

Occupancy Grid Mapping without Ray-Casting for High-resolution LiDAR Sensors

Yixi Cai, Fanze Kong, Yunfan Ren, Fangcheng Zhu, Jiarong Lin, Fu Zhang

Abstract—Occupancy mapping is a fundamental component of robotic systems to reason about the unknown and known regions of the environment. This article presents an efficient occupancy mapping framework for high-resolution LiDAR sensors, termed D-Map. The framework introduces three main novelties to address the computational efficiency challenges of occupancy mapping. Firstly, we use a depth image to determine the occupancy state of regions instead of the traditional ray-casting method. Secondly, we introduce an efficient on-tree update strategy on a tree-based map structure. These two techniques avoid redundant visits to small cells, significantly reducing the number of cells to be updated. Thirdly, we remove known cells from the map at each update by leveraging the low false alarm rate of LiDAR sensors. This approach not only enhances our framework’s update efficiency by reducing map size but also endows it with an interesting decremental property, which we have named D-Map. To support our design, we provide theoretical analyses of the accuracy of the depth image projection and time complexity of occupancy updates. Furthermore, we conduct extensive benchmark experiments on various LiDAR sensors in both public and private datasets. Our framework demonstrates superior efficiency in comparison with other state-of-the-art methods while maintaining comparable mapping accuracy and high memory efficiency. We demonstrate two real-world applications of D-Map for real-time occupancy mapping on a handle device and an aerial platform carrying a high-resolution LiDAR. In addition, we open-source the implementation of D-Map on GitHub to benefit society: github.com/hku-mars/D-Map.

Index Terms—Occupancy Mapping, LiDAR Perception, Range Sensing.

I. INTRODUCTION

IN recent years, advancements in light detection and ranging (LiDAR) sensors have led to the commercialization of lightweight, low cost, and high accuracy 3D LiDARs, raising tremendous popularity in various applications such as robotics [1]–[4], autonomous driving [5]–[7], 3D reconstruction [8]–[10], etc. Over the past decade, a clear trend has emerged to develop 3D LiDARs with a smaller size, longer detection range, and higher resolution, approaching image-level quality [11]. This trend has not only opened up new possibilities for deploying LiDARs in diverse applications but has also necessitated the development of new techniques to exploit the potential of LiDAR-based applications.

Manuscript received February 5, 2023; revised July 17, 2023; accepted September 29, 2023. This work is supported by the University Grants Committee of Hong Kong General Research Fund under Project 17206421 and DJI Donation. (Yixi Cai and Fanze Kong contributed equally to this work.) (Corresponding author: Fu Zhang.)

The authors are with Mechatronics and Robotic Systems (MaRS) Laboratory, Department of Mechanical Engineering, University of Hong Kong, Hong Kong SAR, China. (email: yxicai@connect.hku.hk; kongfz@connect.hku.hk; renfy@connect.hku.hk; zhufc@connect.hku.hk; jiarong.lin@connect.hku.hk; fuzhang@hku.hk)




Fig. 1. Our proposed framework D-Map served as a real-time high-resolution occupancy mapping module for an autonomous UAV exploration task in an ancient fortress. (a) The high-fidelity point cloud collected by UAV. (b) A bird-view of the scene. (c) The aerial platform carried a 128-channel LiDAR (OSI-128) to conduct the exploration task. The accompanying video of this paper is available on YouTube: youtu.be/m5QQPbkYYnA.

Occupancy grid mapping provides an efficient means for autonomous systems to navigate in unknown environments, which is crucial in a variety of applications such as obstacle avoidance [12, 13], path planning [14, 15], and autonomous exploration [16]–[18]. When adapted to LiDARs with increasing performance, occupancy grid mapping is faced with additional challenges in efficiency, especially in the two core steps: ray-casting and occupancy state updates.

Ray-casting is challenging because LiDARs provide increasingly dense depth measurements (e.g., over one million points per second) with long ranges (e.g., over 300 m). Therefore, the number of rays to be processed increases with the number of points. Secondly, the number of cells traversed by a single ray is affected by the sensor’s detection range, with LiDARs traversing dozens of times more cells than depth cameras with smaller detection range (e.g., less than 5 m). Consequently, the tremendous amount of cells results in a costly computation load. Finally, the dense LiDAR measurements can lead to redundant visits to the same cells in ray-casting, which can consume unnecessary computation resources in the subsequent occupancy state updates.

The second challenge in LiDAR-based occupancy mapping arises from the need for high-resolution occupancy maps to fully exploit the high accuracy of LiDAR sensors. This challenge involves the selection of map structures to achieve high-resolution occupancy mapping where a trade-off between computational and memory efficiency has to be balanced. The two commonly used map structures are grid-based [19] and tree-based [20]. Grid-based maps offer highly efficient updates in $\mathcal{O}(1)$ time complexity but suffer from extensive memory consumption when applied in large-scale environments or with high resolutions. Tree-based maps are more memory-efficient than grid-based maps, but the trade-off is the higher computational complexity when updating the tree. The update efficiency of a tree-based map structure is directly influenced by the tree height, which is determined by both the mapping environment and map resolution.

In addition to the challenges mentioned above, some essential features of LiDAR sensors have not yet been fully utilized in existing occupancy mapping approaches, such as the low false alarm rate which provides high confidence in identifying free and occupied space. The existing occupancy mapping approaches commonly utilize occupancy probabilities to handle the false detection of depth measurements. Nevertheless, in LiDAR-based occupancy mapping, the low false alarm rate (e.g., three out of a million points [21]) offers an opportunity to enhance efficiency by directly eliminating known space from the map without requiring any probability updates.

In this work, we present a novel mapping framework that effectively addresses the aforementioned issues in LiDAR-based occupancy mapping. Our contributions can be summarized as follows:

- We propose an occupancy state determination method based on depth image projection to alleviate the computation load in the traditional ray-casting technique. This projection-based approach enables occupancy state determination of cells at any size, allowing subsequent efficient updates in large-scale environments.
- We present a novel on-tree update strategy for updating occupancy states based on a hybrid map structure, providing a superior balance between computation and memory efficiency. The hybrid map structure stores unknown space on an octree, which enables a memory-efficient representation for the large unknown space, while the occupied space is stored on a hashing grid map. In terms of efficiency, the proposed strategy allows occupancy state determination of large cells on an octree, thereby avoiding unnecessary updates on small cells and increasing efficiency.
- We leverage the low false alarm rate of LiDAR measurements to directly remove cells with determined states (i.e., occupied or free) at each update. This approach renders our map structure a decremental property, for which we term our framework as D-Map, providing higher computational efficiency and less memory usage.
- We conduct an in-depth analysis of the accuracy of the proposed occupancy state determination method and the time complexity of updates and queries in D-Map. Specifically, we derive an analytical function to quantify

the accuracy loss relative to depth image resolution. The time complexity analysis of updates on D-Map provides theoretical support to our superior performance over state-of-the-art methods that rely on ray-casting.

- We make the implementation of D-Map available on Github: github.com/hku-mars/D-Map to promote the reproducibility and further development of our work.

The remainder of the paper is organized as follows. Section II provides an overview of related works in the field of occupancy mapping. We then present the complete mapping framework and the details of each key component in Sections III, IV, and V, respectively. Section VII presents the benchmark experiments and ablation studies conducted on open datasets. In Section VIII, we demonstrate two real-world experiments, followed by a discussion in Section X. Finally, we conclude this article in Section XI.

II. RELATED WORKS

In this section, we review the previous research on occupancy mapping and discuss their update methods.

A. Occupancy Mapping Approaches

We categorize the occupancy mapping approaches into two classes: continuous maps and discrete maps, based on their assumptions of modeling the environments.

Continuous maps assume an implicit correlation of locality in space and model environments through continuous occupancy distribution. In the early years, Thrun *et al.* introduced the concept of learning inverse sensor models to generate local occupancy maps by means of sampling [22]. In more recent years, O'Callaghan *et al.* proposed using Gaussian Process regression to learn a continuous representation of 2-dimensional environments, which has since been extended to 3-dimensional environments using Gaussian mixture models or sparse Gaussian processes to reduce computational load during training [23]–[25]. However, the time complexity of these methods is often prohibitive for real-time applications, with a computational complexity of $\mathcal{O}(N^3)$ in the number of depth measurements N . Therefore, considerable efforts have been devoted to improving the update efficiency of these methods [26]–[29]. Unlike previous approaches maintaining dense occupancy states, Duong *et al.* developed a sparse Bayesian formulation for occupancy mapping using a sparse set of relevance vectors [30]. While the aforementioned methods can be applied to LiDAR sensors, there is a substantial body of research focused on learning-based occupancy mapping using image sequences, including Occupancy Networks [31], Neural Radiance Fields (NeRF) [32], DeepSDF [33], Semantic Mapping [34, 35], NICE-SLAM [36], among others.

As opposed to continuous maps, occupancy grid mapping approaches assume independent occupancy at different locations, resulting in a discrete representation of environments using grids. Historically, a uniform grid map structure was imposed to discretize the environments and represents the occupancy probabilities independently [19, 37, 38]. However, this approach is impractical for generating high-resolution maps or working in large-scale environments for its significant

memory consumption. To address this limitation, tree-based map structures such as quadtree [39] and octree [40] were applied in occupancy grid mapping approaches with great effect. These structures recursively divide space into equal subspaces for updates and merge subspaces with equivalent states for a more compact representation [20, 41]. However, the updates on the tree-based maps are much more expensive. A voxel hashing technique [42] has recently been adapted to occupancy mapping [43, 44]. Though alleviating the memory consumption in uniform grids, the memory consumption remains prohibitive to high-resolution maps and large-scale environments. The use of discrete representations of environments has also been applied in Euclidean Signed Distance Field (ESDF), an alternative approach to occupancy mapping [45, 46]. However, the construction of ESDF heavily relies on ray-casting and encounters similar challenges to other existing occupancy mapping techniques.

While the continuous mapping approaches provide dense occupancy maps that can reason about measured and unmeasured areas, concerns remain about the reliability of the inferred regions, particularly in safety-critical robotics applications. Besides, continuous mapping approaches face significant challenges in real-time ability due to the complicated training and querying process, as integration over non-trivial space is required to update and query on a continuous map which necessitates timely numerical evaluation. Compared with the continuous mapping approaches methods, occupancy grid mapping is preferable in robotic applications with limited computation resources because of its higher update efficiency resulting from less complex map representation. Additionally, occupancy grid mapping is regarded as a more reliable technique, as it updates the occupancy map solely based on actual depth measurements without any inferences. Our D-Map is one kind of occupancy grid mapping approach. It achieves higher computational efficiency than existing grid-based approaches while maintaining comparable memory consumption with tree-based methods.

B. Update Methods

Regardless of different map representations, various research has been conducted to improve the efficiency of map updating, most of which focuses on alleviating the computation load of ray-casting. Wurm *et al.* present a hierarchical structure of octrees with multiple resolutions, allowing not complete but adequate resolution in different sub-maps depending on objects to reduce the computation load in ray-casting [47]. Similarly, [48] propose an adaptive resolution mapping method based on statistical measurements. However, the resultant maps generated from this method are affected by the parameters of adaptive resolution, which may differ from the original results. In Octomap [20], a batch-based method is utilized to avoid redundant cells resulting from ray-casting, which effectively reduces the effort in subsequent tree updates. Besides, a clamping policy is used in Octomap, initially proposed to handle dynamic environments [49], allowing the acceleration of map updates by pruning the octree. Duberg *et al.* proposed two possible approaches to simplify the ray-casting process, including downsampling point clouds to map

resolution and ray marching with adaptive steps. However, both of these approaches introduce accuracy loss and require handcraft parameters [50]. In contrast, motivated by the idea of coherent rays [51], Kwon *et al.* combine rays traversing equivalent cells into a super ray to reduce the number of rays in ray-casting [52] while maintaining full accuracy. Besides super rays, the number of cells to be traversed is reduced by culling regions. However, these two techniques provide less benefit at high resolution due to the ineffectiveness of merging rays into a super ray and the inefficiency of culling region construction.

In comparison with the aforementioned update approaches, D-Map discards the timely ray-casting for occupancy mapping. Alternatively, we design an occupancy state determination method based on depth image projection. We develop an on-tree update strategy for efficient map updates. This strategy determines the occupancy states of cells at various resolutions by projecting cells to the depth image, which is more reliable compared to the adaptive resolution mapping approach [48] that relies on handcraft parameters. Moreover, unlike existing occupancy mapping methods, D-Map does not require occupancy probability updates. Instead, we assume the environment is static and directly remove known space (i.e., free and occupied) from the tree map by leveraging the low false alarm rate of LiDAR sensors. This approach results in a decreasing size of the map during the update process, thus further enhancing the efficiency of D-Map.

III. OVERVIEW

Figure 2 presents an overview of the proposed D-Map framework. The map structure delineated in the orange block is explained in Section V-A. It consists of two parts: the occupied map and the unknown map. The green block represents the pipeline of the occupancy update strategy. At each update, a depth image is rasterized from the incoming point clouds at the sensor pose (see Section IV-A). Subsequently, a 2-D segment tree is constructed on the depth image to enable efficient occupancy state determination (see Section IV-B and Section IV-C). The entire procedure to update the unknown map is described in Section V-B and summarized as follows. The cell extraction module retracts the unknown cells on the octree from the largest to the smallest size, projects them to the depth image, and determines their occupancy states. The cells determined as known are directly removed from the map, while the unknown ones remain, and the undetermined ones are split into smaller cells for further occupancy state determination. This coarse-to-fine process facilitates updates on large cells directly without queries each small cell. Moreover, the removal of known cells endows our framework with a decremental property, providing high efficiency in both computation and memory.

IV. OCCUPANCY STATE DETERMINATION ON DEPTH IMAGE

This section describes how to determine the occupancy states on a depth image rasterized from incoming point clouds.




Fig. 2. The framework overview of D-Map. The blue block shows the input to D-Map, including the point clouds and the corresponding sensor odometry. The orange block is the occupancy map structure of D-Map, which is composed of a hashing grid map for maintaining occupied space and an octree for maintaining unknown space. The occupancy update strategy is presented in the green block, which extracts the cells inside the sensing area on the octree and conducts operations depending on the occupancy state determination method using a depth image.

A. Depth Image Rasterization

In preparation for occupancy state determination, the point cloud captured by a LiDAR sensor is rasterized into a depth image at the current sensor pose. To ensure the accuracy of state determination, the depth image resolution should be sufficiently small so that the projected area of a cell from the map to the depth image is larger than one pixel. As illustrated in Fig. 3, we determine the depth image resolution ψ_{map} relative to the map resolution d and LiDAR's detection range R using the following equation:

$$\psi_{\text{map}} = 2 \arcsin\left(\frac{d}{2R}\right) \approx \frac{d}{R} \quad (1)$$

However, high-resolution maps would result in a high-resolution depth image of enormous size, with many empty pixels due to the much smaller number of point clouds than the size of the depth image. To address this issue, we bound the depth image resolution by the LiDAR angular resolution, which is the minimum angle between two laser pulses emitted and received in a rotating manner. Specifically, we define the standard depth image resolution ψ_I as

$$\psi_I = \max\{\psi_{\text{map}}, \psi_{\text{lidar}}\} \approx \max\left\{\frac{d}{R}, \psi_{\text{lidar}}\right\} \quad (2)$$

where ψ_{lidar} is the angular resolution of the LiDAR. Note that we do not distinguish between the vertical and horizontal angular resolution for simplicity of definition. Moreover, we keep the minimum depth value when neighboring points are projected to the same pixel.

B. 2-D Segment Tree

To determine the occupancy state of a cell in the map, a two-step process is employed, whereby the cell is first projected onto the depth image, followed by a comparison of the projected depth against the minimum and maximum depth values of the corresponding area. Since the projected area of a cell on the depth image varies with the cell location, a 2-D segment tree structure is employed to expedite efficient queries of the minimum and maximum values on the depth image, as detailed below.




Fig. 3. This figure illustrates the spatial relationship among the map resolution d , the detection range R , and the depth image resolution ψ_{map} .

A segment tree is a perfectly balanced binary tree that efficiently provides range queries by representing a set of intervals [53]. Figure 4 describes the process of querying the minimum value by a 1-D segment tree. The segment tree is constructed by recursively splitting the array in half until each node contains a single element. As each node on the segment tree represents an interval of the array, the summarized information of the values in the interval, such as the minimum and maximum, is preprocessed during the construction procedure to accelerate the subsequent query. When querying, the segment tree retrieves a minimal representation of the queried interval using a subset of nodes (colored nodes in Fig. 4). The result is obtained by summarizing information from the retrieved nodes with fewer operations than direct queries. The time complexity of query on a 1-D segment tree is $\mathcal{O}(\log N)$ where N is the number of elements in the discrete array, while direct query leads to a time complexity of $\mathcal{O}(N)$.

The approach to extending a 1-Dimensional segment tree to a 2-Dimensional structure involves constructing a “segment tree of segment trees”, as proposed in [54]. The segment tree in the outer layer splits the 2-D array by row, and on each node of the outer segment tree, a 1-D inner segment tree is constructed to maintain the column information on the covered rows. The query on a 2-D segment tree first searches the outer tree for the nodes representing rows covered by the queried region and then traverses the inner segment trees on the corresponding




Fig. 4. This figure illustrates an example of a fast query of the minimum value in the pixel range of [2, 7] on a 1-D segment tree. Starting from the root of the segment tree, the range query searches along the tree recursively until the current node range is completely covered by the queried range, where the minimum value of the node range that has been saved on the node during the tree construction will be returned. In this example, the range [2, 7] leads to four nodes representing the range of [2, 2], [3, 4], [5, 6], and [7, 7], respectively. The minimum value of the range is efficiently obtained from these four nodes instead of counting the six elements in the array.

nodes to retrieve the covered columns. Finally, the result over the queried region is summarized from the information stored on retrieved nodes. The time complexity of a 2-D segment tree to query on a 2-D array of size $N \times M$ is $\mathcal{O}(\log N \log M)$, while direct queries lead to a time complexity of $\mathcal{O}(N^2)$.

Given a depth image rasterized from point clouds, we build a 2-D segment tree to maintain the minimum and maximum depth values on each tree node, denoted as d_{Min} and d_{Max} , respectively. Additionally, we keep track of the number of pixels occupied by point clouds within the covered area of each node, denoted as d_{Sum} .

C. Occupancy State Determination

We introduce the principle of our method for determining the occupancy state of a cell using five cells as an example, depicted in Fig. 5 and numbered from 1 to 5. We commence by classifying the cells based on whether they are completely located inside the LiDAR's sensing area. As illustrated in the top-down view in Fig. 5(b), Grid 1, Grid 2, and Grid 3 are entirely located within the sensing area, while Grid 4 and Grid 5 only have part of them inside. Among the cells completely inside, Grid 3 is determined as known since it is situated in front of all objects in the observed environment; Grid 1 is determined as unknown due to its location behind the objects. The occupancy state of Grid 2 remains undetermined since part of it lies in front of the objects while the other part lies behind. Regarding the cells with part of them inside the sensing area, Grid 5 is determined as unknown because it is located behind the objects. Though lying in front of the object, the occupancy

Algorithm 1: Occupancy State Determination

```

Params: Completeness threshold  $\varepsilon$ 
Input : Grid Center C, Grid Size L, LiDAR Pose T,
          2D Segment Tree  $\mathcal{T}$ 
Output: State S
1 Function DetermineOccupancy (T,C,L, $\mathcal{T}$ )
2    $r, \theta, \phi \leftarrow \text{Projection}(T, C);$ 
3    $\text{BoxMin} = r - L/2, \text{BoxMax} = r + L/2;$ 
4    $xL, xR, yL, yR \leftarrow \text{Coverage}(r, \theta, \phi, L);$ 
5    $\text{ProjPixels} = (yR - yL + 1)(xR - xL + 1);$ 
6    $d_{\text{Min}}, d_{\text{Max}}, d_{\text{Sum}} \leftarrow \mathcal{T}.\text{query}(xL, xR, yL, yR);$ 
7    $\alpha = d_{\text{Sum}}/\text{ProjPixels};$ 
8   if  $d_{\text{Sum}} == 0$  then return (Unknown);
9   if  $\alpha > \varepsilon$  then
10     if  $d_{\text{Max}} < \text{BoxMin}$  then
11       |  $S \leftarrow \text{Unknown};$ 
12     else
13       if  $d_{\text{Min}} > \text{BoxMax}$  then
14         |  $S \leftarrow \text{Unknown};$ 
15       else
16         |  $S \leftarrow \text{Undetermined};$ 
17       end
18     end
19   else
20     if  $d_{\text{Max}} < \text{BoxMin}$  then
21       |  $S \leftarrow \text{Unknown};$ 
22     else
23       |  $S \leftarrow \text{Undetermined};$ 
24     end
25   end
26   return (S);
27 End Function

```

state of Grid 4 remains undetermined owing to its location not completely inside.

Following the principles above, D-Map projects the cells that are either inside or intersected with the LiDAR's sensing area onto the current depth image rasterized from recent point clouds which represent the objects in the environment. The relative location between the cells and the objects is determined by comparing their depth values. The procedure of occupancy state determination on the depth image is described in Alg. 1 and explained below.

In consideration of consistency and computational efficiency, we project a cell to the depth image by its inscribed circle, as shown in Fig. 5 (b) and (c). The cell center, denoted as C , is first transformed into a spherical coordinate as (r, θ, ϕ) (Line 2). The depth range covered by the cell is described as its minimum depth BoxMin and maximum depth BoxMax with respect to its center C and cell size L (Line 3). The boundary of the projected area on the depth image (i.e., queried pixels in Fig. 5(c)) is obtained as follows

$$\begin{aligned}
 x_L &= \lfloor \frac{\theta - \arcsin(\frac{L}{2r})}{\psi_I} \rfloor, \quad x_R = \lceil \frac{\theta + \arcsin(\frac{L}{2r})}{\psi_I} \rceil \\
 y_L &= \lfloor \frac{\phi - \arcsin(\frac{L}{2r})}{\psi_I} \rfloor, \quad y_R = \lceil \frac{\phi + \arcsin(\frac{L}{2r})}{\psi_I} \rceil
 \end{aligned} \tag{3}$$




Fig. 5. This figure demonstrates an example of occupancy state determination on five cells in the map. (a) The 3D view shows the relative position of the five cells with respect to (w.r.t.) the LiDAR and the objects in the environment. (b) The top-down view helps to understand the occlusion between the cells and the objects when seeing from LiDAR. (c) The cells are projected to the depth image by their circumsphere radius, after which the projected areas are queried for the depth values in the depth image, which are finally used to determine the cells' occupancy states.




Fig. 6. The relative position between the cells and objects in the environments, along the one-pixel direction of the depth image, is determined by comparing the depth range of the cell (represented by the minimum depth $BoxMin$ and maximum depth $BoxMax$) with the depth range on the depth image (represented by the minimum depth $dMin$ and maximum depth $dMax$). Grids 1-3 completely locate inside the LiDAR's sensing area, while Grids 4-5 are partially inside.

Then, we prepare the following information for determining the occupancy states of the cell. Firstly, the number of queried pixels covered by the projected area is counted in $ProjPixels$ (Line 5). Secondly, the number of occupied pixels $dSum$, minimum depth $dMin$, and maximum depth $dMax$ within the projected area is provided by querying the 2-D segment tree as described in Section IV-B (Line 6). Thirdly, we define the observation completeness α of a cell to determine whether a cell is completely observed by the current depth image, which is calculated by dividing the number of occupied pixels $dSum$ over the number of queried pixels $ProjPixels$ on the depth image (Line 7). The computed α quantifies if the cell's projected area has been sufficiently observed by points in the current depth image. A large α indicates that a major part of the cell lies inside the sensing area, with most of the pixels in its projected area on the depth image actively occupied by LiDAR points. Only cells with large α have their occupancy




Fig. 7. A special case when the pixel size is larger than the projected area of a cell. This happens when the depth image resolution is computed from LiDAR's angular resolution.

state updated.

The occupancy state of a cell is acquired by comparing the depth range of the depth image (represented by the minimum depth $dMin$ and maximum depth $dMax$) against the depth range of the cell (represented by minimum depth $BoxMin$ and maximum depth $BoxMax$), as shown in Fig. 6 and detailed below. We first categorize the cells by comparing the observation completeness α with a completeness threshold ε . For completely observed cells (i.e., $\alpha > \varepsilon$ in Line 9), the occupancy state is determined as unknown if $dMax$ of the depth image is smaller than $BoxMin$ of the cell (Line 10~12). This condition indicates that the cell is located behind the object in the environment, as shown in Fig. 6(a). When $dMin$ of the depth image is greater than $BoxMax$ of the cell, the occupancy state of the cell is determined as known (Line 13~15), indicating that it is located in front of the objects as shown in Fig. 6(b). Otherwise, the occupancy state of the cell cannot be determined, as shown in Fig. 6(c). For cells that are not completely observed (i.e., $\alpha \leq \varepsilon$), the occupancy state is determined as unknown if $dMax$ of the depth image is less than $BoxMin$ of the cell (Line 20~22), indicating that the cell lies behind the objects as shown in Fig. 6(d). If this condition is not met, the occupancy state of the cell is considered undetermined due to incomplete observation (Line 22~24), as shown in

Fig. 6(e).

In some situations, the projected area of a cell may be smaller than the one-pixel area on the depth image, as shown in Fig. 7. This occurs when the depth image resolution ψ_I is computed from the LiDAR's angular resolution, as defined by (2). These cases are referred to as “single pixel cases”. In such cases, recalling that a pixel only saves one LiDAR point with the smallest depth as presented in Section IV-A, we should further determine whether the LiDAR point actually traverses the cell by examining if the point lies within the cell's projected area. If the LiDAR point lies outside the projected area, the occupancy state of the cell remains unknown. Otherwise, we use a similar logic to Alg. 1 to determine the occupancy state.

D. Depth Image Resolution Analysis


The resolution of a depth image is a critical parameter that plays a crucial role in balancing the accuracy and efficiency of occupancy state determination. A low-resolution depth image is more convenient to query but at the cost of increased information loss from the original point clouds. Conversely, a high-resolution depth image provides better accuracy but requires a larger computational effort, leading to reduced efficiency. When considering a depth image \mathcal{M}_I with a standard resolution of ψ_I determined by (2), it is possible to relax the resolution to $\psi = \gamma\psi_I$ by a relaxed factor γ to achieve higher efficiency while sacrificing some accuracy. To evaluate the accuracy loss with the resolution relaxation factor γ , we derive a function $f(\gamma)$ that describes the retained accuracy from the original depth image \mathcal{M}_I as a function of the factor γ .

We assume a LiDAR is employed to scan a large-scale open environment free of obstacles at a static pose. When determining the free space by projecting to the depth image \mathcal{M}_I with the standard resolution of ψ_I , the resultant occupancy map should show complete free space in a 3D spherical domain, the volume of which is denoted as V_I . However, when relaxing the depth image resolution to $\psi = \gamma\psi_I$ for a relaxed depth image \mathcal{M} , some free space is falsely determined as unknown when projecting to \mathcal{M} due to the single pixel cases explained in Section IV-C. We denote the volume of this case as V . An illustration of V_I and V is presented in Fig. 1 in the supplementary materials [55]. The accuracy function $f(\gamma)$ is defined as the ratio of space correctly determined as free:

$$f(\gamma) = \frac{V}{V_I} \quad (4)$$

We made the following assumption to ease the theoretical analysis:

Assumption 1. *We assume that the LiDAR is an idealized representation, which assumes a horizontal FoV of 360° and a symmetric vertical FoV in the range of $[-\alpha, \alpha]$ with infinitely small angular resolution (e.g., $\psi_{\text{lidar}} \rightarrow 0$) and a detection range R . The occupancy mapping resolution is d .*

Fig. 8. The simulation results for validation on accuracy function $f(\gamma)$.

Theorem 1. *The function $f(\gamma)$ is derived based on Assumption 1 as follows:*

$$f(\gamma) = \begin{cases} \frac{\gamma_0^3}{\gamma} + A(1 - \frac{\gamma_0}{\gamma}) \frac{1}{\gamma^2} & \gamma > \gamma_0 \\ 1 & 0 < \gamma \leq \gamma_0 \end{cases} \quad (5)$$

where A and γ_0 are constant factors given by

$$A = \frac{3(1 - \cos(\alpha)) + \frac{12}{\pi} \sin(\alpha)}{1 - \frac{1}{2}(\sin(\alpha) \cos^2(\alpha) + (1 - \sin(\alpha))^2(2 + \sin(\alpha)))} \quad (6a)$$

$$\gamma_0 = \sqrt{3} \sin(\min\{\text{atan}(\frac{1}{\sqrt{2}}) + \alpha, \frac{\pi}{2}\}). \quad (6b)$$

Proof. See Supplementary I [55]. ■

According to Theorem 1, the relaxed factor γ can be determined by an expected accuracy value $\omega \in (0, 1]$ using the trigonometric form of Cardano's Formula [56], given as:

$$\gamma = \begin{cases} 2\sqrt{\frac{A}{3\omega}} \cos(\beta) & 0 < \omega < 1 \\ 1 & \omega = 1 \end{cases} \quad (7)$$

where β is given as

$$\beta = \frac{1}{3} \arccos\left(-\frac{A\gamma_0 - \gamma_0^3}{2\omega} \left(\frac{A}{3\omega}\right)^{\frac{3}{2}}\right) \quad (8)$$

It is worth noting that for $\alpha \in [0, \pi/2]$, γ_0 always satisfies $\gamma_0 \geq 1$. Therefore, setting $\gamma = 1$ always results in $f(\gamma) = 1$, regardless of the value of α . Besides, it is noted the accuracy function $f(\gamma)$ attains a value of 1 for $\gamma \in (0, \gamma_0]$, as shown in equation (5). Therefore, in equation (7), we select $\gamma = 1$ as the solution for $\omega = 1$.

The accuracy function $f(\gamma)$ is validated through a series of simulation experiments involving various LiDAR detection ranges and map resolutions. Specifically, given a detection range R and map resolution d , a batch of point clouds covering the free space within the detection range is generated on a sphere of radius R . To compute the volume V_I , we count the number of cells in the map with resolution d that are identified as free by projecting the map to the original depth image \mathcal{M}_I and multiplying the result by the unit volume d^3 . Similarly, we obtain the volume V by projecting the map onto the relaxed depth image \mathcal{M} . The accuracy function value $f(\gamma)$ is calculated using the definition in (4). In the simulations, we set the LiDAR's FoV Φ to $[-15^\circ, 15^\circ]$.

The results of the simulation experiments are presented in Fig. 8. The theoretical model fits the experimental curves well.

However, slight errors arise when γ approaches γ_0 due to discretization in counting the number of cells to evaluate the volume of covered regions. These errors are evident in the curves of low resolution and short detection range (e.g., in the setup of $R = 10\text{m}$ and $d = 0.5\text{m}$, which is the green solid line in Fig. 8).

V. OCCUPANCY MAPPING

In this section, we describe how to update and query the occupancy states using the map structure in D-Map.

A. Occupancy Map Structure

With an aim to optimize the balance between computation and memory efficiency, D-Map leverages a hashing grid map to maintain occupied space and an octree to maintain unknown space.

1) *Hashing Grid Map*: As there are typically fewer occupied regions in the environment than free and unknown ones, we maintain the occupied space of the environment in a hashing grid map using the voxel hashing technique [42], which allows efficient update and query operations in a time complexity of $\mathcal{O}(1)$. The hash key value for a given point $p = [x, y, z]$ is computed by a hashing function Hash , defined as follows:

$$\begin{aligned} \text{Hash}(p) &= (P^2 n_z + P n_y + n_x) \bmod Q \\ n_x &= \lfloor \frac{x}{d} \rfloor, \quad n_y = \lfloor \frac{y}{d} \rfloor, \quad n_z = \lfloor \frac{z}{d} \rfloor, \end{aligned} \quad (9)$$

where d represents the resolution of the hashing grid map. P and Q are large prime numbers while Q also serves as the size of the hashing table. \bmod is the modulus calculation between two integers. The value of P and Q are carefully selected to minimize conflict probability [57], with P and Q set to 116101 and 201326611 in our work, respectively.

2) *Octree*: The tree-based map structure is a commonly used approach for occupancy mapping due to its high memory efficiency. Among various tree-based data structures, the octree [58] stands out as it outperforms other spatial data structures for dynamic updates [59]. Additionally, the spatial division on the octree naturally allows for determining the occupancy states of large cells during tree updates. Therefore, we utilize an octree to organize the unknown space.

In D-Map, a node on the octree contains the following elements:

- The point array $\text{ChildNodes}[8]$ contains the address to its eight child nodes. The point array is empty if the node is a leaf node.
- The center C of the cell represented by the node.
- The size L of the cell represented by the node.

3) *Initialization*: To initialize the mapping procedure in D-Map, the occupancy states of the environment are considered to be entirely unknown. In the implementation, given an initial bounding box C_{bbox} of the interested regions, the root node of the octree is initialized to represent an unknown cube C_{root} whose the center C is allocated at the center of C_{bbox} . The size L of the root node is assigned as the longest side length of the bounding box C_{bbox} , and the point array ChildNodes for

Algorithm 2: Occupancy Update

Params: Map Resolution d , Initial Grid Size E ,
LiDAR Detection Range R ,
LiDAR FoV Φ and Θ

Input : Sensor Pose T , Point Cloud \mathcal{P}

1 Algorithm Start

2 $\text{UpdateGridMap}(\mathcal{P})$;
3 $\mathcal{M} \leftarrow \text{Rasterization}(\mathcal{P})$;
4 $\mathcal{T} \leftarrow \text{Build2DSegTree}(\mathcal{M})$;
5 $\mathcal{V} \leftarrow \text{SensingArea}(T, R, \Phi, \Theta)$;
6 $\text{UpdateOctree}(\text{RootNode}, T, \mathcal{V}, \mathcal{T})$;

7 Algorithm End

8 Function $\text{UpdateOctree}(\text{Node}, T, \mathcal{V}, \mathcal{T})$

9 $L, C \leftarrow \text{GetCenterAndSize}(\text{Node})$;
10 **if** $\text{!Intersected}(T, C, \mathcal{V})$ **then return**;
11 **if** $L > E$ **then**
12 $S \leftarrow \text{Undetermined}$;
13 **else**
14 $S \leftarrow \text{DetermineOccupancy}(T, C, L, \mathcal{T})$;
15 **end**

16 **switch** S **do**

17 **case** Unknown **do**
18 **return**;
19 **end**

20 **case** Known **do**
21 $\text{DeleteNode}(\text{Node})$;
22 **end**

23 **case** Undetermined **do**
24 **if** $L \leq d$ **then**
25 $\text{DeleteNode}(\text{Node})$
26 **else**
27 **if** $\text{IsLeaf}(\text{Node})$ **then** $\text{Split}(\text{Node})$;
28 **foreach** Child in ChildNodes **do**
29 $\text{UpdateOctree}(\text{Child}, T, \mathcal{V})$;
30 **end**
31 **end**
32 **end**
33 **end**

34 End Function

child nodes is initialized as empty. Besides, the hash table in the hashing grid map is initialized as an empty table. Notably, the initial bounding box does not restrict the mapping area, as both the octree and the hashing grid map allow for the growth of the mapping space on-the-fly [42, 50].

B. Occupancy Update

The occupancy updates in D-Map involve updating the hashing grid map and the octree, as described in Alg. 2. The point clouds are directly inserted into the hashing grid map using the hash function in (9) (Line 2) and rasterized into a depth image \mathcal{M} (Line 3). A 2-D segment tree is constructed from the depth image \mathcal{M} for the subsequent occupancy state determination as described in Section IV (Line 4). The sensing area \mathcal{V} for cell extraction is obtained from the sensor pose T , LiDAR detection range R and LiDAR FoV $\Phi \times \Theta$ (Line

5). D-Map employs an on-tree update strategy to update the occupancy states of octree nodes within the LiDAR sensing area \mathcal{V} using a function named `UpdateOctree` (Line 6).

The function `UpdateOctree` is described in Lines 9~33 and explained as follows. Starting from the root node of the octree, the corresponding cell center C and its size L are obtained (Line 9). The algorithm checks if the corresponding cell intersects with LiDAR sensing area \mathcal{V} , terminating further updates if there is no intersection (Line 10). To avoid meaningless occupancy state determination on a too-large cell that always returns undetermined, we directly split those cells larger than an initial cell size E for efficiency (Line 11~13). Otherwise, the occupancy state S is determined by Alg. 1 (Line 14). The operation on the tree node is decided based on the occupancy state S : If the cell is determined as unknown, the node is kept on the tree, and no further operations are required (Line 17~19). The node is removed from the tree if the cell is known (i.e., free or occupied) (Line 20~22). If the occupancy state of the node cannot be determined, further updates are required (Line 23~32). A special condition is considered when the cell size L achieves the map resolution d (Line 24~26). In this case, at least one pixel with a depth no smaller than the cell's minimum depth BoxMin exists (otherwise, it is determined as unknown for $d\text{Max} < \text{BoxMin}$, see Alg. 1). As a result, the cell can be determined as known since it must have been either hit or traversed by the corresponding point cloud on that pixel. Otherwise, if the cell size is greater than the map resolution, we split the node, if not have done so (Line 27), and performed updates on its eight children (Line 28~30).

The recursive function `UpdateOctree` operates in a manner that visits cells from the largest to the smallest size. This approach allows the determination of occupancy states directly on large cells (if they are unknown or free), thus avoiding unnecessary updates on smaller cells. As a result, occupancy updates on the octree are performed concurrently with the tree traversal, resulting in a so-called “on-tree” update strategy. This strategy is distinct from the classical pipeline [20, 52], which updates cells at the constant map resolution d after ray-casting. Furthermore, D-Map continuously removes known cells from the environment's unknown space, making it a decremental mapping method. This property enhances the efficiency as the mapping process progresses.

C. Occupancy State Query

Since two data structures are used in D-Map, two steps are performed to obtain the occupancy state of a cell at map resolution. Firstly, the cell is determined as unknown if its corresponding node exists on the octree. Otherwise, the region is determined as known. Subsequently, the hashing grid map is queried to determine whether it is occupied. If not, the region is determined as free.

VI. TIME COMPLEXITY ANALYSIS

In this section, we analyze the time complexity of updating and querying the occupancy states on D-Map. To facilitate further analysis of the benchmark experiments in Section VII, we additionally provide a time complexity analysis of the classical




Fig. 9. A 2-D illustration of the worst-case occupancy updates for (a) D-Map and (b) ray-casting-based methods. In the worst-case scenario for D-Map, two neighboring points appear at the minimum and maximum distance to the sensor origin, resulting in a serrated-shaped point cloud. In contrast, the worst-case scenario for ray-casting-based methods results in a spherical-shaped point cloud located at the maximum distance to the sensor origin.

occupancy update pipeline, which involves ray-casting prior to map updates.

A. Occupancy State Update

The occupancy state updates in D-Map are composed of two parts: updating occupied cells on the hashing grid map and updating unknown cells on the octree. The time complexity for occupancy updates is mainly due to the on-tree update in the octree data structure, as the time of maintaining occupied points in a hashing grid map is constant at $\mathcal{O}(1)$. The on-tree update strategy involves searching for nodes within the LiDAR sensing area on the octree and determining their occupancy states. To analyze the time complexity, we first consider the number of nodes to be visited on the octree, as presented in the following lemma:

Lemma 1. *Suppose the maximum side length of environment bounding box C_{bbx} on the octree is D , and map resolution is d . The number of nodes to be visited S_{DMap} on the octree for the on-tree update strategy in D-Map is bounded as:*

$$S_{\text{DMap}} = \mathcal{O}\left(n \frac{R}{d} \log\left(\frac{D}{d}\right)\right) \quad (10)$$

where n is the number of points projected to the depth image, and R is the detection range of the LiDAR sensor.

Proof. See Supplementary II-A [55]. ■

Then we derive the time complexity of occupancy state determination (see Alg. 1) that is frequently used in the update process.

Lemma 2. *The time complexity of occupancy state determination is $\mathcal{O}(\log(\Phi_I/\psi_I) \log(\Theta_I/\psi_I))$, where Φ_I and Θ_I are the FoV angles of the depth image and ψ_I is the depth image resolution.*

Proof. See Supplementary II-B [55]. ■

Finally, the time complexity for occupancy updates on D-Map is given as follows:

Theorem 2. *The time complexity of occupancy updates in D-Map is $\mathcal{O}\left(\frac{nR}{d} \log\left(\frac{D}{d}\right)\right)$, where n is the number of points*

projected to the depth image, d is the map resolution, and R is the LiDAR detection range.

Proof. See Supplementary II-C [55]. \blacksquare

We provide the time complexity of the occupancy updates on a grid-based map and an octree-based map as follows:

Theorem 3. *The time complexity to update the occupancy states on a grid-based map is $\mathcal{O}(\frac{nR}{d})$ and on an octree-based map is $\mathcal{O}(\frac{nR}{d} \log(\frac{D}{d}))$.*

Proof. See Supplementary II-D [55]. \blacksquare

Remark 1. Although D-Map and the octree-based map have the same time complexity of $\mathcal{O}(\frac{nR}{d} \log(\frac{D}{d}))$, their sources of time complexity are entirely different. Since D-Map updates the occupancy states concurrently with the tree traversal, as discussed in Section V-B, the time complexity of D-Map depends solely on the number of visited nodes $S_{D\text{Map}}$ during the on-tree update process, which contributes to $\mathcal{O}(\frac{nR}{d} \log(\frac{D}{d}))$ directly. In contrast, occupancy updates on an octree-based map consist of two consecutive procedures: ray-casting, which contributes to $\mathcal{O}(\frac{nR}{d})$, and map updates, which contribute to $\mathcal{O}(\log(D/d))$, resulting in the same total time complexity of $\mathcal{O}(\frac{nR}{d} \log(\frac{D}{d}))$.

Remark 2. The time complexity of D-Map is calculated without considering the removal of known cells, which could reduce the number of visited nodes $S_{D\text{Map}}$ substantially. In contrast, existing ray-casting-based methods (either grid-based map or octree-based map) do not have such decremental properties.

Remark 3. We derive the time complexity of occupancy updates for our D-Map and ray-casting-based methods (i.e., the grid-based and octree-based maps) based on their respective worst-case scenario, as illustrated in Fig. 9. However, the occurring chances of their worst-case scenarios are different. Generally speaking, the worst-case scenario for ray-casting-based methods is more likely to occur in the real world (e.g., large open space) than for our D-Map.

B. Occupancy State Query

The time complexity of the occupancy state query in D-Map is provided as follows:

Theorem 4. *The time complexity of querying the occupancy state of a cell in D-Map is $\mathcal{O}(\log(D/d))$.*

Proof. See Supplementary II-E [55]. \blacksquare

VII. BENCHMARK RESULTS

In this section, exhaustive benchmark experiments are conducted on various datasets with different LiDARs to evaluate the computational efficiency, accuracy, and memory consumption against other state-of-the-art occupancy mapping approaches. Besides the time performance evaluation, we conduct an ablation study based on the time complexity to provide deeper insight into the efficiency of our approach.

TABLE I
DETAILS OF DATASETS FOR BENCHMARK EXPERIMENT

	Mapped Area (m ³)	Number of Scans	Average Point Number (per scan)
<i>FR_079</i>	$48 \times 37 \times 5$	66	89445.9
<i>Freiburg</i>	$293 \times 168 \times 29$	81	247817.1
<i>Workshop</i>	$40 \times 40 \times 8$	7340	23540.8
<i>MainBuilding</i>	$636 \times 251 \times 26$	1402	22124.9
<i>Kitti_04</i>	$551 \times 160 \times 37$	271	125265.2
<i>Kitti_06</i>	$616 \times 181 \times 47$	1101	122189.4
<i>Kitti_07</i>	$362 \times 347 \times 40$	1101	121225.5

A. Datasets

The experiments are conducted on three open datasets and one private dataset. The first public dataset is *Kitti* dataset which captured depth measurements by a Velodyne HDL-64E rotating 3D laser scanner at 10 Hz [60]. Considering the tremendous memory consumption associated with grid-based maps that will be utilized in our benchmark experiments, we chose sequences *Kitti_04*, *Kitti_06*, and *Kitti_07* to conduct the benchmark evaluation. The second public dataset is an outdoor sequence *MainBuilding* provided in FAST-LIO [61], which used a semi-solid state 3D LiDAR sensor Livox Avia to collect data at 10 Hz. The third public dataset, provided in the work Octomap [20], includes an indoor sequence *FR-079* and an outdoor sequence *Freiburg*. In addition, we evaluate the benefit of decremental property on a private dataset *Workshop*, where a cluttered indoor environment is completely scanned manually using 3D LiDAR Livox Avia at 10 Hz. The reconstruction result of *Workshop* dataset is available in [62]. It is noticed that the odometry estimation for occupancy mapping in sequence *Workshop* and *MainBuilding* was acquired by our LiDAR-inertial odometry framework FAST-LIO2 [59]. Table I provides further details on the aforementioned datasets.

B. Evaluation Setup

We compared D-Map with state-of-the-art occupancy mapping methods commonly used in robotics applications, including a grid-based method (Grid Map updated by 3DDDA [63]), a tree-based method (Octomap [20]), and the extended versions of Grid Map and Octomap using super rays and culling regions [52] (denoted with “SR&CR(Grid)” and “SR&CR(Octo)”). These occupancy mapping methods were selected for their efficient updates and widespread usage in the field. For the Octomap, SR&CR(Grid), and SR&CR(Octo), we used their open-source implementations available on GitHub repositories^{1,2}. For the grid map, we modify the open-source SR&CR(Grid) by disabling the super rays and culling regions and termed it as “GridMap”. It is noticed that all grid-based maps employed voxel hashing to achieve better memory efficiency. The benchmark experiments are conducted at various map resolutions, ranging from a rough resolution of 1.0 m to a finer resolution of 1 cm for indoor sequences. Due to system memory limitations, the minimum map resolution for outdoor sequences is set to 5 cm.

¹<https://github.com/OctoMap/octomap>

²<https://github.com/PinocchioYS/SuperRay>

The computation platform for evaluation is an Intel NUC computer with CPU Intel i7-10710U and 64 GB RAM. To facilitate evaluation on high-resolution maps in large-scale environments, a 64 GB swap space on Solid State Drive (SSD) is allocated to account for extensive memory usage.

D-Map uses a fixed set of parameters for all benchmark experiments. The completeness threshold ε in Alg. 1 is set to 0.8, while the initial cell size E in Alg. 2 is set to 5.0 m. The depth image resolution is determined using (2) without relaxation (i.e., $\gamma = 1$, which is not greater than γ_0 and leads to $f(\gamma) = 1$, as explained in Section IV-D). The LiDAR detection range R and LiDAR angular resolution ψ_{lidar} are acquired from LiDARs' manual sheets, except for *Workshop* indoor sequence whose detection range is assigned as 60 m in consideration of the indoor environment.

Given the high accuracy and low false alarm rate of LiDAR sensors, the space that is hit by a LiDAR pulse can be reliably considered occupied, and that passed by the laser ray can be considered free. We use the following parameter setup for other occupancy mapping approaches in the benchmark experiment to ensure consistency with this assumption that our D-Map has utilized. We set the probabilities for hit and miss by a ray (i.e., for occupied and free space) to $P_{\text{occ}} = 0.9999$ and $P_{\text{free}} = 0.4999$, respectively, and the minimum and maximum clamping probabilities to $P_{\text{min}} = 0.499$ and $P_{\text{max}} = 0.9999$. Moreover, we set the initial occupancy probabilities of all cells to 0.5, which represents unknown cells without any update (or observation) from LiDAR points, and we regard the cells with occupancy probability values smaller than 0.5 (due to subsequent LiDAR points update) as free and those with occupancy probability values larger than 0.5 as occupied. Finally, the methods SR&CR(Grid) and SR&CR(Octo) both require a threshold parameter k for the minimum number to generate super rays, which is set to the default value $k = 20$ as in [52].

C. Efficiency Evaluation and Analysis

1) *Benchmark Results*: We summarize the average update time of each occupancy mapping approach and report the results in Table II. The update time for SR&CR(Grid) and SR&CR(Octo) includes the time required for preprocessing, which involves generating super rays and constructing culling regions, as well as the time for occupancy updates involving ray-casting and map updates. In contrast, for GridMap and Octomap, only the time for ray-casting and map updates is considered, as no preprocessing is required. As for our D-Map, we count the preprocessing time for depth image rasterization and 2-D segment tree construction as well as the occupancy update time for updating the octree and hashing grid map. As presented in Table II, our D-Map consistently outperforms the other mapping approaches across various map resolutions in three sequences of the *Kitti* dataset. Particularly, at a map resolution of 5 cm in *Kitti_06*, D-Map achieves a remarkable speedup of 7.66 times and 4.78 times faster than GridMap and Octomap, respectively. However, in the *MainBuilding* datasets, GridMap shows the highest efficiency at different map resolutions, followed by our D-Map and SR&CR(Grid).




Fig. 10. Update time in *Workshop* indoor sequence at a high map resolution of 1 cm.

The lower efficiency of D-Map is caused by the sparse clouds in *MainBuilding* sequence, which only has an average number of 22 k per scan, while *Kitti* datasets have over 120 k points per scan on average. This finding suggests that our method is less efficient in updating sparse point clouds since a sparse depth image rasterized from the sparse point clouds causes redundant occupancy state queries on those unobserved regions during the updating process. In contrast, ray-casting-based mapping approaches exhibit better performance due to their precise traversal of observed regions. However, this disadvantage of D-Map could be compensated by the decremental property as shown in the *Workshop* indoor sequence, where the scanning process frequently visited the previously explored areas with an aim to map the entire indoor environment exhaustively. This scanning process, which is often used in mapping applications, enables our D-Map to continuously remove known space from the octree map, achieving higher efficiency than the other methods. Fig. 10 presents the update time consumption of different occupancy approaches at a 1 cm map resolution, clearly illustrating the decreasing update time of D-Map as the environment is explored.

D-Map demonstrates superior performance compared to other approaches at high map resolutions ($d \leq 0.25$ m) in the indoor sequence *FR_079*, as well as at various resolutions in the outdoor sequence *Freiburg*. However, at low resolutions ($d \geq 0.5$ m) in *FR_079* sequence, the performance is limited by the preprocessing time required for depth image rasterization (approximately 5.7 ms) and the processing time to insert occupied points into the hashing grid map (about 4.2 ms).

Across all benchmark experiments conducted in this study, we have observed that SR&CR(Grid) and SR&CR(Octo) exhibit low efficiency at high resolution compared to GridMap and Octomap, respectively, without super rays or culling regions. This phenomenon has been previously noted and discussed in [52], where it is attributed to the low grouping ratio of super rays and the need for the timely rebuilding of culling regions at high resolution. Additionally, we have also observed anomalous time performances when comparing the grid-based method (i.e., GridMap and SR&CR(Grid)) against the tree-based methods (i.e., Octomap and SR&CR(Octo)). Normally, grid-based maps are expected to offer superior update efficiency compared to tree-based maps. However, we

TABLE II
COMPARISON OF UPDATE TIME (MS) AT DIFFERENT RESOLUTION

Resolution (m)	Kitti_04					Kitti_06					Kitti_07				
	1.0	0.5	0.25	0.1	0.05	1.0	0.5	0.25	0.1	0.05	1.0	0.5	0.25	0.1	0.05
GridMap	47.32	89.76	190.78	720.33	3636.54	50.22	97.70	203.56	1110.29	10505.13	37.04	68.64	140.28	536.51	3431.81
Octomap	398.52	637.52	964.52	2459.13	7623.91	437.43	664.70	1076.62	2629.29	6558.01	299.72	469.21	775.76	1777.65	3877.65
SR&CR(Grid)	43.08	129.51	298.76	1398.72	7360.95	48.86	116.96	322.30	2422.06	17977.50	30.17	72.91	180.42	995.12	6290.95
SR&CR(Octo)	44.47	125.31	389.37	2330.39	9287.51	56.49	161.41	554.35	2974.88	12361.56	30.05	85.18	307.54	1634.40	6153.84
Ours	22.67	39.05	101.14	470.67	1905.32	22.71	36.08	85.91	367.08	1371.69	19.12	27.27	59.33	209.46	779.32
Workshop															
Resolution (m)	1.0	0.5	0.25	0.1	0.05	0.025	0.01	1.0	0.5	0.25	0.1	0.05	1.0	0.5	0.05
GridMap	4.94	8.42	15.97	50.17	127.43	346.68	13702.461	9.63	19.91	45.88	162.65	640.83			
Octomap	42.21	59.62	92.76	175.17	297.93	589.07	7568.64	64.82	106.43	199.05	529.66	1333.22			
SR&CR(Grid)	2.64	7.35	18.65	72.68	244.95	1514.41	39624.96	10.86	27.65	72.61	315.12	1384.12			
SR&CR(Octo)	2.64	9.29	29.96	162.54	717.84	3352.30	50174.95	12.61	44.12	156.80	908.08	3783.49			
Ours	3.02	4.45	9.26	20.10	42.12	88.16	915.20	11.16	31.56	112.52	332.83	991.66			
FR_079															
Resolution (m)	1.0	0.5	0.25	0.1	0.05	0.025	0.01	1.0	0.5	0.25	0.1	0.05	1.0	0.5	0.05
GridMap	10.14	15.47	26.47	63.77	152.04	438.84	4082.14	61.41	109.99	230.90	1021.84	6005.56			
Octomap	3.89	7.90	17.45	66.48	347.57	1514.97	7568.64	356.43	548.27	990.92	3357.36	10637.87			
SR&CR(Grid)	3.76	7.91	16.34	53.34	237.25	816.58	4805.09	40.72	110.27	285.78	1522.21	7911.33			
SR&CR(Octo)	3.87	7.70	18.28	68.21	375.14	1613.46	8402.84	43.68	112.59	354.07	2186.24	8709.99			
Ours	10.09	11.29	15.11	29.01	65.22	205.06	756.33	31.40	40.48	82.55	678.09	3599.96			

have observed that grid-based methods are slower than tree-based methods at high resolution ($d = 5$ cm) in the sequences *Kitti_06* and *Kitti_07*. The inefficiency of grid-based methods can be attributed to their reliance on swap space on the SSD during the update process since their memory usage exceeds the capacity of RAM. As memory access on SSD is significantly slower than RAM, this reliance on swap space can lead to the observed slower update performance.

2) *Efficiency Analysis*: We present a comprehensive analysis of our superior efficiency performance based on the time complexity analysis in Section VI. As demonstrated in Section VI, the efficiency of D-Map is solely determined by the number of visited cells $S_{D\text{Map}}$ due to its “on-tree” update strategy. However, GridMap, Octomap, SR&CR(Grid), and SR&CR(Octo), which follow the classical occupancy update pipeline using ray-casting, have an efficiency that depends on two factors: the number of cells S_{rc} traversed by rays and the time complexity for map updates, which is $\mathcal{O}(1)$ for grid-based methods and $\mathcal{O}(\log(D/d))$ for octree-based methods. Therefore, we conduct a comparison between the number of visited cells $S_{D\text{Map}}$ in D-Map and the number of updated cells S_{rc} in the ray-casting-based methods, followed by a comparison of the corresponding occupancy update time (without counting the preprocessing time). Specifically, we opt to compare D-Map with SR&CR(Grid) and SR&CR(Octo), which incorporate super rays and culling region techniques to reduce the number of cells. The comparison is carried out on *FR_079* and *Freiburg* sequences, with the results presented in Fig. 11.

In comparison between $S_{D\text{Map}}$ and S_{rc} , we conduct an ablation study to investigate the performance of all methods (i.e., SR&CR(Grid), SR&CR(Octo), and D-Map) without any removal of cells. For our D-Map, we disable the removal of known cells from the octree. For SR&CR(Grid) and SR&CR(Octo), we disable the culling region technique and

denote these modified methods as “SR(Grid)” and “SR(Octo)”, respectively. As shown in Fig. 11, SR(Grid) exhibits the highest count of cells across all resolutions, followed by SR(Octo) and our D-Map without removal. The difference in the number of updated cells between SR(Grid) and SR(Octo) is attributed to the batching-based method in SR(Octo) that batches the uniform cells traversed by multiple rays to reduce the cells. The difference between the two methods becomes smaller at higher resolution as there are fewer uniform cells traversed by multiple rays. It is worth noticing that, although batching is essential to tree-based methods to account for the cost of tree updates, it is unnecessary for grid-based maps, which have an efficient update in simple time complexity of $\mathcal{O}(1)$. As analyzed in Section VI, D-Map has a theoretical number of visited nodes $S_{D\text{Map}} = \mathcal{O}(\frac{nR}{d} \log(\frac{D}{d}))$ in its worst case, which is larger than the theoretical number of traversed cells $S_{rc} = \mathcal{O}(\frac{nR}{d})$ in the worst case of ray-casting-based methods. However, in this ablation study, D-Map without removal visits a comparable number of nodes to SR(Octo) because the worst case of D-Map rarely occurs. Specifically, D-Map without removal processes fewer nodes than SR(Octo) at low resolutions (i.e., $d \geq 0.25$ m) in indoor sequence *FR_079* and at all resolutions except 5 cm in outdoor sequence *Freiburg*. When approaching high resolutions (e.g., $d < 0.25$ m in indoor sequence and $d = 0.05$ m in outdoor sequence), the number of visited nodes $S_{D\text{Map}}$ on D-Map is closer to its worst case, which has a higher increasing rate on a logarithmic scale than S_{rc} .

The removal of known cells and the culling region technique are then enabled in corresponding methods for further investigation. D-Map demonstrates the most significant reduction in the number of cells among all approaches when the removal technique is enabled, while the culling region technique only slightly lowers the number of updated cells in SR&CR(Grid)




Fig. 11. The number of cells to be updated in sequences *FR_079* and *Freiburg*.




Fig. 12. The update time in sequences *FR_079* and *Freiburg*.

and SR&CR(Octo).

Finally, we compare the time consumption for occupancy updates in both sequences, as shown in Fig. 12. When disabling the removal of known cells in D-Map and the culling region technique, the results reveal that D-Map without removal outperforms the others across all map resolutions in sequence *Freiburg*, as well as at low resolutions (i.e., $d > 5\text{cm}$) in sequence *FR_079*. At high resolutions (i.e., $d \leq 0.5\text{m}$) in sequence *FR_079*, SR(Grid) achieves better performance than D-Map without removal due to a smaller number of cells to be updated. However, despite having a slightly larger number of visited nodes, D-Map without removal maintains its leading position over SR(Octo), which is hindered by the timely updates on the tree structure. When enabling the removal of known cells, D-Map demonstrates the highest efficiency across different map resolutions, owing to the fewest cells requiring to be updated.

D. Accuracy and Memory Evaluation

1) *Accuracy Benchmark*: We conduct accuracy benchmark experiments using the mapping results from Octomap [20] as ground truth. Specifically, we calculate the accuracy of D-Map by counting the cells with correct occupancy states over

the total number of cells inside the mapping area and with the corresponding states. The benchmark results of unknown space and free space are presented in Table III. It is noted that the accuracy of occupied space is very close to 100 % in all experiments and thus not presented in the table. This is due to our high confidence in occupied space based on the assumption of high accuracy and low false alarm rate of LiDAR sensors.

The overall accuracy performance of D-Map shows a slight accuracy loss compared to Octomap, as reported in Table III. The inaccuracy mainly arises from the occupancy state determination module on the depth image. At low resolutions (i.e., $d \geq 0.5\text{m}$), the accuracy of unknown space is relatively lower due to the more severe discretization error introduced in (3) when determining the projected area on a lower-resolution depth image. The results also show a decreasing accuracy of free space when the map resolution grows smaller. This decreasing accuracy is introduced by projecting a cell by its circumsphere on the depth image. When a ray crosses a large cell only at its corner but not intersects with its circumsphere, this large cell is incorrectly decided as not intersecting with any rays without further splitting and query. Thus, the small cells at the corner space are falsely determined as unknown. With a higher map resolution, more small free

TABLE III
MAPPING ACCURACY AT DIFFERENT RESOLUTION

Resolution (m)	Unknown Space							Free Space						
	1.0	0.5	0.25	0.1	0.05	0.025	0.01	1.0	0.5	0.25	0.1	0.05	0.025	0.01
<i>Kitti_04</i>	97.33%	98.29%	99.67%	99.77%	99.46%	-	-	98.54%	98.49%	96.73%	94.61%	93.34%	-	-
<i>Kitti_06</i>	97.77%	98.62%	99.87%	99.84%	99.24%	-	-	98.80%	98.77%	97.21%	94.99%	94.53%	-	-
<i>Kitti_07</i>	97.68%	99.50%	99.64%	99.87%	99.40%	-	-	98.40%	98.36%	96.91%	94.93%	94.24%	-	-
<i>FR_079</i>	83.53%	96.46%	97.86%	98.96%	99.66%	99.17%	95.39%	95.06%	95.72%	96.55%	97.88%	95.14%	91.68%	93.56%
<i>Freiburg</i>	97.79%	98.93%	99.39%	99.91%	99.96%	-	-	98.45%	98.61%	98.80%	92.60%	83.28%	-	-
<i>Workshop</i>	98.68%	99.19%	99.61%	99.82%	99.87%	99.95%	99.99%	95.11%	96.74%	98.08%	98.80%	99.01%	98.90%	99.00%
<i>MainBuilding</i>	99.51%	99.71%	99.77%	99.58%	99.80%	-	-	97.65%	97.56%	97.08%	95.10%	87.76%	-	-

cells in the corner space are determined as unknown, leading to a decreasing accuracy. To improve the accuracy of free space for such cases, it is recommended to use a smaller initial cell size E to avoid missing the cell corner. However, a trade-off must be considered as a smaller E might miss the opportunity to update on large cells, leading to slower occupancy updates.

Notably, since we utilize Octomap as ground truth, the discretization error in Octomap also contributes to the accuracy loss in our benchmark experiment. Besides, despite the quantization error, Octomap has been widely used in real-world applications. Therefore, the accuracy achieved in our experiment, as reported in Table III, is adequate for various practical scenarios, as demonstrated in Section VIII.

2) *Memory Consumption*: We evaluated the memory consumption of each mapping approach and reported the results in Table IV. Given that SR&CR(Grid) and SR&CR(Octo) utilize the same map structure as GridMap and Octomap with equivalent mapping results, we only compare the memory consumption of D-Map against GridMap and Octomap. As shown in Table IV, Octomap exhibits the lowest memory consumption in most experiments, while GridMap consumes the highest. The memory consumption of D-Map follows that of Octomap closely. The excess memory consumption of D-Map mainly arises from using hashing grid map for maintaining occupied space. When high-resolution maps are built in large-scale environments (e.g., 5 cm in *Kitti* datasets), our D-Map exhibits better memory efficiency than Octomap due to its decremental property as well as the proper combination of grid-based and tree-based structures for different occupancy states.

VIII. REAL-WORLD APPLICATIONS

We demonstrate the high efficiency of D-Map in two applications that require real-time high-resolution occupancy mapping on a high-resolution LiDAR.

A. Interactive Guidance for High-resolution Real-time 3D Mapping

With the recent emergence of 3D applications such as metaverse [64, 65], virtual and augmented reality [66], and physical simulators [62, 67], the demand for accurate and detailed 3D reconstructions of real-world environments has increased. The accuracy and completeness of such reconstructions depend heavily on the quality of the data collection process. To overcome this challenge, we have developed




Fig. 13. (a) Hardware setup of the handheld device for high-resolution 3D mapping, including an onboard computer (blue dashed block), a high-resolution LiDAR (orange dashed block), and a screen for online visualization (green dashed block). (b) A screenshot of the online visualization for interactively guiding the mapping process. The red cubes are the frontiers that users need to eliminate by scanning. The yellow arrow indicates the direction to the next suggested frontier for scanning. The depth measurements are accumulated and visualized by height value on the screen. (c) The first-person view of the user to conduct 3D mapping using our handheld device.

an interactive guidance system that leverages our D-Map to achieve high-resolution real-time 3D mapping. The system offers users information on the explored and unexplored areas, along with suggestions for the next mapping region. By utilizing this information, users can avoid rescanning the same areas repeatedly or skipping any unscanned areas, thereby improving their overall work efficiency. A video demonstrating the use of the system is available on YouTube: youtu.be/m5QQPbkYYnA?t=251.

1) *Experiment Setup*: A handheld device is designed to conduct high-resolution 3D mapping, as shown in Fig. 13(a). The handheld device is equipped with an Intel NUC onboard computer with an i9-8550U CPU and 64 GB memory, and an OS1-128 LiDAR integrated with an IMU. The OS1-128 LiDAR can output 2,621,440 high-precision point clouds per second with a maximum detection range of 120 m and a $360^\circ \times 45^\circ$ field of view. The frame rate of the OS1-128 LiDAR is set to 10 Hz.

We develop an interactive guidance system that consists of three modules: localization, mapping, and guidance. The localization module provides 6 DoF sensor pose estimation using our previous work FAST-LIO2 [59]. Additionally, the point cloud acquired at each frame is compensated for motion distortion in FAST-LIO2 and transformed from the LiDAR frame to the world frame. The mapping module leverages our D-Map to update the occupancy map in real-time using the estimated LiDAR pose and the currently registered point cloud.

TABLE IV
COMPARISON OF MEMORY CONSUMPTION (MB) AT DIFFERENT RESOLUTION

Resolution (m)	Kitti_04					Kitti_06					Kitti_07				
	1.0	0.5	0.25	0.1	0.05	1.0	0.5	0.25	0.1	0.05	1.0	0.5	0.25	0.1	0.05
Grid Map	13.00	97.11	751.00	6964.45	33012.05	25.32	194.32	1532.03	14586.55	110589.78	22.49	105.35	894.36	7616.39	56745.76
Octomap	6.49	32.79	181.97	2212.85	14594.72	10.87	57.29	340.08	4587.48	31780.61	8.66	44.24	247.52	2858.18	18505.91
Ours	18.19	75.05	282.98	2377.75	11990.52	26.36	135.79	620.75	4726.50	24082.73	20.80	89.68	402.93	3417.51	15315.92
Workshop															
Resolution (m)	1.0	0.5	0.25	0.1	0.05	0.025	0.01	1.0	0.5	0.25	0.1	0.05	MainBuilding		
Grid Map	0.18	1.32	6.50	102.35	826.43	6719.52	109724.66	22.31	106.70	817.38	7328.23	31258.35			
Octomap	0.11	0.52	2.81	27.39	149.20	803.47	10619.32	5.83	33.28	261.84	3359.45	17566.79			
Ours	0.31	1.80	8.53	72.56	338.94	1516.09	11453.62	7.90	33.45	187.56	3038.41	26168.12			
FR_079															
Resolution (m)	1.0	0.5	0.25	0.1	0.05	0.025	0.01	1.0	0.5	0.25	0.1	0.05	Freibrug		
Grid Map	0.16	0.71	5.36	49.99	391.92	1921.08	26432.80	11.60	56.30	432.26	6594.09	30490.11			
Octomap	0.08	0.34	1.64	14.32	99.17	797.95	9991.86	4.23	21.61	132.60	2012.67	14453.00			
Ours	0.25	1.02	4.38	26.89	120.96	600.16	3289.48	11.00	45.67	189.19	1721.62	14026.21			




Fig. 14. Comparison of the update time among our D-Map, SR&CR(Grid) and SR&CR(Octo). The update time of D-Map is acquired online in the interactive guidance system during the mapping process. The update time of the SR&CR(Grid) and SR&CR(Octo) is obtained by processing the recorded point clouds offline on the same computation platform.

TABLE V
OCCUPANCY MAPPING PERFORMANCE IN INTERACTIVE GUIDANCE SYSTEM

	Average Update Time	Processed Scan	Completion Rate ¹
D-Map	36.50 ms	9980	99.46%
SR&CR(Grid)	348.58 ms	2587	25.78%
SR&CR(Octo)	721.79 ms	1252	12.48%

¹ The completion rate is calculated as the ratio of the processed scan over the total scan. In this experiment, LiDAR generated 10034 scans of point clouds in total.

Finally, the guidance module visualizes the frontiers, which correspond to the unknown space adjacent to the free space in the occupancy map, and suggests a direction to the user for complete scanning, as shown in Fig. 13(b).

The experiment aims to reconstruct an area measuring 43 m × 19 m × 9 m located in Main Building, University of Hong Kong. We use a map resolution of 10 cm for occupancy mapping. The relax factor for D-Map is set to $\gamma = 1$ for full accuracy. In addition, we set the completeness threshold ε in Alg. 1 to 0.8, and the initial cell size E to 1.6 m.

2) *Results*: As Fig. 15(a) illustrates, the building reconstruction was completed with high fidelity. The mapping area was thoroughly scanned, except for an inaccessible classroom

located in the middle of the map.

We conducted a performance evaluation of our D-Map against the super ray and culling region-based methods (i.e., SR&CR(Grid) and SR&CR(Octo)) on the recorded LiDAR data file that contains 10,034 scans of LiDAR points in total. As shown in Fig. 14 and Table V, D-Map achieved an average update time of only 36.50 ms to update the occupancy map, which is about 9.6 times faster than SR&CR(Grid) and 19.8 times faster than SR&CR(Octo). In addition, D-Map successfully processed the LiDAR data in real-time, except for the initial 54 s when the device was stationary. Consequently, D-Map processed 99.46 % of the total scan, generating a high-resolution occupancy map with high fidelity. In contrast, SR&CR(Grid) and SR&CR(Octo) failed to process in real-time, processing only 25.78 % and 12.48 % of the total scan, respectively. As a result, they missed a significant amount of environment information required for interactive guidance.

B. Autonomous UAV Exploration

Unmanned aerial vehicles (UAVs) are becoming increasingly popular for autonomous exploration and scanning of real-world environments due to their unrestricted flight view and accessibility to hard-to-reach locations, such as caves [68] and ancient remains [69]. However, the limited onboard computing power of UAVs presents a higher demand for efficient mapping modules compared to handheld devices. To address this demand, we have embedded our D-Map into a LiDAR-based UAV system, enabling it to autonomously explore complex scenes with higher-resolution mapping. This integration enables UAVs to achieve real-time mapping and guidance, thereby providing high-fidelity results even in challenging environments.

1) *Hardware System Setup*: The UAV hardware platform includes an OS1-128 LiDAR, a NORA+ flight controller, and an onboard computer NUC 12 Pro Kit (i7-1260P CPU, maximum 4.70 GHz, 12-core, 32 GB RAM), as shown in Fig. 1(c). The extrinsic parameters between the LiDAR and the IMU on the flight controller are calibrated by LI-Init [70]. Additionally, we




Fig. 15. The high-fidelity point cloud map reconstructed from data collected by a high-resolution LiDAR. (a) Main Building in the University of Hong Kong, collected by a handheld device. (b) A forest in Hong Kong, autonomously collected by a UAV platform.

installed an action camera (DJI action 2) on the UAV to provide first-person view (FPV) images for better visualization. The small size of the UAV (40 cm \times 40 cm \times 21 cm), combined with its large detection range and high scanning density, makes it well-suited for exploration and scanning tasks in complex scenes.

2) *Software System Implementation*: Multiple modules operate concurrently on the onboard computer to accomplish autonomous exploration tasks. The localization module employs FAST-LIO2 [59] to estimate the UAV's pose using data from the LiDAR and the IMU. This module generates undistorted point clouds, which are fed into the mapping module of the exploration planner. The mapping module utilizes D-Map for real-time occupancy mapping in place of the traditional ray-casting-based methods. The exploration planner, FUEL [17], reads the information in the mapping module and calculates efficient, collision-free trajectories for exploration. These trajectories are tracked using an on-manifold model predictive controller [71].

In this experiment, we used the same parameters for D-Map as in Section VIII-A. The maximum detection range of the LiDAR was set to 15 m to ease the computation load in the exploration planner.

3) *Results*: The real-flight autonomous exploration and scanning experiments were conducted in two complex scenarios: an abandoned fortress site built about 100 years ago (20.5 m \times 16 m \times 6 m) and a natural forest (17 m \times 35 m \times 3 m). The entire exploration process is available on youtu.be/m5QQPbkYYnA?t=78. We successfully carried out three flight experiments in each scenario. The UAV completed high-precision scanning in these complex environments, and the average flight time for autonomous exploration and scanning by the UAV was 123 s and 163 s, respectively. The acquired point clouds are shown in Fig. 1(a) and Fig. 15(b), exhibiting high precision and completeness.

We compared the running times of the mapping module employing D-Map and a uniform grid map utilizing ray-casting, which was the original method used in FUEL [17], on the same computing platform. As depicted in Fig. 16, D-Map substantially reduced the time required for the mapping module, thereby demonstrating its effectiveness and efficiency for real-time occupancy mapping on massive streaming point clouds of over two million per second.




Fig. 16. The comparison of processing time between a ray-casting-based grid map and our D-Map.

TABLE VI
COMPARISON OF UPDATE TIME (MS) AT DIFFERENT RESOLUTIONS ON A
LARGE-SCALE DATASET

Resolution [m]	ford_1				ford_2			
	1.0	0.5	0.25	0.15	1.0	0.5	0.25	0.15
GridMap	21.0	42.7	153.4	\times^1	18.3	40.5	133.9	\times
Octomap	186.8	305.5	570.9	1312.6	176.5	289.3	533.1	966.0
SR&CR(Grid)	24.7	64.1	233.2	\times	22.9	58.9	208.2	\times
SR&CR(Octo)	34.8	105.4	373.1	1197.9	30.2	103.6	391.8	\times
Ours	19.1	41.1	121.0	277.9	17.2	36.5	111.3	278.3

¹ “ \times ” denotes that this method failed due to exceeding memory limitation (i.e., 64 GB RAM + 64 GB swap memory).

IX. EXTENSIONS

A. Occupancy Mapping in Large-scale Environment

To evaluate the performance of our D-Map in large-scale occupancy mapping, we conduct an experiment using two sequences from Ford Multi-AV Seasonal Dataset [72]. These two sequences, originally namely “2017-10-26-V2-Log1” and “2017-10-26-V2-Log2”, are referred to as *ford_1* and *ford_2*, respectively. The mapping areas for these sequences measure 8090 m \times 11 494 m \times 96 m and 8107 m \times 11 659 m \times 103 m. We compare our D-Map with four benchmark methods, as discussed in Section VII. The corresponding results are presented in Table VI. The findings indicate that D-Map consistently outperforms the other methods across all resolutions in the context of large-scale applications. Notably, D-Map exhibits superior memory efficiency, enabling it to handle a resolution of 0.15 m within such expansive mapping areas. In contrast, grid-based maps (i.e., GridMap and SR&CR(Grid)) fail to meet the memory limitations imposed by the device.




Fig. 17. A demonstration of D-Map with map region sliding. Figures (a), (b), and (c) showcase the mapping process as the vehicle moves from left to right, with the mapping region sliding accordingly. The white point clouds represent the accumulated historical point clouds. The colored space within the mapping region of D-Map represents the occupancy information. An axis-aligned bounding box is employed to provide a visual representation of the mapping region, outlined by orange lines.

B. Map Region Sliding for D-Map

In the context of occupancy mapping in even larger-scale environments, the issue of memory consumption becomes a significant concern. Therefore, we introduce a map region sliding technique in D-Map to address this challenge. This approach allows the removal of distant occupancy information from the current vehicle position. Fig. 17 illustrates the concept of map region sliding, wherein the mapping region of D-Map slides alongside the vehicle. The sliding process adjusts the mapping region when the vehicle has moved beyond a distance threshold and removes information outside of the mapping region from the D-Map. By dynamically updating the mapping region, D-Map optimizes memory usage and only retains the relevant occupancy information. This mechanism enables our method to handle large-scale environments while mitigating the impact of memory constraints.

C. Extension to Range Sensors with Measurement Noise

In D-Map, the on-tree update strategy leverages the high precision of depth measurements from LiDAR sensors to remove known cells on the octree directly. We further extend D-Map to range sensors with higher measurement noise (e.g., depth cameras) by incorporating occupancy probabilities. At each update, the occupancy probabilities on the correspondent nodes decrease for cells determined as known, while the occupancy probabilities in the corresponding cells of point clouds in the hashing grid map increase. The removal of D-Map is disabled to obtain correct occupancy probabilities. When querying the map, D-Map summarizes the occupancy probabilities from the hybrid map structure to determine occupancy states. This adaptation of occupancy probabilities appropriately handles measurement noise similar to other

existing occupancy mapping approaches while retaining the superior efficiency of D-Map. We conduct several experiments to validate the performance of D-Map when incorporating occupancy probabilities.

1) Qualitative Evaluation: We conducted a qualitative comparison of the mapping results obtained from the original D-Map and D-Map with occupancy probability against those obtained from Octomap. To evaluate our approach, we used the *pioneer_slam3* sequence in the TUM dataset [73], which was captured by a Kinect depth camera mounted on a Pioneer robot. The original D-Map used the same parameters as those described in Section VII. For D-Map with occupancy probability and Octomap, the hit and miss probabilities by a ray for occupied and free space are set to 0.7 and 0.4, respectively, while the occupancy threshold for determining a cell as occupied is set to 0.9. The mapping results are presented in Fig. 18. The results of the original D-Map and D-Map with occupancy probability (Fig. 18(a) and (b), respectively) highlight the effective handling of measurement noise from the depth camera. Furthermore, the mapping result obtained by D-Map with occupancy probability (Fig. 18(b)) exhibits few discrepancies from that produced by Octomap (Fig. 18(c)), indicating our accurate mapping performance.

2) Efficiency: We conduct benchmark experiments to evaluate our efficiency on depth cameras. We compare the update time of our original D-Map and D-Map with occupancy probability against the four mapping methods in Section VII, using four sequences in the TUM dataset [73]: *pioneer_360*, *pioneer_slam*, *pioneer_slam2*, and *pioneer_slam3*. The results are presented in Table VII, where the original D-Map and D-Map with occupancy probability are referred to as “Ours” and “Ours(prob)” respectively. Super Ray and Culling Region-based Grid Map (i.e., SR&CR(Grid)) performs the best at resolutions of 1.0 m and 0.5 m. However, at high resolutions (i.e., $d < 0.5$ m), our original D-Map and D-Map with occupancy probability have superior performance. On average, the original D-Map and D-Map with occupancy probability achieve approximately 9.1 and 5.3 times faster than the fastest competing method (i.e., GridMap) at the resolution of 5 cm, respectively. The update time of D-Map with occupancy probability is comparable to that of the original D-Map at low resolutions (i.e., $d \geq 0.25$ m). However, at high resolutions (i.e., $d \leq 0.1$ m), D-Map with occupancy probability is slower than the original D-Map since a large number of cells are kept on the octree to maintain occupancy probabilities.

3) Accuracy: In addition to the previous experiments, we conduct simulation experiments to evaluate the accuracy of D-Map with occupancy probability. The simulation environment has a dimension of $30 \text{ m} \times 20 \text{ m} \times 4 \text{ m}$, as depicted in Fig. 19(a). We evaluate the Area Under the Receiver Operating Characteristic (AUROC) curves of D-Map with occupancy probability, Grid Map, and Octomap by removing different fractions of observed data, with results provided in Fig. 19(b). The results show that our D-Map with occupancy probability has a similar mapping performance as Octomap and Grid Map.




Fig. 18. The mapping results of the sequence *pioneer_slam3* in TUM dataset. (a) The original D-Map (b) D-Map with occupancy probability. (c) Octomap

TABLE VII
COMPARISON OF UPDATE TIME (MS) AT DIFFERENT RESOLUTION ON TUM DATASET

Resolution [m]	pioneer_360					pioneer_slam					pioneer_slam2					pioneer_slam3				
	1.0	0.5	0.25	0.1	0.05	1.0	0.5	0.25	0.1	0.05	1.0	0.5	0.25	0.1	0.05	1.0	0.5	0.25	0.1	0.05
GridMap	29.7	48.3	85.1	195.8	383.2	30.4	47.1	83.1	186.7	377.9	32.7	51.8	90.9	211.5	408.0	28.5	43.2	76.4	173.0	334.3
Octomap	93.3	154.0	273.3	611.9	1138.7	91.3	148.9	262.8	582.1	1135.2	106.5	179.4	318.4	699.7	1237.6	83.7	138.5	246.5	542.2	1001.4
SR&CR(Grid)	6.8	10.2	34.7	180.5	482.9	7.0	11.0	35.2	169.8	472.7	7.5	11.8	38.3	202.7	532.5	6.4	9.5	31.2	159.4	436.6
SR&CR(Octo)	6.7	10.3	31.1	192.7	485.9	7.1	10.5	31.3	184.7	514.7	7.6	11.3	35.2	213.1	508.0	6.2	9.2	28.0	172.4	464.6
Ours(prob)	26.3	25.7	27.5	37.6	73.9	29.0	27.7	29.5	38.5	70.3	28.6	29.2	30.2	39.6	70.8	25.3	26.1	27.3	36.0	66.5
Ours	27.0	26.8	27.8	31.5	41.2	30.5	28.9	29.9	33.2	40.9	30.3	30.4	30.7	34.5	43.2	26.8	27.3	27.7	31.0	38.9




Fig. 19. (a) The simulation environment for accuracy evaluation. (b) The AUROC curves of D-Map with occupancy probability, Grid Map, and Octomap for different fractions of removed data.

X. DISCUSSION

In this section, we first discuss occupancy mapping using a depth image in terms of efficiency and accuracy, followed by the discussion on parallel processing over D-Map.

A. Occupancy Mapping on Depth Image

1) *Efficiency*: ray-casting is an indispensable component in the existing occupancy mapping framework for occupancy updates. However, the computational demands of ray-casting increase with the number of point clouds and the longer detection range of high-resolution LiDARs, which makes it unsuitable for computationally limited robotic applications. To alleviate the increasing computation load in occupancy mapping, we propose an alternative approach that leverages a depth image created from point clouds to determine occupancy states. Moreover, we design a novel on-tree occupancy update strategy that exploits the hierarchical structure of the octree to determine the occupancy state of larger cells, avoiding the need for timely updates on smaller cells as required by ray-casting-based methods. Additionally, we remove known cells from the octree, reducing the map size and further lowering the update cost. The significant reduction in both the number of cells to be

updated and the cost of map update renders substantially high efficiency in D-Map, as verified in the benchmark experiments in Section VII.

2) *Accuracy*: The accuracy of D-Map can be affected by two primary factors: depth map rasterization and occupancy state determination, which are both related to the depth map resolution.

In the rasterization process, using a low-resolution depth image might result in the loss of depth measurements since a pixel only keeps one point with the smallest depth, as described in Section IV-A. In the occupancy state determination process, D-Map determines occupancy states in spherical coordinates on the depth image rather than Cartesian coordinates. Several design choices affect the accuracy of D-Map, including 1) projecting cells onto a depth image using their circumsphere radius, which can result in information loss at the corners of the cells; 2) discretizing the projected area into pixels on a depth image, which can result in a distorted shape in 3D space; and 3) using an occupancy state determination method that allows for early determination of large regions if the observation completeness α surpasses a threshold ϵ , which may cause errors in the unobserved regions.

Despite a slight decrease in accuracy, D-Map provides comparable accuracy with the existing mapping approaches while achieving higher efficiency thanks to the comprehensive analysis and appropriate selection of depth image resolution.

B. Parallel Processing over D-Map

Although the primary focus during the design of D-Map does not prioritize parallel processing, certain aspects of the update process in D-Map lend themselves well to parallel implementation. Specifically, 1) The projection of a point cloud onto a depth image can be parallelized at the pixel level. 2) The update process of the hashing grid map naturally lends itself to parallel processing, allowing concurrent updates to different grid cells. However, it is worth noting that the octree structure employed in D-Map poses challenges for

parallelization. Parallelizing the octree's operations is not straightforward due to its hierarchical nature. As an alternative, we suggest utilizing a B-tree [74] to exploit parallel processing on efficient management of 3D data within D-Map.

XI. CONCLUSION

This paper proposes a novel framework for occupancy mapping termed D-Map, which aims to provide efficient occupancy updates for high-resolution LiDAR sensors. Our proposed framework consists of three key techniques. Firstly, a method has been proposed to determine the occupancy state of a cell at arbitrary size through depth image projection. Secondly, a hybrid map structure has been developed with an efficient on-tree update strategy. Thirdly, a removal strategy has been introduced, which utilizes the low false alarm rate of LiDARs to remove known cells from the map. These techniques work in conjunction to reduce the number of cells that need updating and lower the cost of map updates, resulting in a significant improvement in efficiency.

To validate our proposed framework, we provide theoretical analyses of its accuracy and efficiency and conduct extensive benchmark experiments on various LiDAR datasets. The results show that D-Map substantially improves efficiency against other state-of-the-art mapping methods while maintaining comparable accuracy and high memory efficiency. Two real-world applications were demonstrated to showcase the effectiveness and efficiency of D-Map for high-resolution LiDAR-based applications.

In the future, we could extend our D-Map framework to support occupancy mapping in dynamic environments and exploit parallel processing.

ACKNOWLEDGEMENT

The authors would like to thank DJI Innovation Technology for the equipment support during the whole work. They would like to thank Ms. Yuhua Xie for the helpful discussions and Prof. Ximin Lyu for the support of the experiment site. They would also like to thank Mr. Wendi Dong for his help with the handheld equipment.

REFERENCES

- [1] F. Gao, W. Wu, W. Gao, and S. Shen, "Flying on point clouds: Online trajectory generation and autonomous navigation for quadrotors in cluttered environments," *Journal of Field Robotics*, vol. 36, no. 4, pp. 710–733, 2019.
- [2] Y. Ren, F. Zhu, W. Liu, Z. Wang, Y. Lin, F. Gao, and F. Zhang, "Bubble planner: Planning high-speed smooth quadrotor trajectories using receding corridors," in *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2022, pp. 6332–6339.
- [3] Y. Ren, S. Liang, F. Zhu, G. Lu, and F. Zhang, "Online whole-body motion planning for quadrotor using multi-resolution search," in *2023 IEEE International Conference on Robotics and Automation (ICRA)*, 2023, pp. 1594–1600.
- [4] F. Zhu, Y. Ren, F. Kong, H. Wu, S. Liang, N. Chen, W. Xu, and F. Zhang, "Swarm-lio: Decentralized swarm lidar-inertial odometry," in *2023 IEEE International Conference on Robotics and Automation (ICRA)*, 2023, pp. 3254–3260.
- [5] J. Leonard, J. How, S. Teller, M. Berger, S. Campbell, G. Fiore, L. Fletcher, E. Frazzoli, A. Huang, S. Karaman, et al., "A perception-driven autonomous urban vehicle," *Journal of Field Robotics*, vol. 25, no. 10, pp. 727–774, 2008.
- [6] J. Levinson, J. Askeland, J. Becker, J. Dolson, D. Held, S. Kammel, J. Z. Kolter, D. Langer, O. Pink, V. Pratt, et al., "Towards fully autonomous driving: Systems and algorithms," in *2011 IEEE intelligent vehicles symposium (IV)*. IEEE, 2011, pp. 163–168.
- [7] Y. Li, L. Ma, Z. Zhong, F. Liu, M. A. Chapman, D. Cao, and J. Li, "Deep learning for lidar point clouds in autonomous driving: A review," *IEEE Transactions on Neural Networks and Learning Systems*, vol. 32, no. 8, pp. 3412–3432, 2020.
- [8] J. Zhang and S. Singh, "Loam: Lidar odometry and mapping in real-time," in *Robotics: Science and Systems*, vol. 2, no. 9. Berkeley, CA, 2014, pp. 1–9.
- [9] J. Lin, C. Yuan, Y. Cai, H. Li, Y. Zou, X. Hong, and F. Zhang, "Immesh: An immediate lidar localization and meshing framework," *arXiv preprint arXiv:2301.05206*, 2023.
- [10] X. Liu, Z. Liu, F. Kong, and F. Zhang, "Large-scale lidar consistent mapping using hierarchical lidar bundle adjustment," *IEEE Robotics and Automation Letters*, vol. 8, no. 3, pp. 1523–1530, 2023.
- [11] Y. Li and J. Ibanez-Guzman, "Lidar for autonomous driving: The principles, challenges, and trends for automotive lidar and perception systems," *IEEE Signal Processing Magazine*, vol. 37, no. 4, pp. 50–61, 2020.
- [12] B. T. Lopez and J. P. How, "Aggressive 3-d collision avoidance for high-speed navigation," in *ICRA*, 2017, pp. 5759–5765.
- [13] F. Kong, W. Xu, Y. Cai, and F. Zhang, "Avoiding dynamic small obstacles with onboard sensing and computation on aerial robots," *IEEE Robotics and Automation Letters*, vol. 6, no. 4, pp. 7869–7876, 2021.
- [14] S. Liu, M. Watterson, S. Tang, and V. Kumar, "High speed navigation for quadrotors with limited onboard sensing," in *2016 IEEE international conference on robotics and automation (ICRA)*. IEEE, 2016, pp. 1484–1491.
- [15] J. Tordesillas, B. T. Lopez, and J. P. How, "Faster: Fast and safe trajectory planner for flights in unknown environments," in *2019 IEEE/RSJ international conference on intelligent robots and systems (IROS)*. IEEE, 2019, pp. 1934–1940.
- [16] Z. Zhang, T. Henderson, S. Karaman, and V. Sze, "Fsmi: Fast computation of shannon mutual information for information-theoretic mapping," *The International Journal of Robotics Research*, vol. 39, no. 9, pp. 1155–1177, 2020.
- [17] B. Zhou, Y. Zhang, X. Chen, and S. Shen, "Fuel: Fast uav exploration using incremental frontier structure and hierarchical planning," *IEEE Robotics and Automation Letters*, vol. 6, no. 2, pp. 779–786, 2021.
- [18] W. Tabib, K. Goel, J. Yao, C. Boirum, and N. Michael, "Autonomous cave surveying with an aerial robot," *IEEE Transactions on Robotics*, vol. 38, no. 2, pp. 1016–1032, 2021.
- [19] A. Elfes, *Robot navigation: Integrating perception, environmental constraints and task execution within a probabilistic framework*. Springer Berlin Heidelberg, 1996, pp. 91–130.
- [20] A. Hornung, K. M. Wurm, M. Bennewitz, C. Stachniss, and W. Burgard, "Octomap: An efficient probabilistic 3d mapping framework based on octrees," *Autonomous robots*, vol. 34, no. 3, pp. 189–206, 2013.
- [21] *Livox Avia User Manual*, Livox Technology Company Limited, Oct. 2020, <https://www.livoxtech.com/avia>.
- [22] S. Thrun, W. Burgard, and D. Fox, *Probabilistic Robotics*. MIT Press, 2005.
- [23] S. T. O'Callaghan and F. T. Ramos, "Gaussian process occupancy maps," *The International Journal of Robotics Research*, vol. 31, no. 1, pp. 42–62, 2012.
- [24] S. Kim and J. Kim, "Building occupancy maps with a mixture of gaussian processes," in *2012 IEEE International Conference on Robotics and Automation*. IEEE, 2012, pp. 4756–4761.
- [25] ———, "Gpmap: A unified framework for robotic mapping based on sparse gaussian processes," in *Field and service robotics*. Springer, 2015, pp. 319–332.
- [26] F. Ramos and L. Ott, "Hilbert maps: Scalable continuous occupancy mapping with stochastic gradient descent," *The International Journal of Robotics Research*, vol. 35, no. 14, pp. 1717–1730, 2016.
- [27] K. Doherty, J. Wang, and B. Englot, "Probabilistic map fusion for fast, incremental occupancy mapping with 3d hilbert maps," in *2016 IEEE international conference on robotics and automation (ICRA)*. IEEE, 2016, pp. 1011–1018.
- [28] C. O'Meadhra, W. Tabib, and N. Michael, "Variable resolution occupancy mapping using gaussian mixture models," *IEEE Robotics and Automation Letters*, vol. 4, no. 2, pp. 2015–2022, 2018.
- [29] K. Doherty, T. Shan, J. Wang, and B. Englot, "Learning-aided 3-d occupancy mapping with bayesian generalized kernel inference," *IEEE Transactions on Robotics*, vol. 35, no. 4, pp. 953–966, 2019.

[30] T. Duong, M. Yip, and N. Atanasov, "Autonomous navigation in unknown environments with sparse bayesian kernel-based occupancy mapping," *IEEE Transactions on Robotics*, vol. 38, no. 6, pp. 3694–3712, 2022.

[31] L. Mescheder, M. Oechsle, M. Niemeyer, S. Nowozin, and A. Geiger, "Occupancy networks: Learning 3d reconstruction in function space," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2019, pp. 4460–4470.

[32] B. Mildenhall, P. P. Srinivasan, M. Tancik, J. T. Barron, R. Ramamoorthi, and R. Ng, "Nerf: Representing scenes as neural radiance fields for view synthesis," *Communications of the ACM*, vol. 65, no. 1, pp. 99–106, 2021.

[33] J. J. Park, P. Florence, J. Straub, R. Newcombe, and S. Lovegrove, "Deepsdf: Learning continuous signed distance functions for shape representation," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2019, pp. 165–174.

[34] A. Asgharivaskasi and N. Atanasov, "Semantic octree mapping and shannon mutual information computation for robot exploration," *IEEE Transactions on Robotics*, vol. 39, no. 3, pp. 1910–1928, 2023.

[35] A. Rosinol, M. Abate, Y. Chang, and L. Carlone, "Kimera: an open-source library for real-time metric-semantic localization and mapping," in *2020 IEEE International Conference on Robotics and Automation (ICRA)*, 2020, pp. 1689–1696.

[36] Z. Zhu, S. Peng, V. Larsson, W. Xu, H. Bao, Z. Cui, M. R. Oswald, and M. Pollefeys, "Nice-slam: Neural implicit scalable encoding for slam," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 12786–12796.

[37] M. C. Martin and H. P. Moravec, *Robot evidence grids*. Carnegie Mellon University, the Robotics Institute, 1996.

[38] Y. Roth-Tabak and R. Jain, "Building an environment model using depth information," *Computer*, vol. 22, no. 6, pp. 85–90, 1989.

[39] G. M. Hunter and K. Steiglitz, "Operations on images using quad trees," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, no. 2, pp. 145–153, 1979.

[40] C. L. Jackins and S. L. Tanimoto, "Oct-trees and their use in representing three-dimensional objects," *Computer Graphics and Image Processing*, vol. 14, no. 3, pp. 249–270, 1980.

[41] P. Payeur, P. Hébert, D. Laurendeau, and C. M. Gosselin, "Probabilistic octree modeling of a 3d dynamic environment," in *Proceedings of International Conference on Robotics and Automation*, vol. 2. IEEE, 1997, pp. 1289–1296.

[42] M. Nießner, M. Zollhöfer, S. Izadi, and M. Stamminger, "Real-time 3d reconstruction at scale using voxel hashing," *ACM Transactions on Graphics (ToG)*, vol. 32, no. 6, pp. 1–11, 2013.

[43] X. Zhou, Z. Wang, H. Ye, C. Xu, and F. Gao, "Ego-planner: An esdf-free gradient-based local planner for quadrotors," *IEEE Robotics and Automation Letters*, vol. 6, no. 2, pp. 478–485, 2020.

[44] B. Zhou, H. Xu, and S. Shen, "Racer: Rapid collaborative exploration with a decentralized multi-uav system," *IEEE Transactions on Robotics*, 2023.

[45] H. Oleynikova, Z. Taylor, M. Fehr, R. Siegwart, and J. Nieto, "Voxblox: Incremental 3d euclidean signed distance fields for on-board mav planning," in *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2017, pp. 1366–1373.

[46] L. Han, F. Gao, B. Zhou, and S. Shen, "Fiesta: Fast incremental euclidean distance fields for online motion planning of aerial robots," in *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2019, pp. 4423–4430.

[47] K. M. Wurm, D. Hennes, D. Holz, R. B. Rusu, C. Stachniss, K. Konolige, and W. Burgard, "Hierarchies of octrees for efficient 3d mapping," in *2011 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2011, pp. 4249–4255.

[48] E. Einhorn, C. Schröter, and H.-M. Gross, "Finding the adequate resolution for grid mapping-cell sizes locally adapting on-the-fly," in *2011 IEEE International Conference on Robotics and Automation*. IEEE, 2011, pp. 1843–1848.

[49] M. Yguel, O. Aycard, and C. Laugier, "Update policy of dense maps: Efficient algorithms and sparse representation," in *Field and service robotics*. Springer, 2008, pp. 23–33.

[50] D. Duberg and P. Jensfelt, "Ufomap: An efficient probabilistic 3d mapping framework that embraces the unknown," *IEEE Robotics and Automation Letters*, vol. 5, no. 4, pp. 6411–6418, 2020.

[51] I. Wald, T. Ize, A. Kensler, A. Knoll, and S. G. Parker, "Ray tracing animated scenes using coherent grid traversal," in *ACM SIGGRAPH 2006 Papers*, 2006, pp. 485–493.

[52] Y. Kwon, D. Kim, I. An, and S.-e. Yoon, "Super rays and culling region for real-time updates on grid-based occupancy maps," *IEEE Transactions on Robotics*, vol. 35, no. 2, pp. 482–497, 2019.

[53] J. L. Bentley and D. Wood, "An optimal worst case algorithm for reporting intersections of rectangles," *IEEE Transactions on Computers*, vol. 29, no. 07, pp. 571–577, 1980.

[54] V. K. Vaishnavi, "Computing point enclosures," *IEEE Transactions on Computers*, vol. 31, no. 01, pp. 22–29, 1982.

[55] Y. Cai, F. Kong, Y. Ren, F. Zhu, J. Lin, and F. Zhang. (2023) Supplementary materials for occupancy grid mapping without ray casting for high-resolution lidar sensors. [Online]. Available: <https://github.com/hku-mars/D-Map/blob/main/documents/SupplementaryMaterials.pdf>

[56] M. R. Spiegel, S. Lipschutz, and J. Liu, *Schaum's Outlines: Mathematical Handbook of Formulas and Tables*. McGraw-Hill, 2009, vol. 2.

[57] M. Teschner, B. Heidelberger, M. Müller, D. Pomerantes, and M. H. Gross, "Optimized spatial hashing for collision detection of deformable objects," in *Vmv*, vol. 3, 2003, pp. 47–54.

[58] D. Meagher, "Geometric modeling using octree encoding," *Computer graphics and image processing*, vol. 19, no. 2, pp. 129–147, 1982.

[59] W. Xu, Y. Cai, D. He, J. Lin, and F. Zhang, "Fast-lio2: Fast direct lidar-inertial odometry," *IEEE Transactions on Robotics*, vol. 38, no. 4, pp. 2053–2073, 2022.

[60] A. Geiger, P. Lenz, C. Stiller, and R. Urtasun, "Vision meets robotics: The kitti dataset," *The International Journal of Robotics Research*, vol. 32, no. 11, pp. 1231–1237, 2013.

[61] W. Xu and F. Zhang, "Fast-lio: A fast, robust lidar-inertial odometry package by tightly-coupled iterated kalman filter," *IEEE Robotics and Automation Letters*, vol. 6, no. 2, pp. 3317–3324, 2021.

[62] F. Kong, X. Liu, B. Tang, J. Lin, Y. Ren, Y. Cai, F. Zhu, N. Chen, and F. Zhang, "Marsim: A light-weight point-realistic simulator for lidar-based uavs," *IEEE Robotics and Automation Letters*, vol. 8, no. 5, pp. 2954–2961, 2023.

[63] J. Amanatides and A. Woo, "A fast voxel traversal algorithm for ray tracing," in *Eurographics*, vol. 87, no. 3, 1987, pp. 3–10.

[64] S. Mystakidis, "Metaverse," *Encyclopedia*, vol. 2, no. 1, pp. 486–497, 2022.

[65] Y. Wang, Z. Su, N. Zhang, R. Xing, D. Liu, T. H. Luan, and X. Shen, "A survey on metaverse: Fundamentals, security, and privacy," *IEEE Communications Surveys & Tutorials*, 2022.

[66] P. Cipresso, I. A. C. Giglioli, M. A. Raya, and G. Riva, "The past, present, and future of virtual and augmented reality research: a network and cluster analysis of the literature," *Frontiers in psychology*, vol. 8, p. 2086, 2018.

[67] S. Shah, D. Dey, C. Lovett, and A. Kapoor, "Airsim: High-fidelity visual and physical simulation for autonomous vehicles," in *Field and Service Robotics: Results of the 11th International Conference*. Springer, 2018, pp. 621–635.

[68] W. Tabib, K. Goel, J. Yao, C. Boirum, and N. Michael, "Autonomous cave surveying with an aerial robot," *IEEE Transactions on Robotics*, vol. 38, no. 2, pp. 1016–1032, 2021.

[69] K. Themistocleous, C. Mettas, E. Evagorou, and D. Hadjimitsis, "The use of uavs and photogrammetry for the documentation of cultural heritage monuments: the case study of the churches in cyprus," in *Earth resources and environmental remote sensing/GIS applications X*, vol. 11156. SPIE, 2019, pp. 85–96.

[70] F. Zhu, Y. Ren, and F. Zhang, "Robust real-time lidar-inertial initialization," in *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2022, pp. 3948–3955.

[71] G. Lu, W. Xu, and F. Zhang, "On-manifold model predictive control for trajectory tracking on robotic systems," *IEEE Transactions on Industrial Electronics*, pp. 1–10, 2022.

[72] S. Agarwal, A. Vora, G. Pandey, W. Williams, H. Kourous, and J. McBride, "Ford multi-av seasonal dataset," *The International Journal of Robotics Research*, vol. 39, no. 12, pp. 1367–1376, 2020.

[73] J. Sturm, N. Engelhard, F. Endres, W. Burgard, and D. Cremers, "A benchmark for the evaluation of rgb-d slam systems," in *Proc. of the International Conference on Intelligent Robot Systems (IROS)*, Oct. 2012.

[74] M. A. Awad, S. Ashkiani, R. Johnson, M. Farach-Colton, and J. D. Owens, "Engineering a high-performance gpu b-tree," in *Proceedings of the 24th symposium on principles and practice of parallel programming*, 2019, pp. 145–157.



Yixi Cai received a B.E. degree in Automation, in 2020, from the School of Automation Science and Electrical Engineering, Beihang University (BUAA), Beijing, China. He is currently a Ph.D. candidate in the Department of Mechanical Engineering, University of Hong Kong (HKU), Hong Kong, China.

Mr. Cai is currently engaged in robotics research, with a particular emphasis on developing efficient mapping techniques using light detection and ranging (LiDAR) for purposes such as localization, navigation, and exploration.



Fanze Kong received the B.S. degree in Flight Vehicle Design and Engineering from Harbin Institute of Technology, Harbin, China, in 2020. He is currently working toward the Ph.D. degree in Mechanical Engineering with the Department of the University of Hong Kong, Hong Kong, China.

Mr. Kong's research interests include UAV design, UAV motion planning and multi-UAV autonomous navigation.



Yunfan Ren received a B.E. degree in Automation in 2021 from the School of Mechanical Engineering and Automation, Harbin Institute of Technology (Shenzhen). He is currently a Ph.D. candidate in the Department of Mechanical Engineering, University of Hong Kong (HKU), Hong Kong, China.

Mr. Ren's research interests include robotic state estimation, motion planning, and aerial swarm systems.



Fangcheng Zhu received the B.E. degree in Automation in 2021 from the School of Mechanical Engineering and Automation, Harbin Institute of Technology (Shenzhen). He is currently a Ph.D. candidate in the Department of Mechanical Engineering, University of Hong Kong (HKU), Hong Kong, China.

Mr. Zhu's research interests include LiDAR-based simultaneous localization and mapping (SLAM), sensor calibration, and aerial swarm systems.



Jiarong Lin received a B.S. degree in Optical Information Science and Technology from the University of Electronic Science and Technology of China (UESTC) in 2015. He is currently a Ph.D. candidate in the Department of Mechanical Engineering, the University of Hong Kong (HKU), Hong Kong, China.

Mr. Lin's research interests include light detection and ranging (LiDAR) mapping and sensor fusion.



Fu Zhang received the B.E. degree in automation from the University of Science and Technology of China (USTC), Hefei, Anhui, China, in 2011, and the Ph.D. degree in controls from the University of California at Berkeley, Berkeley, CA, USA, in 2015.

Dr. Zhang joined the Department of Mechanical Engineering at the University of Hong Kong (HKU), Hong Kong, as an Assistant Professor in August 2018. Dr. Zhang's current research interests are on robotics and controls, with a focus on unmanned aerial vehicle (UAV) design, navigation, control, and

light detection and ranging (LiDAR)-based simultaneous localization and mapping (SLAM).