. Furthermore, whenever paths are added to the spanner, the algorithm always attempts to directly connect points, and if it cannot, it will try smoothing the paths. For nearest neighbors queries, kd-trees were used. Nevertheless, it is often much more practical to consider the neighbors of nodes already in the spanner or the graph and then prune them according to their distance to the selected node. This technique was used heavily, especially for querying the dense graph.

Experiments were run in four environments in $SE(2)$ and $SE(3)$, as shown in Figure 7 for kinematic rigid bodies. The 2D Maze used a disk system, while the Teeth environment used a ‘Z’ shaped body. For both cases in $SE(3)$, a variation of the ‘Z’ robot was used which has one of its legs coming off in an orthogonal direction to the other leg. Runs in $SE(2)$ were tested with $M = 4000$, $t = 3$, and $\Delta = 20$, while runs in $SE(3)$ were tested with $M = 1400$, $t = 3$, and $\Delta = 25$. All runs were terminated if the stopping criterion was not satisfied after 60 minutes.

5.2 Results and Comparison

SPARS was compared against asymptotically optimal δ -PRM and near-optimal planners IRS2. The algorithms were compared in terms of path quality, query resolution time, and number of nodes and edges in the final roadmaps. Results shown for δ -PRM were extracted from the dense graph D constructed by the SPARS algorithm. In this way, the stopping criterion for δ -PRM relates to the stopping condition of SPARS. Figure 8 shows an overview of the algorithms’ performance.

SPARS is the best alternative in terms of aggressively pruning the number of nodes in the final roadmap and is able to retain many-orders of magnitude smaller roadmaps than δ -PRM. The related IRS2 method retains slightly more nodes than SPARS but actually returns fewer edges. Nevertheless, given that nodes have higher space requirements relative to edges, the overall space requirements of SPARS are more limited than IRS2. This had a significant effect on query resolution times which were significantly shorter for SPARS relative to δ -PRM. In terms of path length, the asymptotically near-optimal algorithms were able to return solutions




Fig. 8 A comparison of various aspects of the algorithms. (upper left) The number of nodes in the final roadmaps. (upper right) The number of edges in each roadmap. (lower left) Path length for SPARS and IRS2 as a percentage of paths from δ -PRM. (lower right) Average query resolution time for 100 queries.

which were very close to those computed by δ -PRM. Note that the part of δ -PRM suffers from the side-effect of having multiple nodes along the solution path. Smoothing the final paths results in comparable path lengths between the various methods.

6 Discussion

This paper describes the first approach, to the best of the author’s knowledge, that provides a finite size data structure, which can be used to answer path queries in continuous spaces with near-optimality guarantees. Simulations indicate that the resulting graph is many orders sparser (i.e., it has fewer nodes and edges) than graphs with asymptotic optimality guarantees. This results in significantly shorter query resolution times. The quality of paths computed on the sparse roadmap spanner is shown to be in practise significantly better than the theoretical guarantees.

Future research will address the dependency of the current approach to the input parameters, such as the visibility ranges Δ , δ and the maximum allowed number of failures M . An interesting question is whether the proposed method can discover the important homotopic classes in a C-space and how does it compare against methods that aim to identify homotopic classes [7, 22]. It is also important to consider alternative methods that do not need to store during the construction of the spanner the entire graph returned by the asymptotically optimal planner. Such a development could reduce the space and time requirements of the spanner’s construction. In a similar direction, it will be helpful to introduce steps that improve computational efficiency, such as tools for reducing the cost of A* searches on the dense graph or the cost of nearest-neighbor queries.

Acknowledgements Work by the authors has been supported by NSF CNS 0932423. Any conclusions expressed here are of the authors and do not reflect the views of the sponsors.

References

- [1] Agarwal, P.: Compact Representations for Shortest-Path Queries (2011). Appeared at the IROS 2012 Workshop on Progress and Open Problems in Motion Planning
- [2] Amato, N.M., Bayazit, O.B., Dale, L.K., Jones, C., Vallejo, D.: OBPRM: An Obstacle-based PRM for 3D Workspaces. In: WAFR, pp. 155–168 (1998)
- [3] Baswana, S., Sen, S.: A Simple and Linear Time Randomized Algorithm for Computing Spanners in Weighted Graphs. *Random Structures and Algorithms* **30**(4), 532–563 (2007)
- [4] Geraerts, R., Overmars, M.H.: Creating High-Quality Roadmaps for Motion Planning in Virtual Environments. In: IROS, pp. 4355–4361. Beijing, China (2006)
- [5] Hsu, D., Kavraki, L., Latombe, J.C., Motwani, R., Sorkin, S.: On Finding Narrow Passages with Probabilistic Roadmap Planners. In: WAFR. Houston, TX (1998)
- [6] Hsu, D., Kindel, R., Latombe, J.C., Rock, S.: Randomized Kinodynamic Motion Planning with Moving Obstacles. *IJRR* **21**(3), 233–255 (2002)
- [7] Jaillet, L., Simeon, T.: Path Deformation Roadmaps. In: WAFR. New York City, NY (2006)
- [8] Kallman, M., Mataric, M.: Motion Planning Using Dynamic Roadmaps. In: ICRA, vol. 5, pp. 4399–4404. New Orleans, LA (2004)
- [9] Karaman, S., Frazzoli, E.: Sampling-based Algorithms for Optimal Motion Planning. *IJRR* **30**(7), 846–894 (2011)
- [10] Kavraki, L.E., Kolountzakis, M.N., Latombe, J.C.: Analysis of Probabilistic Roadmaps for Path Planning. *IEEE TRA* **14**(1), 166–171 (1998)
- [11] Kavraki, L.E., Svestka, P., Latombe, J.C., Overmars, M.: Probabilistic Roadmaps for Path Planning in High-Dimensional Configuration Spaces. *IEEE TRA* **12**(4), 566–580 (1996)
- [12] Ladd, A.M., Kavraki, L.E.: Measure Theoretic Analysis of Probabilistic Path Planning. *IEEE TRA* **20**(2), 229–242 (2004)
- [13] LaValle, S.M., Kuffner, J.J.: Randomized Kinodynamic Planning. *IJRR* **20**, 378–400 (2001)
- [14] Li, Y., Bekris, K.E.: Learning Approximate Cost-to-Go Metrics To Improve Sampling-based Motion Planning. In: IEEE ICRA. Shanghai, China (2011)
- [15] Marble, J.D., Bekris, K.E.: Asymptotically Near-Optimal is Good Enough for Motion Planning. In: ISRR. Flagstaff, AZ (2011)
- [16] Marble, J.D., Bekris, K.E.: Towards Small Asymptotically Near-Optimal Roadmaps. In: IEEE ICRA. Minnesota, MN (2012)
- [17] Nechushtan, O., Raveh, B., Halperin, D.: Sampling-Diagrams Automata: a Tool for Analyzing Path Quality in Tree Planners. In: WAFR. Singapore (2010)
- [18] Nieuwenhuisen, D., Overmars, M.H.: Using Cycles in Probabilistic Roadmap Graphs. In: IEEE ICRA, pp. 446–452 (2004)
- [19] Peleg, D., Schäffer, A.: Graph Spanners. *Journal of Graph Theory* **13**(1), 99–116 (1989)
- [20] Raveh, B., Enosh, A., Halperin, D.: A Little More, a Lot Better: Improving Path Quality by a Path-Merging Algorithm. *IEEE TRO* **27**(2), 365–370 (2011)
- [21] Sanchez, G., Latombe, J.C.: A Single-Query, Bi-Directional Probabilistic Roadmap Planner with Lazy Collision Checking. In: ISRR, pp. 403–418 (2001)
- [22] Schmitzberger, E., Bouchet, J.L., Dufaut, M., Wolf, D., Husson, R.: Capture of Homotopy Classes with Probabilistic Roadmap. In: IEEE/RSJ IROS, pp. 2317–2322 (2002)
- [23] Simeon, T., Laumond, J.P., Nissoux, C.: Visibility-based Probabilistic Roadmaps for Motion Planning. *Advanced Robotics Journal* **41**(6), 477–494 (2000)
- [24] Varadhan, G., Manocha, D.: Star-shaped Roadmaps: A Deterministic Sampling Approach for Complete Motion Planning. *IJRR* (2007)
- [25] Xie, D., Morales, M., Pearce, R., Thomas, S., Lien, J.L., Amato, N.M.: Incremental Map Generation (IMG). In: WAFR. New York City, NY (2006)