THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

HDRM: A Resolution Complete Dynamic Roadmap for Real-Time Motion Planning in Complex Environments

Citation for published version:

Yang, Y, Merkt, W, Ivan, V, Li, Z & Vijayakumar, S 2018, 'HDRM: A Resolution Complete Dynamic Roadmap for Real-Time Motion Planning in Complex Environments', *IEEE Robotics and Automation Letters*, vol. 3, no. 1, pp. 551-558. <https://doi.org/10.1109/LRA.2017.2773669>

Digital Object Identifier (DOI):

[10.1109/LRA.2017.2773669](https://doi.org/10.1109/LRA.2017.2773669)

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Peer reviewed version

Published In:

IEEE Robotics and Automation Letters

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

HDRM: A Resolution Complete Dynamic Roadmap for Real-Time Motion Planning in Complex Environments

Yiming Yang, Wolfgang Merkt, Vladimir Ivan, Zhibin Li, and Sethu Vijayakumar
Institute of Perception, Action and Behaviour, School of Informatics, University of Edinburgh
Email: {yiming.yang, wolfgang.merkt, v.ivan, zhibin.li, sethu.vijayakumar}@ed.ac.uk

Abstract—We present the Hierarchical Dynamic Roadmap (HDRM), a novel resolution complete motion planning algorithm for solving complex planning problems. A unique hierarchical structure is proposed for efficiently encoding the configuration-to-workspace occupation information that allows the robot to check the collision state of tens of millions of samples on-the-fly—the number of which was previously strictly limited by available memory. The hierarchical structure also significantly reduces the time for path searching, hence the robot is able to find feasible motion plans in real-time in extremely constrained environments. The HDRM is theoretically proven to be resolution complete, with a rigorous benchmarking showing that HDRM is robust and computationally fast, compared to classical dynamic roadmap methods and other state-of-the-art planning algorithms. Experiments on the 7 degree-of-freedom KUKA LWR robotic arm integrated with real-time perception of the environment further validate the effectiveness of HDRM in complex environments.

I. INTRODUCTION

Motion planning is one of the fundamental problems in robotics, involving automatically finding a sequence of configurations that take the robot from a start to a goal pose. Generally, motion planning can be categorized into optimization-based and sampling-based algorithms. Optimization-based methods [1], [2] generate optimal trajectories with respect to cost functions, but may get stuck in local minima and fail to produce a valid solution when the problem is non-convex or ill-defined. On the other hand, sampling-based approaches [3], [4], [5], [6] promise to solve complex problems by sampling globally in the configuration space.

In sampling-based algorithms, collision checking is usually the most expensive operation and reportedly consumes up to 90-95% of the planning time [7]. Lazy collision checking is used to delay the collision checking until it is needed [8]. One can also define possible collision regions and limit collision checking to these regions [9]. However, these techniques only reduce the collision checking time indirectly by reducing the number of calls, but not the actual computation time of the collision checking function. Parallel implementations have been proposed to speed up collision checking and motion planning [10], but these approaches focus on parallelization and system implementation based on existing algorithms.

In contrast, the Dynamic Roadmap (DRM) algorithm [11], an extension to the probabilistic roadmap (PRM) [4], algorithmically reduces the collision checking time by encoding configuration-to-workspace occupation information. Given different environments, the DRM can efficiently remove invalid edges and form a valid subset of the full roadmap. Subsequently, search algorithms can proceed without considering collision checking since the remaining vertices and edges are

Fig. 1: The 7-DoF KUKA LWR robot with SCHUNK Dexterous Hand operating inside a cage. Left: grasping the target from upright posture; right: dropping the object to the side.

all collision-free. However, encoding the occupation information requires to store a significant amount of data, which needs to be loaded into memory during run-time. In the early work [11], [12], the low amount of available memory allowed storing only small roadmaps with limited number of vertices and edges. Without enough vertices and edges to densely cover the configuration space, the DRM achieves a very low planning success rate [12], [13]. The success rate is closely tied with the term *Completeness*. An algorithm is considered complete if for any input it correctly reports whether there is a solution or not. If a solution exists, it must return one in finite time [14]. Optimization-based methods are incomplete due to local minima. Unfortunately, sampling-based algorithms, such as PRM and Rapidly-exploring Random Tree (RRT), are also incomplete. A weaker notion of completeness called *probabilistic completeness* is used to describe random sampling-based algorithms. This means that with enough samples (possibly infinite number of samples), the probability of finding a solution asymptotically converges to one. Another term, *resolution completeness*, is used if an algorithm guarantees to find a solution in finite time; however, if a solution does not exist, the algorithm may run forever by incrementally increasing the sampling resolution; or, the algorithm may terminate in finite time by reporting no solution at a certain resolution, although a solution may exist at a finer resolution. In [15], a DRM consisting of up to a million vertices can be stored and updated efficiently for achieving good performance when more memory and powerful GPUs are available. However, even one million vertices are not enough

Fig. 2: Preliminaries of DRM. Left: probabilistic roadmap in configuration space; right: workspace swept volume of an edge.

for dense coverage of 6-7 dimensional configuration space. Instead of claiming completeness, a small size DRM was built on customized hardware for solving very specific tasks [16]. Considering only vertices but not edges makes the DRM a Dynamic Reachability Map that finds valid reaching poses for high-dimensional floating-base robots [17], [18]. No matter whether we store a small roadmap with limited vertices and edges, or only the vertices, these methods eventually need to sacrifice completeness for a manageable storage size.

In order for robust motion planning in complex environments to be practical, a tremendous number of vertices and edges are required; yet storing such a roadmap is infeasible on commodity computers with current technology. However, we observe that the memory consumption for storing a DRM can be greatly reduced by exploiting the topology of robot arms. On this basis, a new resolution complete planning algorithm, the *Hierarchical Dynamic Roadmap* (HDRM), is proposed with the following contributions:

- Theoretical proof of resolution completeness of HDRM as a deterministic roadmap with a discretized workspace;
- A novel formulation for encoding the occupancy information of roadmap vertices and eliminating the necessity of computing/storing edges, which enables efficient storing of roadmaps with tens of millions of vertices;
- A novel technique that efficiently identifies/removes colliding samples, which in turn enables real-time planning for single-chain robot manipulators in complex environments by exploiting the proposed hierarchical roadmap.

Extensive benchmarking shows that the HDRM is able to find valid solutions in extremely constrained conditions within a few milliseconds or less, which could not be achieved previously by classical DRM and other state-of-the-art algorithms. Experiments on a KUKA LWR arm further demonstrate HDRM's capability to solve real-world problems.

II. DYNAMIC ROADMAP

A. Preliminaries

Let $\mathcal{C} \in \mathbb{R}^N$ be the configuration space of a N -DoF robot and $\mathbf{q} \in \mathcal{C}$ be a state in configuration space. Let $\mathcal{C}_{obs}$ represent the obstacles and $\mathcal{C}_{free} = \mathcal{C} \setminus \mathcal{C}_{obs}$ the collision-free region. A classical PRM contains a connected graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, where $\mathcal{V} \in \mathcal{C}_{free}$ are the vertices and $\mathcal{E}$ are the edges that connect two neighboring vertices, as highlighted in Fig. 2 (left). These vertices and edges are generated during off-line pre-processing. During the on-line planning phase, given start and goal states $\mathbf{q}_{start}, \mathbf{q}_{goal}$, we first find these two vertices

$\mathcal{V}_{start}$ and $\mathcal{V}_{goal}$ which are closest to the start and goal states respectively. Then a graph search algorithm, such as A^* [19], is deployed to find a path in the roadmap connecting $\mathcal{V}_{start}$ and $\mathcal{V}_{goal}$. However, the pre-generated vertices and edges may not be valid in unknown and non-static environments. The validity of the stored vertices and edges must be checked, and in many cases we need to sample new collision-free configurations during the on-line phase which is very time consuming.

The dynamic roadmap (DRM) is a variation of the PRM proposed by Leven and Hutchinson [11]. The DRM is *dynamic* in the sense that the graph $\mathcal{G}$ can be dynamically updated in different environments. The invalid vertices and edges can be efficiently identified and removed, with the remaining ones forming a new graph of only valid vertices and edges. This reduced graph is ready for path searching algorithms without considering collision checking. The key feature of DRM is a configuration-to-workspace mapping, as highlighted in Fig. 2 (right). One can find the list of discretized workspace voxels which an edge occupies, referred to as the swept volume. If one or more of the voxels in the swept volume are in collision, then the corresponding edge becomes invalid. In practice, it is inefficient to check the swept volume of all edges exhaustively. Instead, the occupation information is stored per each workspace voxel, i.e. each voxel stores a list of edges that sweep through this voxel. In a new environment, we first find all the voxels that are occupied by the obstacles in the environment. Then, by iterating through the occupation lists of these invalid voxels, all the invalid edges can be found and then removed accordingly.

The main observed limitation of the existing DRM method is its low success rate [12], [13]. Although the success rate can be improved for particular tasks by carefully selecting the samples [16], those approaches can not solve generic problems, therefore they are not (resolution) complete.

B. Resolution Completeness of a Deterministic Roadmap with a Discretized Workspace

In this section, we provide theoretical proof of the conditions and boundaries of resolution completeness for deterministic DRM methods with a discretized workspace.

The work in [20] has proven that a deterministic roadmap is resolution complete. Here, *deterministic* refers to the property of the sampling distribution. The authors proved that uniform sampling (e.g. a Sukharev grid) results in a resolution complete planning algorithm. To show this, let Ψ be the subset of the power set of $\mathcal{C}$ corresponding to all open subsets that can be constructed with algebraic constraints as defined in [21], and $\Psi(x)$ for $x \in (0, \infty)$ be the set of all $\mathcal{C}_{free}$ with the *width* of $\mathcal{C}_{free}$, $w(\mathcal{C}_{free}) \geq x$ (see [20]). The width x can be viewed as the minimum width of a passable corridor in the collision free portion of the configuration space. Fig. 3a illustrates the corridor (with solid color background). All queries lie within this corridor, therefore, if all the queries of the deterministic roadmap also lie within the same corridor, the roadmap is hence resolution complete. The minimal width of the corridor x required for completeness is defined in [20]. We extend their proof to roadmaps with discretized work space.

Fig. 3: Illustration of additional volume an obstacle in a discretized workspace occupies in the configuration space. An algorithm is resolution complete if it accounts for the additional increase corridor width $f(s)$ due to discretization.

Lemma 1: After M iterations, a deterministic DRM is resolution complete for all $C_{free} \in \Psi(4b(N)M^{-\frac{1}{N}} + f(s))$, where M is the number of samples, N is the dimension of the configuration space, s is the resolution of the workspace, $b(N)$ is a factor that depends on the sampling method ($b(N) = 1$ for HDRM) and $f(s)$ is a robot-dependent function.

Proof: It has been proven in [20] that, after M iterations, a deterministic roadmap planner is resolution complete for all $C_{free} \in \Psi(4b(N)M^{-\frac{1}{N}})$, without workspace discretization. However, as shown in Fig. 3, with a discretized workspace with voxel size $s > 0$, the corresponding $C_{corridor}$ and C_{obs} are both inflated due to the workspace discretization, where $C_{corridor}$ is the narrowest corridor in the configuration space. Let $C'_{corridor}$ and C'_{obs} denote the inflated $C_{corridor}$ and C_{obs} , respectively, and they must not intersect. After discretizing the workspace, the algorithm is able to solve problems for C'_{free} where $w(C'_{free}) \geq w(C_{free}) + f(s)$. Thus, the algorithm is resolution complete for all $C_{free} \in \Psi(4b(N)M^{-\frac{1}{N}} + f(s))$. ■

To calculate $f(s)$, let $V(e)$ denote the voxelized swept volume of an edge e , and $C_{V(e)}$ be the C space region occupied by $V(e)$, then the width of C'_{free} can be defined as

$$w(C'_{free}) = 4b(N)M^{-\frac{1}{N}} + f(s) = \sup_{e \in \mathcal{E}} \{w(C_{V(e)})\}, \quad (1)$$

which yields

$$f(s) = \sup_{e \in \mathcal{E}} \{w(C_{V(e)})\} - 4b(N)M^{-\frac{1}{N}}. \quad (2)$$

It is practically difficult to pre-determine $f(s)$ before sampling as it depends not only on the number of samples M and resolution s , but also on the robot's geometric shape.

C. Limitations: Curse-of-Dimensionality

According to **Lemma 1**, although certain level of resolution completeness can be guaranteed for any given workspace and configuration space resolution, the algorithm would not be practically useful if $w(C'_{free})$ is too wide. To achieve a smaller $w(C'_{free})$, one could either increase the number of samples or the workspace resolution. However, both options are restricted by the available memory size.

Consider a 6-DoF robotic manipulator where for example $K = 15$ different discretization values are chosen for each

Fig. 4: A 2-DoF example of HDRM with $K_1 = K_2 = 3$, the number of roadmap vertices stored in the structure is $K_1 \times K_2 = 9$.

joint. In this case we would generate $15^6 \approx 11.4$ million vertices to build a DRM. If each vertex occupies around 400 workspace voxels, that means the occupation lists may contain billions of indices (unsigned int¹), which is too expensive to store on commodity hardware. Note that this is only a 6-DoF case and the problem scales exponentially with the dimension of the configuration space (curse-of-dimensionality).

III. HIERARCHICAL DYNAMIC ROADMAP

In this section, we introduce the Hierarchical Dynamic Roadmap which reduces $w(C'_{free})$ by exploiting the inherit hierarchical structure of the robot, which in turn enables us to store a higher resolution deterministic roadmap, e.g. $s \leq 5\text{cm}$, given a limited amount of available memory.

A. Hierarchical Configuration Structure

Let $[b_{n,l}, b_{n,u}]$ be the lower and upper bounds of joint $n \in N$ of a N -DoF robot. An even discretization of the n -th joint to $K_n \in \mathbb{N}$ values results in configurations

$$q_n(k_n) = b_{n,l} + (k_n - 1) \times \frac{b_{n,u} - b_{n,l}}{K_n - 1} \quad (3)$$

where $k_n \in \{1, \dots, K_n\}$. Let

$$\mathbf{k}(n) = [k_1, \dots, k_n] \quad (4)$$

be a n -dimensional vector containing the joint value indices for the first n joints, and

$$\mathbf{q}(\mathbf{k}(n)) = [q_1(k_1), \dots, q_n(k_n)] \quad (5)$$

be a n -dimensional vector contains the actual joint values corresponding to $\mathbf{k}(n)$. The full N -dimensional robot configuration can be retrieved given $\mathbf{k}(N)$. For example, as shown in Fig. 4, consider a 2-DoF robot, where the range of motion of each joint is $[-\pi, \pi]$. Given $K_1 = K_2 = 3$, we have $q_1(1) = q_2(1) = -\pi$, $q_1(2) = q_2(2) = 0$, and $q_1(3) = q_2(3) = \pi$. Then, $\mathbf{k}(2) = [1, 1]$ gives the robot configuration $\mathbf{q} = [-\pi, -\pi]$, $\mathbf{k}(2) = [2, 3]$ gives another robot configuration $\mathbf{q} = [0, \pi]$, and $\mathbf{k}(1) = [2]$ gives the first joint value $q_1 = 0$.

The data structure stores the equivalent of $M = \prod_{n=1}^N K_n$ vertices. The robot configurations can be accessed with $\mathbf{k}(n)$, however, the vertices in the roadmap are indexed with one integer index $i \in M$. Let $\mathcal{H} : (n, i) \mapsto \mathbf{k}(n)$ be the map from pair (n, i) to $\mathbf{k}(n)$, and $\mathcal{H}^{-1} : \mathbf{k}(n) \mapsto (n, i)$ be the corresponding inverse map. Given an index i and level n , the first n indices $\mathbf{k}(n) = \mathcal{H}(n, i)$ can be efficiently calculated using Algorithm 1. Similarly, given hierarchical indices $\mathbf{k}(n)$, the corresponding (n, i) can also be found by Algorithm 2.

¹The size of one unsigned int is 4 Byte on 64-bit operating systems. Hence storing indices for the occupation lists for these vertices in this example requires $4 \times 400 \times 11.4 \times 10^6$ Byte ≈ 17 GB of memory. 400 is the number of voxels occupied by one sample of the LWR robot at $s = 5\text{cm}$ resolution.

Algorithm 1 Obtain hierarchical indices from integer index**Require:** Dimension level n , vertex index i **Ensure:** Hierarchical indices $\mathbf{k}(n)$

```

1: Quotient =  $i$ 
2: while  $n > 1$  do
3:   Quotient, Remainder = Division(Quotient,  $\prod_{l=1}^n K_n$ )
4:    $k_n$  = Remainder
5:  $k_1$  = Quotient
   return  $\mathbf{k}(n) = [k_1, \dots, k_n]$ 

```

Algorithm 2 Obtain integer index from hierarchical indices**Require:** Hierarchical indices $\mathbf{k}(n) = [k_1, \dots, k_n]$ **Ensure:** Dimension level n , vertex index i

```

1:  $i = 0$ 
2: for  $l \in \{1, \dots, n-1\}$  do
3:   Counter = 1
4:   for  $j \in \{l+1, \dots, n-1\}$  do
5:     Counter = Counter  $\times K_j$ 
6:    $i = i + \text{Counter} \times k_l$ 
7:  $i = i + k_n$ 
   return  $n, i$ 

```

The hierarchical structure is a multi-resolution Sukharev grid [20] with $b(N) \equiv 1$, thus the hierarchical roadmap is resolution complete for $\mathcal{C}_{free} \in \Psi(4M^{-\frac{1}{N}} + f(s))$. As we will show in Section III-C, such structure enables a hierarchical way for storing the configuration-to-workspace occupation information that dramatically reduces the memory consumption.

B. Removal of Swept Volumes

There are two types of occupation information: the occupation voxels of a vertex and the swept volume of an edge (dark and light grey voxels in Fig. 5 respectively). In the classical DRM algorithm, the edge is invalidated if one or more voxels in the swept volume are in collision. However, there will be many sub-edges still valid in the cases where only very few of the voxels are in collision. For example, in Fig. 5a, if only the red voxel is in collision, the long edge $\mathcal{E}(\mathbf{q}_1, \mathbf{q}_4)$ is invalid while the sub-edge $\mathcal{E}(\mathbf{q}_3, \mathbf{q}_4)$ is still valid. Yet, the whole edge is considered invalid as these sub-edges are not stored in the roadmap. This is the underlying reason for planning failures and the low success rate of the classical DRM method, i.e. the $f(s)$ in (2) is too large to make it practically useful. In fact, we argue that storing swept volumes does not utilize available resources well as this information is used only for collision checking and not the actual path planning. Therefore, we instead propose to use the memory for storing more vertices and edges instead.

In our method, we store only the occupation voxels of the vertices excluding the swept volumes of the edges, while still being able to check the collision status of both vertices and edges. Let O_a be the occupation voxels of a vertex a , and $O_{a,b}$ be the swept volume of edge $\mathcal{E}(a, b)$. If two vertices a, b

Fig. 5: Comparison of classical and hierarchical DRM. (a) A long edge $\mathcal{E}(\mathbf{q}_1, \mathbf{q}_4)$ in classical DRM sweeps through a large number of workspace voxels. (b) Dense vertices and short edges in HDRM.

are very close and the edge is so short that

$$O_{a,b} = O_a \cup O_b, \quad (6)$$

then we do not need to store the swept volume of the edge since it can be represented by the occupation voxels of these two end-point vertices, as illustrated in Fig. 5b. The edge $\mathcal{E}(a, b)$ is collision-free if vertices a, b are collision-free, and vice versa. This ensures that a colliding workspace voxel only affects those corresponding short edges without invalidating others. A lower bound of K_n needs to be met in order to achieve such roadmap density.

Let $\theta_n = b_{n,u} - b_{n,l}$ be the range of motion of joint n , l_n and r_n be the approximate length and radius of n -th robot link, as illustrated in Fig. 6. For joint n , set all the subsequent/child joints such that the rest of the robot kinematic chain is fully extended. We assume that only one joint moves at a time, in which case, the distance between the n -th link of two neighboring configurations must not be greater than $s + \sqrt{2}r_n$, so that two end-effector links occupy the same or neighboring workspace voxels in order to satisfy (6). Hence, the inequality constraint shall be

$$\sup_{n \leq k \leq N} \left\{ \frac{s + \sqrt{2}r_k}{2\pi L_{n,k}} \right\} \geq \frac{\Delta_n}{2\pi}, \quad (7)$$

where $L_n^k = \sum_{j=n}^k l_j$ is the fully extended length from link n to link k , $k \geq n$. Rearranging terms yields

$$\Delta_n \leq \sup_{n \leq k \leq N} \left\{ \frac{s + \sqrt{2}r_k}{L_n^k} \right\}. \quad (8)$$

So joint n should have evenly distributed values of

$$K_n = \lceil \frac{\theta_n}{\Delta_n} + 1 \rceil = \lceil \theta_n \left(\sup_{n \leq k \leq N} \left\{ \frac{s + \sqrt{2}r_k}{L_n^k} \right\} \right)^{-1} + 1 \rceil \quad (9)$$

within its range of motion. We choose the minimum valid value for each K_n which already guarantees resolution completeness for a certain workspace voxel resolution s . Greater K_n only introduce more vertices and edges that increase memory consumption and slow down the search process.

The swept volumes can be removed if (9) is true for all joints. Furthermore, the information of the hundreds of millions of edges itself can be removed as well, because all the edges can be calculated analytically from the hierarchical

Fig. 6: Illustration of the maximum discretization step Δ_n the n -th joint can take without violating (6).

structure. Additionally, we add an assumption that one joint moves at a time. A N -dimensional configuration $\mathbf{k}(N)$ has $2 \times N$ neighbors (apart from the ones at the boundary of the range of motion), each of which forms an edge with $\mathbf{k}(N)$. As the edges have no direction, a N -dimensional HDRM with M vertices contains roughly $N \times M$ edges. It is also possible to allow full connectivity to all neighbors and thus the ability to move multiple joints at a time, but it will require a higher sampling density to achieve the same $\Psi(x)$.

C. Hierarchical Occupation Lists

Section III-A describes how to create a hierarchical structure to efficiently store tens of millions of configurations, and Section III-B explains why and how to remove the swept volumes and edges. The final step involves processing and storing the occupation lists of all vertices. However, when the roadmap contains tens or potentially hundreds of millions of vertices, their occupation lists are too expensive to store using classical methods. Next, we discuss how to exploit the hierarchical structure to resolve this problem.

Let $\mathcal{B}_n$ be the collision body between joint n and $n + 1$. Consider K_N configurations with identical values for the first $N - 1$ joints but only differing at the last joint, as illustrated in Fig. 7. These K_N configurations are invalid if $\mathcal{B}_{N-1}$ is in collision at the red voxel. In the classical DRM method, the red voxel's occupation list needs to store K_N indices to encode this information where each index corresponds to a particular configuration—which is very inefficient.

Instead of storing integer indices $i \in M$ for each configuration, we store a list of pairs (n, i) , where $i \in \prod_{1}^n K_n$ and $n \in N$. A pair (n, i) is added to a workspace voxel v 's occupation list if $\mathcal{B}_n$ at configuration $\mathbf{k}(n) = \mathcal{H}(n, i)$ occupies this voxel. In Fig. 7, when the red voxel is in collision with the environment, based on the pair $(N - 1, i)$, we can invalidate the i -th vertex of level $N - 1$ of the hierarchical structure. It is clear that all these K_N configurations are invalid since the first $N - 1$ joints have already caused body $\mathcal{B}_{N-1}$ to be in collision. Hence, we can encode the occupation information of K_N configurations using only two rather than K_N indices. Consider another case with $K_2 \times \dots \times K_N$ vertices, which could be millions, that have same value k_1 for the first joint but differ at all other joints. If k_1 puts $\mathcal{B}_1$ to a colliding position with the environment, then the millions of vertices with same k_1 index are all invalid. In such case, we can more efficiently use only a pair $(1, k_1)$ instead of millions of indices to encode

Fig. 7: Illustration of hierarchical occupation lists.

the occupation information of all these vertices. As we will show later in Section IV-B, using the hierarchical structure and this novel indexing technique, we can significantly reduce the memory required for storing the occupation information.

Algorithm 3 shows the details of generating the full occupation lists for all workspace voxels. First, given the size of the workspace and grid resolution s , a set of workspace voxels $\mathbb{V}$ can be generated. Each voxel $v \in \mathbb{V}$ is associated with an empty occupation list $\mathcal{O}_v$. Lines 3-12 generate the initial hierarchical occupation lists, but we can compress these to further reduce memory storage (lines 13-23). The compression is based on the fact that some robots, or part of the robots, are axially symmetric, which means that rotating a joint will not change the occupation list of the subsequent links. More generally, if the collision body $\mathcal{B}_n$ of K_n vertices ($xK_n + 1$ to $xK_n + K_n$, $x \in \mathbb{N}$) from the same sub-tree of level n occupies a voxel v , then the occupation list of v needs to store only one pair of $(n - 1, \cdot)$ rather than K_n pairs of $(n, \cdot)$, because the first $n - 1$ joints already make $\mathcal{B}_n$ unavoidably occupy voxel v . We “promote” the occupation list from level n to $n - 1$ if such axial symmetry occurs.

D. Motion Planning using HDRM

With the HDRM created and loaded, our goal is to efficiently solve motion planning queries online in different environments. The three main steps are as follows.

1) *Collision update*: First, we create a voxelized environment to represent the discretized workspace, and then apply conventional collision checking on this voxelized environment against the real collision environment to find the list of voxels that are occupied by the obstacles. For each occupied voxel, we iterate through its occupation lists and invalidate vertices in the hierarchical structure accordingly.

2) *Connecting start/goal to roadmap*: The start and goal vertices $\mathcal{V}_{start}, \mathcal{V}_{goal}$ are required for the graph search algorithm, which are the closest valid vertices to the start and goal configurations $\mathbf{q}_{start}, \mathbf{q}_{goal}$. Traditionally, this involves comparing the distance between a given configuration $\mathbf{q}$ and all vertices in the roadmap and finding the one with the shortest distance. As we will show later, such process could be very slow for a roadmap with a large number of vertices. In our approach, instead of searching through all vertices, we can analytically compute the closest one. Given a configuration $\mathbf{q}$, we can easily get the closest hierarchical configuration $\mathbf{k}_{closest}(N)$. Then, the index of closest vertex $\mathcal{V}_{closest}$ can be found using Algorithm 2.

Algorithm 3 Generate hierarchical occupation lists**Require:** Robot model $\mathcal{R}$, voxelized workspace $\mathbb{V}$ **Ensure:** Hierarchical occupation lists $\mathcal{O}_v, v \in \mathbb{V}$

```

1: for  $v \in \mathbb{V}$  do
2:   Occupation list  $\mathcal{O}_v = \emptyset$ 
3: for  $n \in \{1, \dots, N\}$  do
4:   for  $i \in \{1, \dots, \prod_1^n K_n\}$  do
5:      $\mathbf{k}(n) = \mathcal{H}(n, i)$ 
6:     Set first  $n$  joints of  $\mathcal{R}$  to  $\mathbf{q}(\mathbf{k}(n))$ 
7:     if  $\{\mathcal{B}_1, \dots, \mathcal{B}_n\}$  are NOT in self-collision then
8:        $V = \text{findBodyOccupiedVoxels}(\mathbb{V}, \mathcal{R}, \mathcal{B}_n)$ 
9:       for  $v \in V$  do
10:         $\mathcal{O}_v = \mathcal{O}_v \cup \{(n, i)\}$ 
11:   else
12:     Set vertex  $(n, i)$  as default invalid
13: for  $v \in \mathbb{V}$  do
14:   for  $n = N$  to 1 do
15:      $\mathbf{O} = \text{extractListOfDimension}(\mathcal{O}_v, n)$ 
16:     Remove duplicated indices and sort  $\mathbf{O}$ 
17:     for  $O_i \in \mathbf{O}$  do
18:       if  $O_i \bmod K_n = 0$  &  $O_{i+K_n} = O_i + K_n$  then
19:          $\mathcal{O}_v = \mathcal{O}_v \setminus \{(n, p) | p \in [O_i, \dots, O_{i+K_n}]\}$ 
20:         if  $n > 1$  then
21:            $\mathbf{k}(n) = [k_1, \dots, k_n] = \mathcal{H}(n, O_i)$ 
22:            $\mathcal{O}_v = \mathcal{O}_v \cup \{(n-1, \prod_{p=1}^{n-1} k_p)\}$ 
23:          $i = i + K_n - 1$ 

```

3) *Shortest path searching*: The last step is to find a valid path on the roadmap connecting $\mathcal{V}_{start}$ and $\mathcal{V}_{goal}$. The A* shortest path searching algorithm is used. We implemented the sequential version of A* using a single CPU-thread. Parallelization is not the main focus of this paper, however, we believe that parallel version of Dijkstra or A* algorithms would be more efficient [15].

IV. EXPERIMENTS

The proposed HDRM method is benchmarked against the classical DRM approach and standard sampling-based planners (SBP) in various scenarios using two different fixed-base manipulators: a 6-DoF Universal Robot UR5 and a 7-DoF KUKA LWR. The evaluation was performed using an Intel Core i7-6700K 4.0 GHz CPU with 32 GB 2133 MHz RAM, and the KUKA LWR for hardware experiments.

A. Experimental Setup

Given the robot model, first K_n is calculated using equations 7–9, as shown in Table I. Two different workspace voxel resolutions are used, $s = 0.1m$ and $s = 0.05m$. Smaller s leads to greater K_n and more samples are required to densely cover the space. We have also implemented classical DRM methods for comparison. To achieve resolution completeness, we generate the vertices by uniformly sampling in the configuration space and apply no roadmap compression technique. Three classical DRM datasets are created with different number of vertices: 1,000 (DRM_a), 10,000 (DRM_b) and 200,000 (DRM_c). A K-nearest neighbor search based on

TABLE I: Robot kinematic analysis for creating HDRM.

Robot	Range of motion θ (rad), extend length R (m)	s (m)	K_n
UR5	$\theta_n = \{6.28, 6.28, 6.28, 6.28\}$	0.1	$\{37, 36, 21, 9, 7, 1\}$
	$R_n = \{0.98, 0.97, 0.57, 0.23, 0.17, 0.0\}$	0.05	$\{52, 51, 30, 12, 9, 1\}$
LWR	$\theta_n = \{5.86, 4.12, 5.86, 4.12, 5.86, 4.12, 5.86\}$	0.1	$\{35, 20, 21, 10, 7, 2, 1\}$
	$R_n = \{0.99, 0.79, 0.59, 0.39, 0.19, 0.05, 0\}$	0.05	$\{49, 27, 29, 14, 10, 2, 1\}$

TABLE II: Comparison between classical DRM and HDRM.

Robot & Method	No. Vertices	No. Edges	s (m)	Memory size (MB)
UR5 DRM	1,000	6,336	0.1	2.8
			0.05	13.6
	10,000	61,274	0.1	22.3
			0.05	104
	200,000	1,200,956	0.1	356
			0.05	1593
UR5 HDRM	1,762,236	10,573,416	0.1	8.5
	8,592,480	51,554,880	0.05	145
LWR DRM	1,000	6,369	0.1	7.9
			0.05	33.4
	10,000	62,031	0.1	70
			0.05	280
	200,000	1,216,755	0.1	1239
			0.05	4793
LWR HDRM	2,058,000	14,406,000	0.1	16.7
	10,742,760	75,199,320	0.05	266

the Euclidean distance with $K = 10$ is then applied to find the edges in the roadmap.

B. Memory Requirements

As highlighted in Table II, the HDRM scales exponentially with roadmap size while the memory requirement is much lower compared to classical DRM. In the case of the UR5 robot with 0.1m voxel size, HDRM can store over 1.7 million vertices and 10 million edges using only 8.5 MB of memory, which is even less than the memory required for classical DRM to store only 10,000 vertices. In the LWR scenario with 0.05m voxel size, the HDRM stores over 10 million vertices and up to 75 million edges using only 266 MB of memory whereas the memory size for classical DRM to store 200,000 vertices is over 4.7 GB.

C. Motion Planning Evaluation

We benchmarked the performance of eight candidate methods: three classical DRM methods DRM_a, DRM_b, DRM_c; four standard sampling-based planners (SBP), i.e. RRT, PRM, Single-Query Bi-directional Planning with Lazy Collision Checking (SBL) [9], and RRTConnect; and finally the proposed HDRM. For DRM/HDRM, the datasets of the LWR robot with 0.1m voxel resolution are used. For SBP methods, we use the standard implementations from the OMPL library [22] and the FCL library [23] for explicit online collision

Fig. 8: Random problems in environments with different workspace obstacle densities. The highlighted trajectories are valid solutions found by HDRM.

checking with all parameters set to library defaults.

Five different categories of environments with random obstacles were created. From simple to complex, these environments have 0%, 0.1%, 0.5%, 1% and 5% of the whole workspace occupied by obstacles, where the latter four are illustrated in Fig. 8. The environments with 1% and 5% obstacle densities are extremely complicated for any kind of motion planning algorithm. We have created 1000 random valid problems for each category. Each problem has a valid start and goal states, and at least one trajectory connecting these two. To guarantee this, a random self-collision-free trajectory is generated in an empty space, which is then populated with obstacles that are not colliding with the trajectory. All algorithms are given 10 seconds to solve each problem. Since all problems are solvable, reporting no solution or exceeding the time limit is considered a failure. Solutions that differ from the one used to create the problem are also valid.

Evaluation results are detailed in Table III. As a baseline, all algorithms achieve 100% success rate in free space. The success rate of the classical DRMs falls below 100% when the environment is populated with only a few obstacles (0.1% obstacle density). The SBP methods are generally slower due to explicit collision-checking for every sample which is very time consuming. In more complicated environments (0.05% and 1% obstacle densities), the success rate of classical DRM methods decreases significantly. SBP methods still achieve reasonable success rates, however, the planning time increases considerably. HDRM performs better compared to all other methods in complicated scenarios in terms of both success rate and planning time. In the extreme cases with 5% obstacle density, we do not show the average planning time for the classical DRM methods as the success rate is too low. All SBP methods also report lower success rates and much longer planning times. On the contrary, HDRM constantly achieves 100% success rate in these extremely constrained environments.

It is interesting that DRM_c has a much smaller roadmap than HDRM, but takes longer time to find a solution even in free space. We break down the DRM/HDRM planning time into separate components, as in Table IV, and the time is given in microseconds. The collision update takes $141.6\mu\text{s}$ in free space, which is basically the overhead of

TABLE IV: Breakdown of computational time (in **microseconds**).

Method	Roadmap update		Planning		Total
	Find coll. voxels	Remove invalids	Connect to roadmap	A* search	
Obstacle density 0%					
DRM _a	141.6	0	14.17	0.581	156.4
DRM _b			170.1	0.626	312.3
DRM _c			3273	0.698	3415
HDRM			0.229	134.1	275.9
Obstacle density 0.1%					
DRM _a	694.8	1.665	13.774	0.557	710.8
DRM _b		18.71	163.8	0.601	877.9
DRM _c		435.0	3212	0.711	4343
HDRM		17.48	0.267	144.9	857.4
Obstacle density 1%					
DRM _a	1212	13.92	13.32	0.607	1240
DRM _b		145.5	163.5	0.745	1522
DRM _c		4728	3290	1.021	9231
HDRM		177.9	0.233	2325	3715

(a) Reaching into confined shelf. (b) Fetching object in distance.

Fig. 9: Experiments on a 7-DoF KUKA LWR robotic arm fitted with SCHUNK Dexterous Hand.

communication and function calls. We use FCL for explicit collision checking where the time increases as expected in more complicated environments. Classical DRMs with more vertices and edges require much longer time to remove invalid roadmap parts, whereas the HDRM is able to do so relatively faster, considering the enormous number of vertices and edges. Another expensive step of classical DRMs is connecting the roadmap, which increases exponentially with the number of vertices. After connecting the roadmap, running A* search is actually very fast since the roadmap size is relatively small. On the other hand, the time for connecting the roadmap is negligible for HDRM since the closest vertices can be analytically computed as elaborated in Section III-D.2. However, the searching takes longer due to the enormous roadmap size.

D. Experimental Validation on Robot Hardware

We further validate the HDRM method on a 7-DoF KUKA LWR manipulator fitted with the SCHUNK Dexterous Hand 2.0. The HDRM dataset was generated with $s = 0.05m$ voxel resolution. Four Microsoft Kinect One RGB-D sensors were fused to sense the environment and create an OctoTree representation for collision checking. It shall be noted that due to the limitation of the out-of-the-box sensing devices and algorithms, the real experiment setup was limited to simpler environments compared with the simulation bench-

TABLE III: Evaluation of 8 candidate motion planners, showing the success rate of solving 1000 problems in random valid environments, followed by the average solving time over the successful cases (in **milliseconds**). All algorithms are given 10 seconds to solve each problem.

		Obstacle density 0%		Obstacle density 0.1%		Obstacle density 0.5%		Obstacle density 1%		Obstacle density 5%	
Method		Success rate	Time (ms)	Success rate	Time (ms)	Success rate	Time (ms)	Success rate	Time (ms)	Success rate	Time (ms)
Standard	RRT	100%	13.392	99%	22.708	92%	325.21	82%	1036.1	36%	1893.6
	PRM	100%	5.7416	100%	4.5041	100%	322.86	99%	656.09	34%	3386.4
	SBP	100%	6.8909	100%	14.775	100%	82.473	100%	273.96	31%	4439.4
	RRT _{Connect}	100%	1.1930	100%	2.1220	100%	10.117	100%	48.926	74%	1723.7
Classical DRM	DRM _a	100%	0.1564	92.2%	0.7108	65.6%	0.7911	39.0%	1.2403	1.6%	-
	DRM _b	100%	0.3123	93.9%	0.8779	69.5%	1.0203	48.9%	1.5221	3.6%	-
	DRM _c	100%	3.4152	95.7%	4.3431	74.7%	6.5206	53.0%	9.2316	3.3%	-
Hierarchical DRM		100%	0.2759	100%	0.8574	100%	1.5813	100%	3.7152	100%	15.506

mark. In our supplementary video (<https://youtu.be/4AzbmiTIIIE>), we demonstrate challenging motions in three different, highly constrained environments: reaching into a confined shelf space and grasping a target object (Fig. 9a); retrieving an object through a frame (Fig. 9b); and moving an object from within a cage (Fig. 1).

V. CONCLUSION

This paper presents a novel method, the Hierarchical Dynamic Roadmap (HDRM), for real-time motion planning in high dimensions. The HDRM, through our proposed novel indexing scheme, is able to encode large numbers of vertices and edges (up to tens of millions) in a memory efficient manner that also allows the algorithm to be resolution complete. An extensive benchmarking shows that HDRM can find valid motion plans in extremely complicated environments in real-time and empirically validates that the algorithm is resolution complete. Experiments on the KUKA LWR robot further demonstrate that our method is capable of incorporating live sensing information and providing collision-free trajectories suitable for tackling practical problems.

Both DRM and HDRM compute solutions in static environments, which can be different between planning queries but need to be static during execution—they inherently do not adapt to runtime changes. Since HDRM guarantees resolution completeness and is able to plan in real-time (few milliseconds or less), the future work will focus on implementing a closed-loop online adaptation/re-planning framework for applications such as real-time interaction between human and robot in a shared workspace similar to [24].

REFERENCES

- [1] J. Schulman, Y. Duan, J. Ho, A. Lee, I. Awwal, H. Bradlow, J. Pan, S. Patil, K. Goldberg, and P. Abbeel, "Motion planning with sequential convex optimization and convex collision checking," *Intl. J. of Robotics Research*, vol. 33, no. 9, pp. 1251–1270, 2014.
- [2] V. Ivan, D. Zarubin, M. Toussaint, T. Komura, and S. Vijayakumar, "Topology-based representations for motion planning and generalization in dynamic environments with interactions," *Intl. J. of Robotics Research*, vol. 32, no. 9–10, pp. 1151–1163, 2013.
- [3] J. J. Kuffner and S. M. LaValle, "RRT-connect: An efficient approach to single-query path planning," in *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, vol. 2, 2000, pp. 995–1001.
- [4] L. E. Kavraki, P. Svetska, J.-C. Latombe, and M. H. Overmars, "Probabilistic roadmaps for path planning in high-dimensional configuration spaces," *IEEE Trans. Robotics and Automation*, vol. 12, no. 4, pp. 566–580, 1996.
- [5] M. Elbanhawi and M. Simic, "Sampling-based robot motion planning: A review," *IEEE Access*, vol. 2, pp. 56–77, 2014.
- [6] Y. Yang, V. Ivan, W. Merkt, and S. Vijayakumar, "Scaling Sampling-based Motion Planning to Humanoid Robots," in *IEEE Intl. Conf. on Robotics and Biomimetics (ROBIO)*, 2016, pp. 1448–1454.
- [7] D. Hsu and Z. Sun, "Adaptively combining multiple sampling strategies for probabilistic roadmap planning," in *IEEE Conf. on Robotics, Automation and Mechatronics*, vol. 2, 2004, pp. 774–779.
- [8] R. Bohlin and L. E. Kavraki, "Path planning using lazy PRM," in *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, vol. 1, 2000, pp. 521–528.
- [9] G. Sánchez and J.-C. Latombe, "On delaying collision checking in PRM planning: Application to multi-robot coordination," *Intl. J. of Robotics Research*, vol. 21, no. 1, pp. 5–26, 2002.
- [10] J. Pan and D. Manocha, "GPU-based parallel collision detection for fast motion planning," *Intl. J. of Robotics Research*, vol. 31, no. 2, pp. 187–200, 2012.
- [11] P. Leven and S. Hutchinson, "A framework for real-time path planning in changing environments," *Intl. J. of Robotics Research*, vol. 21, no. 12, pp. 999–1030, 2002.
- [12] M. Kallman and M. Mataric, "Motion planning using dynamic roadmaps," in *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, vol. 5, 2004, pp. 4399–4404.
- [13] A. Voelz and K. Graichen, "Distance Metrics for Path Planning with Dynamic Roadmaps," in *Intl. Symposium on Robotics*, 2016, pp. 1–7.
- [14] S. M. LaValle, *Planning algorithms*. Cambridge University Press, 2006.
- [15] H. Schumann-Olsen, M. Bakken, Ø. H. Holhjem, and P. Risholm, "Parallel dynamic roadmaps for real-time motion planning in complex dynamic scenes," *IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, 2014.
- [16] S. Murray, W. Floyd-Jones, Y. Qi, D. Sorin, and G. Konidaris, "Robot motion planning on a chip," in *Robotics: Science and Systems*, 2016.
- [17] Y. Yang, V. Ivan, Z. Li, M. Fallon, and S. Vijayakumar, "iDRM: Humanoid motion planning with realtime end-pose selection in complex environments," in *IEEE Intl. Conf. Humanoid Robots*, 2016, pp. 271–278.
- [18] Y. Yang, W. Merkt, H. Ferrolho, V. Ivan, and S. Vijayakumar, "Efficient humanoid motion planning on uneven terrain using paired forward-inverse dynamic reachability maps," *IEEE Robotics and Automation Letters*, vol. 2, no. 4, pp. 2279–2286, 2017.
- [19] P. E. Hart, N. J. Nilsson, and B. Raphael, "A formal basis for the heuristic determination of minimum cost paths," *IEEE Trans. on Systems Science and Cybernetics*, vol. 4, no. 2, pp. 100–107, 1968.
- [20] S. M. LaValle, M. S. Branicky, and S. R. Lindemann, "On the relationship between classical grid search and probabilistic roadmaps," *Intl. J. of Robotics Research*, vol. 23, no. 7–8, pp. 673–692, 2004.
- [21] J.-C. Latombe, *Robot Motion Planning*. Springer, 1991.
- [22] I. A. Sucan, M. Moll, and L. E. Kavraki, "The Open Motion Planning Library," *IEEE Robotics & Automation Magazine*, vol. 19, no. 4, pp. 72–82, 2012.
- [23] J. Pan, S. Chitta, and D. Manocha, "FCL: A general purpose library for collision and proximity queries," in *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, 2012, pp. 3859–3866.
- [24] W. Merkt, Y. Yang, T. Stouraitis, C. E. Mower, M. Fallon, and S. Vijayakumar, "Robust shared autonomy for mobile manipulation with continuous scene monitoring," in *IEEE Intl. Conf. on Automation Science and Engineering (CASE)*, 2017, pp. 130–137.