

ROG-Map: An Efficient Robocentric Occupancy Grid Map for Large-scene and High-resolution LiDAR-based Motion Planning

Yunfan Ren, Yixi Cai, Fangcheng Zhu, Siqi Liang and Fu Zhang

Abstract—Recent advances in LiDAR technology have opened up new possibilities for robotic navigation. Given the widespread use of occupancy grid maps (OGMs) in robotic motion planning, this paper aims to address the challenges of integrating LiDAR with OGMs. To this end, we propose ROG-Map, a uniform grid-based OGM that maintains a local map moving along with the robot to enable efficient map operation and reduce memory costs for large-scene autonomous flight. Moreover, we present a novel incremental obstacle inflation method that significantly reduces the computational cost of inflation. The proposed method outperforms state-of-the-art methods on various public datasets. To demonstrate the effectiveness and efficiency of ROG-Map, we integrate it into a complete quadrotor system and perform autonomous flights against both small obstacles and large-scale scenes. During real-world flight tests with a 0.05 m resolution local map and 30 m×30 m×6 m local map size, ROG-Map takes only 29.8% of frame time on average to update the map at a frame rate of 50 Hz (*i.e.*, 5.96 ms in 20 ms), including 0.33% (*i.e.*, 0.66 ms) to perform obstacle inflation, which represents only half of the total map updating time when compared to the state-of-the-art baseline. We release ROG-Map as an open-source ROS package¹ to promote the development of LiDAR-based motion planning.

I. INTRODUCTION

LiDAR-based autonomous unmanned aerial vehicles (UAVs) have seen significant advancements in various applications, such as search and rescue, inspection, and autonomous exploration. Compared to depth cameras with a sensing range of about 3 ~ 5 m, LiDAR sensors provide more precise and long-range (typically ranging from tens to hundreds of meters) three-dimensional measurements, extending the perception range for autonomous UAVs. Additionally, precise LiDAR points enable autonomous UAVs to avoid small obstacles [1], making it possible for them to operate in more challenging areas.

Occupancy mapping is an essential component of an autonomous aerial system that navigates in unknown environments. The occupancy grid map (OGM) is one of the most promising map structures for real-time occupancy mapping, as verified in [2]–[4], with three critical abilities carefully designed for motion planning

Y. Ren, Y. Cai, F. Zhu, and F. Zhang are with the Department of Mechanical Engineering, University of Hong Kong {renyf, yixicai, zhufc}@connect.hku.hk, {fuzhang}@hku.hk, S. Liang is with School of Mechanical Engineering and Automation, Harbin Institute of Technology sqliang@stu.hit.edu.cn.

¹<https://github.com/hku-mars/ROG-Map>

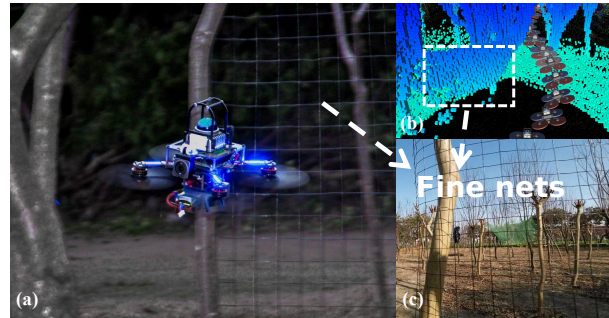


Fig. 1. (a) Fast autonomous navigation against fine metal nets in cluttered environments using Livox Mid360 LiDAR. The FPV camera is only for visualization. (b) The occupied grids in ROG-Map with a resolution of 0.05 m during the flight. (c) Fine wire nets with a width of 3 mm, which require high-resolution and efficient mapping algorithms for accurate online detection. The UAV is equipped with fully onboard sensing, computation, and control. More details can be found in the attached video at <https://youtu.be/eDkwGXCa7w>.

as follows. Firstly, ray casting is used to distinguish the occupied, free, and unknown regions of the environment. Secondly, an inflation technique is applied to the occupied space to guarantee safety in consideration of the UAV’s size. Finally, OGM with high enough resolution is aware of small obstacles, enabling navigation and obstacle avoidance in complex environments.

Recent research in LiDAR-based motion planning has posed new challenges. To fully utilize the capabilities of LiDAR, which provides long-range and accurate measurements, an ideal OGM for LiDAR is expected to maintain high resolution and update over a long distance with high efficiency. However, these requirements significantly increase the computation load for map updates and obstacle inflation. As a result, traditional octree-based [5]–[7], hashing-based [8] or hybrid structure with bosh hashing and octree-based [9] OGMs encounter great difficulties in their real-time ability for LiDAR-based motion planning with limited onboard computation resources. An alternative is the uniform grid-based approach [2], which has high computational efficiency. However, the extensive memory consumption of uniform grid maps makes them impractical for large-scale environments.

To address these issues in LiDAR-based motion planning, we propose **ROG-Map**, a uniform grid-based OGM that is computationally efficient and enables long-range, high-resolution occupancy mapping, leveraging the advanced sensing capabilities of LiDAR sensors. Our

algorithm takes the LiDAR position and the registered point cloud in the world frame as input and builds and updates the occupancy information in real time. The contributions of this paper are as follows:

- 1) We propose a novel incremental inflation method relying on the change of occupancy state, which achieves a computational complexity of $O(n)$ where n is the number of changed grids. This method enables faster obstacle inflation with consistent accuracy compared to existing methods.
- 2) We adopt a zero-copy map sliding strategy with an accelerated indexing process to maintain a local map moving with the robot, making the proposed map suitable to operate at high resolution and in large-scale environments.
- 3) We compare the ROG-Map against SOTA baselines on public datasets, demonstrating its superiority in computation and memory efficiency. Additionally, ROG-Map is integrated into a LiDAR-based quadrotor to conduct extensive real-world tests to verify its outstanding real-world performance.
- 4) The ROG-Map is implemented as a ROS package with carefully engineered work and is open-sourced to promote LiDAR-based motion planning.

II. RELATED WORKS

A. Occupancy Grid Map

The occupancy grid map (OGM) is a promising navigation map type for robots, capable of distinguishing between occupied, free, and unknown environmental areas through ray casting and handling sensor noise and dynamic objects through probabilistic updates. Additionally, it serves as a foundation for other mapping algorithms, such as the Euclidean distance field map [10], [11]. Existing methods for implementing occupancy maps can be divided into three main streams: octree-based [5], hash table-based [8], and uniform grid-based [2]. We begin by analyzing these methods from the perspective of time complexity. Octree-based methods have a time complexity of $O(\log n)$ for map operations such as insertion, change, and query, where n represents the number of nodes in the tree. In contrast, hash table-based methods have a theoretical time complexity of $O(1)$, but the worst-case time complexity of the operations on the map is $O(n)$ due to hash conflicts [12]. Moreover, since both occupied and free grid cells are maintained on OGMs, the number of hash conflicts increases with a denser map at a higher resolution, which can reduce overall performance. Considering the large number of map operations required in ray casting and obstacle inflation with LiDAR, both of these map types can suffer from working in real time with onboard computation devices. In contrast, uniform grid-based OGMs [2] ensure that the time complexity of all map operations is $O(1)$ under all circumstances and they are typically several times faster than octree-based and hashing-based methods in map operations.

Next, we focus on the space complexity of those methods. Hashing-based and octree-based OGMs exhibit a space complexity of $O(n)$, where n denotes the number of nodes in the map. Octree-based methods further optimize memory usage by merging adjacent grid cells into larger ones, resulting in higher memory efficiency compared to hashing-based methods. Notably, both methods demonstrate a space complexity that remains independent of the map size. In contrast, uniform grid-based methods allocate memory for all grid cells in the map and have a space complexity of $O(m)$, where m represents the map size. As the memory consumption scales linearly with map size, this method is not suitable for larger-scale missions. To address these limitations, we propose ROG-Map, a uniform grid-based approach for efficient map operations. To address the space complexity issue, we introduce a zero-copy map sliding strategy that keeps a robo-centric local map with a fixed size. In this way, the memory consumption of ROG-Map is constant, making it well-suited for missions in large-scale environments.

B. Obstacle Inflation

In robot motion planning, inflating obstacles is a widely used technique for generating the robot's configuration space [3], [13]. Modeling the robot as a point mass in the configuration space can significantly simplify and accelerate the motion planning algorithm. Traditional obstacle inflation algorithms, such as the one presented in [2], compute the bounding box of all point clouds in each input frame. After ray casting and probability updating, they traverse all the grid cells in the bounding box, marking occupied grid cells and their neighbors as `InflatedOccupied`, and the remaining grid cells as `non-InflatedOccupied`. However, for LiDAR point clouds, the bounding box is often large, making map traversal and inflation time-consuming. To address this issue, Li *et al.* proposed an incremental update algorithm in their work [14]. This algorithm leverages a queue-based approach, specifically the *Rising Queue* and *Falling Queue*, to monitor grid cells that transition from non-occupied to occupied and from occupied to non-occupied, respectively. The algorithm sets all grid cells in the *Rising Queue* and their neighbors as `InflatedOccupied`. Then, for each grid cell in the *Rising Queue*, it traverses all its neighbors and the neighbors' neighbors to decide if the grid cell and its neighbor should be set as `non-InflatedOccupied`. However, this process has a worst-case complexity of $O(n^2)$, where n is the number of changed grid cells. In contrast, we propose a novel incremental update scheme that ensures $O(n)$ computation complexity for all cases. Our proposed method reduces the number of traversed grid cells by 70%~97% on public datasets compared to [14], significantly accelerating the obstacle inflation process.

III. OCCUPANCY GRID MAPPING

In this section, we introduce the fundamental concept of the occupancy grid map, including the probabilistic update process and the definition of grid cell states.

At the k -th update, the occupancy grid map (OGM) takes the LiDAR position $\mathbf{x}_k \in \mathbb{R}^3$ and a scan of LiDAR points \mathcal{P}_k in the world frame (e.g., the inertial frame or the frame of the odometry) as input and fuses the measurements using Bayesian updating [5], [15]. If a LiDAR point falls in a grid cell, it is considered a *hit*, while if the LiDAR beam passes through a grid cell, it is considered a *miss*. Assuming the map update process is Markovian, we estimate the occupancy probability of a grid cell \mathbf{n} given the history of measurements up to time k , $P(\mathbf{n}|\mathbf{x}_{1:k}, \mathcal{P}_{1:k})$, denoted as $P_{1:k}(\mathbf{n})$ for brevity, by

$$P_{1:k}(\mathbf{n}) = [1 + \mathbf{P}]^{-1}$$

$$\mathbf{P} = \frac{1 - P_k(\mathbf{n})}{P_k(\mathbf{n})} \frac{1 - P_{1:k-1}(\mathbf{n})}{P_{1:k-1}(\mathbf{n})} \frac{P(\mathbf{n})}{1 - P(\mathbf{n})} \quad (1)$$

where $P(\mathbf{n})$ is a prior probability, which is commonly assumed as $P(\mathbf{n}) = 0.5$ to indicate that the map has no prior information of the occupancy state. By using the log-odd notation

$$L_{(\cdot)}(\mathbf{n}) = \log \left(\frac{P_{(\cdot)}(\mathbf{n})}{1 - P_{(\cdot)}(\mathbf{n})} \right), \quad (2)$$

the equation (1) can be rewritten as

$$L_{1:k}(\mathbf{n}) = L_{1:k-1}(\mathbf{n}) + L_k(\mathbf{n}) \quad (3)$$

Let the log-odds value of a hit be denoted as l_{hit} and that of a miss as l_{miss} , then $L_k(\mathbf{n})$ can be computed as follows:

$$L_k(\mathbf{n}) = n_{\text{hit}} \cdot l_{\text{hit}} + n_{\text{miss}} \cdot l_{\text{miss}}, \quad (4)$$

where n_{hit} and n_{miss} represent the number of hits and misses, respectively, for the grid cell at the k -th update. It should be noted that l_{miss} is negative, which means that the occupancy probability decreases when a LiDAR beam passes through a grid cell.

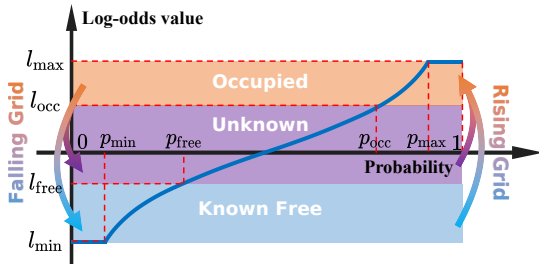


Fig. 2. The log-odds function and the grid cell state corresponding to the log-odds value. A grid cell changes from Unknown or KnownFree to Occupied is called a rising cell (RC), while a grid cell changes from Occupied to Unknown or KnownFree is called a failing cell (FC).

Since the log-odds operation is bijective, we only save $L_{1:k}(\mathbf{n})$ in our map. As described in [16], to ensure the adaptability in both static and dynamic environments, we use a clamping update policy which defines an upper and lower bound on $P_{1:k}(\mathbf{n})$. In this way, the occupancy state is estimated by

$$L_t = L_{1:k-1}(\mathbf{n}) + L_k(\mathbf{n})$$

$$L_{1:k}(\mathbf{n}) = \max(\min(L_t, l_{\text{max}}), l_{\text{min}}) \quad (5)$$

where l_{max} and l_{min} are the upper and lower bound on the log-odds value.

$$\text{state}(\mathbf{n}) = \begin{cases} \text{KnownFree} & l_{\text{min}} \leq L(\mathbf{n}) < l_{\text{free}} \\ \text{Occupied} & l_{\text{occ}} \leq L(\mathbf{n}) \leq l_{\text{max}} \\ \text{Unknown} & \text{otherwise} \end{cases} \quad (6)$$

The clamped log-odds function is shown in Fig. 2. A grid cell whose state changes from Unknown or KnownFree to Occupied is called a rising cell (RC), and a grid cell whose state changes from Occupied to KnownFree or Unknown is called a failing cell (FC). The concept of rising and failing cell detection is essential for our novel incremental updating algorithms, which will be detailed in Sec. IV-B

IV. THE ROG-MAP

The ROG-Map is a uniform grid-based OGM with two main parts in its update process **1) Map sliding** (Sec. IV-A), which involves updating the local map origin and resetting the memory outside of the local map, and **2) Map update** (Sec. IV-B), which includes probabilistic update and incremental obstacle inflation.

A. Robocentric Local Maps

In most robotic navigation missions, the robots only need to consider the surrounding environment. To reduce memory consumption and enable navigation missions within large and unbounded scenes, we introduce a zero-copy map sliding strategy that only maintains a local map around the robot. Unlike existing methods [14], [17] that require knowledge of the local map origin to compute memory addresses for points, we propose a novel indexing method that relies solely on the mapping resolution and map size. The independence of our indexing process from the local map origin allows the pre-construction of a look-up table offline, enhancing indexing performance during map operations.

The indexing process is illustrated in Fig.3. Assume the local map size is $\mathbf{s} = (s_x, s_y, s_z) \in \mathbb{Z}_+^3$ and the discrete resolution is r . Without loss of generality, we assume all elements of \mathbf{s} are odd. All data is saved in an array. For an arbitrary 3D point $\mathbf{p} = (p_x, p_y, p_z) \in \mathbb{R}^3$, we define its global index $\mathbf{i}^g = (i_x^g, i_y^g, i_z^g)$, and its local index $\mathbf{i}^l = (i_x^l, i_y^l, i_z^l)$. In each dimension $d \in \{x, y, z\}$, the indexes can be computed by:

$$i_d^g = \text{round}(p_d/r),$$

$$i_d^l = \text{normalize}(i_d^g \bmod s), \quad (7)$$

where mod is the modulo function and $\text{normalize}()$ is an normalize function to ensure $i_k^l \in [0, s_k - 1]$:

$$\text{normalize}(x, s) = \begin{cases} x - \lfloor \frac{s}{2} \rfloor & x > \lfloor \frac{s}{2} \rfloor \\ x + \lfloor \frac{s}{2} \rfloor & -\lfloor \frac{s}{2} \rfloor \leq x \leq \lfloor \frac{s}{2} \rfloor \\ x + 3\lfloor \frac{s}{2} \rfloor & x < -\lfloor \frac{s}{2} \rfloor \end{cases}, \quad (8)$$

where $\lfloor \cdot \rfloor$ is the floor function. Then we define the indexing function:

$$\text{toAddress}(\mathbf{i}^l) = i_x^l \cdot s_y \cdot s_z + i_y^l \cdot s_z + i_z^l, \quad (9)$$

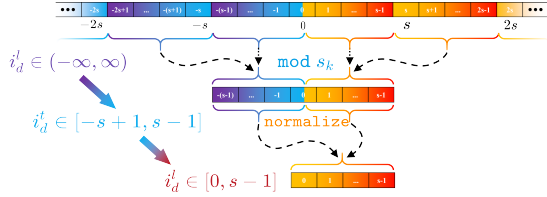


Fig. 3. The indexing process establishes a mapping from the unbounded global index i_d^g to a distinct local index i_d^l without requiring knowledge of the local map origin. The variable $i_d^l = i_d^g \text{ mod } s$ serves as an intermediate variable solely intended to facilitate visualization.

We also provide a visual example of the relationship between the local map and memory, as illustrated in Fig. 4. We only show the map at the robot's height to ease the visualization. As the origin of the map (i.e., the robot) moves, the area within the yellow dashed box at t_{k-1} goes out of the local map at t_k , hence its corresponding space in the memory is reset. The reset space is then used to save newly encountered area shown in the green dashed box. The red dashed boxes indicate that the addresses of the cells remain unchanged in memory before and after the local map's movement. To avoid frequent memory resetting, the map sliding operation for updating the map center and resetting the memory outside the map is only executed when the distance between the robot and the current center of the local map is greater than a threshold d . It is worth noticing that the mapping from the global to the local index is a surjective map, and each local index corresponds to multiple global indexes. However, the global index can be uniquely determined from the local index given the map's origin and the range of the local map.

B. Probabilistic Update and Incremental Inflation

The map updating process is presented in Alg. 1. For the k -th scan of point cloud \mathcal{P}_k and the associated position of the LiDAR \mathbf{x}_k , the first step in the map updating process is ray casting, as shown in Line 3. To accomplish this, we use a fast voxel traversal algorithm [18] to traverse all grid cells between each LiDAR point and \mathbf{x}_k . Since multiple rays may cross the same grid cell, we use a cache \mathcal{C} to save all traversed grid cells and the number of hits c_{hit} and misses c_{miss} following [5]. This

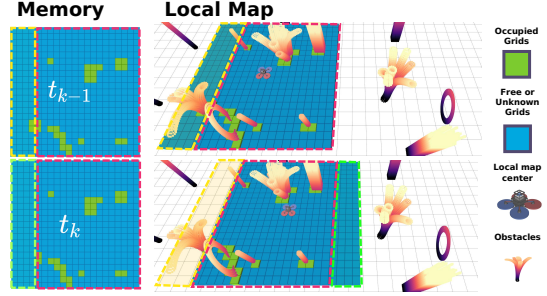


Fig. 4. The visualization of the local map and memory at the robot's height. The Unknown and KnownFree cells are colored in blue, while the Occupied cells are colored in green. Between time steps t_{k-1} and t_k , the map undergoes a rightward shift. At time step t_k , the local map disengages from the yellow area and enters the new green region. Simultaneously, the memory associated with the yellow area is cleared, while the green region is stored in memory at the position of the yellow area at t_{k-1} . The red dashed box delineates the unchanged region and its corresponding memory is also unchanged.

batch update approach significantly reduces the number of map operations in the following probabilistic update process. The ray casting process is encapsulated in the function $\mathcal{C} = \text{raycasting}(\mathbf{x}_k, \mathcal{P}_k)$. After performing ray casting, we process all candidate grid cells in \mathcal{C} (Lines 4-13). The log-odds probability is updated using (3) in Line 6.

Algorithm 1: Local Map Update and Obstacle Inflation

```

1 Notation: The robot's position  $\mathbf{x}_k$ ; Input point cloud  $\mathcal{P}_k$ ; The updating candidates queue  $\mathcal{C}$ ; The spherical update list  $\mathcal{L}$ .
2 Algorithm
3    $\mathcal{C} \leftarrow \text{raycasting}(\mathbf{x}_k, \mathcal{P}_k)$ ;
4   foreach  $\mathbf{n} \in \mathcal{C}$  do
5      $\mathbf{n}_{\text{last}} \leftarrow \mathbf{n}$ ;
6      $\mathbf{n} \leftarrow \text{UpdateProbability}(\mathbf{n})$ ;
7     if  $\text{IsRaisingCell}(\mathbf{n}_{\text{last}}, \mathbf{n})$  then
8       |  $\text{UpdateNeighborCounter}(\mathbf{n}, +1)$ ;
9     end
10    if  $\text{IsFallingCell}(\mathbf{n}_{\text{last}}, \mathbf{n})$  then
11      |  $\text{UpdateNeighborCounter}(\mathbf{n}, -1)$ ;
12    end
13  end
14 Algorithm End
15 Function  $\text{UpdateNeighborCounter}(\mathbf{n}, c)$ 
16   foreach  $\mathbf{p} \in \mathcal{L}$  do
17     |  $\mathbf{n}_{\text{temp}} = \mathbf{n} + \mathbf{p}$ ;
18     |  $\mathbf{n}_{\text{temp}}.\text{inflationCounter} += c$ ;
19   end
20 End Function

```

Then, we perform obstacle inflation by detecting the rising cells (RCs) and falling cells (FCs) defined in Sec. III. For each grid cell in the local map, an integer is used to count how many neighbors of this grid cell are Occupied, with the counter initially set to zero. In order to perform obstacle inflation, it is necessary to traverse all neighbors within the inflation distance of

TABLE I
BENCHMARK COMPARISON

| | New College (Ave. 156.2 points per frame) | | | | | | HKU RSC (Ave. 23157.6 points per frame) | | | | | |
|----------------|---|--------------|----------------|----------------|---------------|---------------|---|---------------|------------------|----------------------|---------------|--------------|
| | t_{tot} (ms) | t_u (ms) | n_{inf} | t_{inf} (ms) | t_q (ms) | m (MB) | t_{tot} (ms) | t_u (ms) | n_{inf} | t_{inf} (ms) | t_q (ms) | m (MB) |
| OctoMap [5] | 2.563 | 0.983 | 34991.262 | 1.625 | 61.463 | 640.809 | 2255.409 | 1502.598 | 27258042.859 | 752.811 | 182.306 | 186.7 |
| HashMap | 1.719 | 0.218 | 34991.262 | 1.501 | 48.771 | 2101.527 | 2010.611 | 1165.105 | 27258042.859 | 845.506 | 152.322 | 2219.953 |
| UniformMap [2] | 0.224 | 0.095 | 34991.262 | 0.129 | 21.688 | 3519.195 | 133.514 | 70.897 | 27258042.859 | 62.617 | 23.896 | 11643.707 |
| FIIMap [14] | 0.119 | 0.106 | 1617.898 | 0.013 | 22.708 | 3519.230 | 267.990 | 63.759 | 1346163.411 | 204.231 ¹ | 22.104 | 11678.153 |
| ROG-Map | 0.101 | 0.097 | 305.917 | 0.004 | 22.981 | 64.387 | 68.187 | 67.416 | 36059.518 | 0.771 | 20.467 | 1046.243 |

¹ Following [14], we implemented incremental inflation for FIIMap using a queue structure. Although the number of accessed grid cells is smaller, using a queue instead of an array significantly reduces the cache hit rate, resulting in a longer overall computation time.

the current grid cell. Since the relative positions of the neighbors and the current grid cell are known, we can build a look-up table \mathcal{L} to save all neighbor's relative indexes and obtain the neighbor's index using Line 17. For each RG, all neighbor counters should increase by 1, and for each FG, the counter should decrease by 1. This way, when a grid cell's inflation count is greater than or equal to one, it indicates that at least one of its neighbors is Occupied and therefore, the grid cell is considered as InflatedOccupied, as shown in Fig. 5. Our incremental inflation strategy has a time complexity of $O(n)$ in all circumstances, where n is the number of changed grid cells (both RCs and FCs). Compared to [14], which has a worst-case complexity of $O(n^2)$, the proposed method reduces the number of grid cell traversals in obstacle inflation by up to 70% in the New College Dataset and 97% in HKU RSC (detailed in Sec. V).

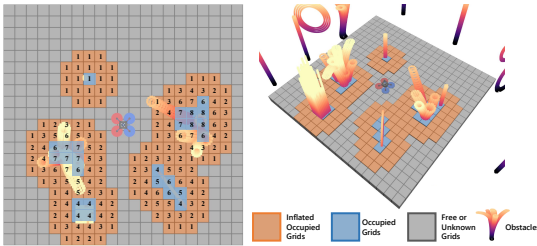


Fig. 5. The visualization of inflationCounter in the local map at the robot's height. The orange and blue cells are InflatedOccupied and the gray cells are KnownFree or Unknown.

V. BENCHMARK COMPARISON

In this section, we compare the proposed ROG-Map with four different types of baselines: **1) OctoMap** [5], which is an octree-based OGM; **2) HashMap**, which is a hash table-based OGM. Since we could not find a suitable open-source hash-table based occupancy grid map, we implemented one based on [8], [19], [20]; **3) UniformMap** [2], which is an OGM with uniform grids and fixed map origin; and **4) FIIMap** [14], a uniform grid-based OGM that features an incremental obstacle inflation approach. As the source code for [14] is not publicly available, we implemented this approach by integrating the incremental inflation algorithm described in the paper into the UniformMap. We compare the average computation time per frame (including map update, and obstacle inflation), the computation time of

probabilistic update t_u , the number of operated grid cells in obstacle inflation n_{inf} , the time consumed for 100,000 random queries t_q , and the memory consumption m . It should be noted that **1) - 3)** inflate obstacles by traversing a local updating box and inflating all occupied grid cells within it, following [2].

We compare the aforementioned methods on two datasets: **1) New College Dataset** an outdoor dataset provided by [5]². The scene size is 250 m × 161 m × 33 m. The resolution of all methods is set to 0.2 m following [5]. The obstacle inflation distance is set to 0.2 m, and the maximum raycast distance is set to 20 m. **2) HKU RSC**³: an indoor dataset proposed by [21]. The size of the dataset is 74 m × 42 m × 20 m, and the mapping resolution is set to 0.05 m with a maximum raycast distance of 20 m. The obstacle inflation distance is set to 0.2 m.

Since ROG-Map's performance is independent of the map size, it only needs to be set based on the available memory size. In both of the two tests, we set it considering the distance used for ray casting (20 m), which is 40 m × 40 m × 12 m.

The benchmark results are presented in Table I. ROG-Map performs similarly to other uniform grid-based approaches in probabilistic update and random query but is several times faster than hashing-based and octree-based methods. In terms of obstacle inflation, ROG-Map requires only 0.13%~0.87% of the map operations of traversal-based methods (*i.e.*, OctoMap, HashMap, and UniformMap) and only 2.67%~18.12% compared to FIIMap, resulting in significantly less time for obstacle inflation and overall map updates. Notably, the HKU RSC dataset was captured at a rate of 10 Hz (*i.e.*, 100 ms per frame), and ROG-Map only requires an average of 68.187 ms to process, highlighting its real-time capability at a high resolution of 0.05 m, while none of the other approaches can run in real-time. From the perspective of memory consumption, OctoMap is the most memory-efficient, followed by HashMap. Uniform grid-based methods have the same space complexity, which increases linearly with the map size and is thus inefficient compared to OctoMap and HashMap. ROG-Map maintains a robocentric local map, resulting in a

²<http://ais.informatik.uni-freiburg.de/projects/datasets/octomap/>

³<https://github.com/hku-mars/MARSIM>

constant space complexity, making it suitable for large-scale missions with acceptable memory consumption.

VI. APPLICATIONS

The ROG-Map was integrated into a LiDAR-based quadrotor platform to assess its real-world performance. The on-board computing unit was an Intel NUC equipped with a CPU i7-10710U. The platform weighed 1.5 kg and boasted a thrust-to-weight ratio over 4.0.

Our perception module utilized the Livox Mid360 LiDAR and Pixhawk’s built-in IMU, running a modified version [22] of FAST-LIO2 [23], which provided high-accuracy state estimation at a frequency of 100 Hz and point clouds at a frequency of 50 Hz. We utilized the method proposed in [24] to calibrate the extrinsic and time-offset between the LiDAR and IMU. For trajectory tracking control, we employed an on-manifold model predictive controller proposed in [25]. To achieve collision avoidance, we utilized a modified version of our previous work [26] as the local planner. Specifically, we employed ROG-Map to identify collision-free paths and generate safe flight corridors.

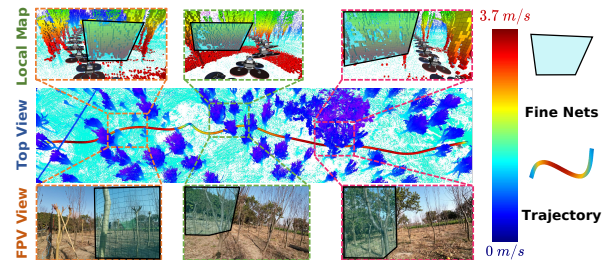


Fig. 6. The *Test 1* was conducted in a small dense forest of short trees where three fine nets made of 3 mm thin metal wires was placed on the flight course. Using ROG-Map, the autonomous quadrotor successfully fly through the scene at a maximum speed over 3.7 m/s.

We conducted eight successful experiments, which were divided into two different categories. Due to space constraints, we only present two of these tests, and other tests can be found in the attached video⁴. In all of the tests, ROG-Map was set to a size of 30 m × 30 m × 6 m with a resolution of 0.05 m. The local map update range was set to 15 m.

The first experiment, labeled as *Test 1*, was conducted within a dense forest, as illustrated in Fig. 6. To create a more challenging environment, we positioned fine metal nets. It should be emphasized that although the LiDAR sensor can detect these delicate nets, the majority of the LiDAR beams traverse through them, with only a few actually hitting the nets. Consequently, it is crucial to ensure that the map resolution is fine enough to prevent the inadvertent exclusion of the fine nets during the ray-casting process of the occupancy mapping. By utilizing ROG-Map, we were able to set the mapping resolution as small as 0.05 m, maintaining a substantial detection

⁴<https://youtu.be/eDkwGXCa7w>

range of 15 m while achieving real-time processing on the onboard computer. As a result, the UAV successfully detected the nets and avoided collisions.

In *Test 2*, we presented a multi-waypoints inspection mission in a large-scale scene. The inspection goal was specified by the user, and the UAV automatically navigated and avoided collisions while moving toward the designated goal positions. The total travel distance was 502.23 m. The UAV successfully traversed all waypoints without a crash, demonstrating the usefulness of our local map-shifting strategy. The generated point cloud and the path of the UAV are shown in Fig. 7.

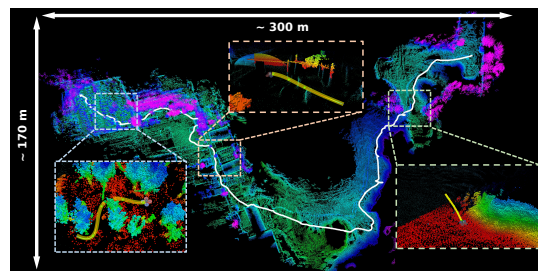


Fig. 7. The point cloud view of *Test 2*. The planned trajectory is highlighted in yellow. The total distance covered during the autonomous flight was 502.23 m.

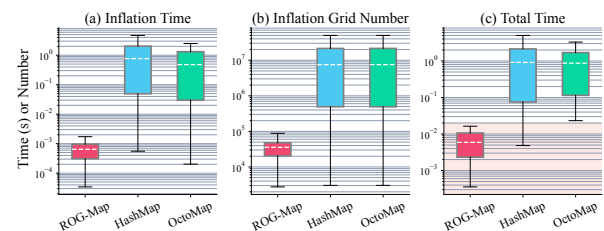


Fig. 8. The computation time of ROG-Map in real-word flight and the baseline running the recorded data on the same computer. The red shaded area in (c) indicated the total time is less than 20 ms, showing ROG-Map can work with LiDAR input in real-time.

The computation time of ROG-Map is presented in Fig. 8. We also run the flight dataset of *Test 2* using Octomap and Hashmap. We didn’t test UniformMap since the memory consumption is beyond the system memory on the onboard computer. As the point cloud input has a frequency of 50 Hz (i.e., 20 ms per frame), neither HashMap nor Octomap can achieve real-time as shown in Fig 8(c), where the red shaded area is computation time less than 20 ms. Our **ROG-Map** only takes an average of 5.96 ms in *Test 1* and 3.46 ms in *Test 2* per frame for local map updates, including only 0.172 ms in *Test 1* and 0.06 ms in *Test 2* for obstacle inflation, demonstrating its high efficiency for robotic planning missions.

VII. CONCLUSION AND FUTURE WORK

In this paper, we propose ROG-Map as a solution for utilizing OGMs with LiDAR sensors. We implement a uniform grid-based OGM to maintain a local

map around the robot, employing a zero-copy map sliding strategy to ensure computational and memory efficiency. Additionally, we introduce a novel incremental obstacle inflation method that significantly reduces computation time and enhances overall mapping performance. Benchmark comparisons on various public datasets demonstrate its advantages over state-of-the-art methods in terms of computation time and memory consumption. Finally, we integrate ROG-Map into a complete quadrotor system, showcasing its capability for high-resolution LiDAR-based motion planning in large-scale environments.

One limitation of ROG-Map is that as the robot moves away from a region, the occupancy information about the region is cleared. This spatial forgetting mechanism does not affect the robot's performance in obstacle avoidance. However, for other applications, such as autonomous exploration that requires maintaining information about the entire environment, ROG-Map is no longer suitable. One possible solution is to use sparse data structures to record the global map information while using ROG-Map to achieve real-time updates and maintain local map information. In the future, we plan to explore a hybrid map framework that combines global and local maps to support a broader range of tasks.

ACKNOWLEDGMENT

This work is funded by the RGC General Research Fund (Grant No. 17204523) and DJI. The authors extend their gratitude to Sifan Tang for recording the experiments and to Longji Yin, Yuhan Xie, and Fanze Kong for their valuable discussions.

REFERENCES

- [1] Fanze Kong, Wei Xu, Yixi Cai, and Fu Zhang. Avoiding dynamic small obstacles with onboard sensing and computation on aerial robots. *IEEE Robotics and Automation Letters*, 6(4):7869–7876, 2021.
- [2] Xin Zhou, Zhepei Wang, Hongkai Ye, Chao Xu, and Fei Gao. EGO-Planner: An ESDF-free gradient-based local planner for quadrotors. *IEEE Robotics and Automation Letters*, 6(2):478–485, 2021.
- [3] Yunfan Ren, Fangcheng Zhu, Wenyi Liu, Zhepei Wang, Yi Lin, Fei Gao, and Fu Zhang. Bubble planner: Planning high-speed smooth quadrotor trajectories using receding corridors. In *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 6332–6339. IEEE, 2022.
- [4] Boyu Zhou, Fei Gao, Jie Pan, and Shaojie Shen. Robust real-time UAV replanning using guided gradient-based optimization and topological paths. In *2020 IEEE International Conference on Robotics and Automation*, pages 1208–1214. IEEE, 2020.
- [5] Armin Hornung, Kai M Wurm, Maren Bennet, Cyrill Stachniss, and Wolfram Burgard. OctoMap: An efficient probabilistic 3d mapping framework based on octrees. *Autonomous Robots*, 34(3):189–206, 2013.
- [6] Daniel Duberg and Patric Jensfelt. Ufomap: An efficient probabilistic 3d mapping framework that embraces the unknown. *IEEE Robotics and Automation Letters*, 5(4):6411–6418, 2020.
- [7] Nils Funk, Juan Tarrío, Sotiris Papatheodorou, Marija Popović, Pablo F Alcantarilla, and Stefan Leutenegger. Multi-resolution 3d mapping with explicit free space representation for fast and accurate mobile robot motion planning. *IEEE Robotics and Automation Letters*, 6(2):3553–3560, 2021.
- [8] Matthias Nießner, Michael Zollhöfer, Shahram Izadi, and Marc Stamminger. Real-time 3d reconstruction at scale using voxel hashing. *ACM Transactions on Graphics (ToG)*, 32(6):1–11, 2013.
- [9] Yixi Cai, Fanze Kong, Yunfan Ren, Fangcheng Zhu, Jiarong Lin, and Fu Zhang. Occupancy grid mapping without ray-casting for high-resolution lidar sensors. *IEEE Transactions on Robotics*, 2023.
- [10] Helen Oleynikova, Zachary Taylor, Marius Fehr, Roland Siegwart, and Juan 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)*, pages 1366–1373. IEEE, 2017.
- [11] Luxin Han, Fei Gao, Boyu Zhou, and Shaojie 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)*, pages 4423–4430. IEEE, 2019.
- [12] Christer Ericson. *Real-time collision detection*. Crc Press, 2004.
- [13] Sikang Liu, Kartik Mohta, Nikolay Atanasov, and Vijay Kumar. Search-based motion planning for aggressive flight in se (3). *IEEE Robotics and Automation Letters*, 3(3):2439–2446, 2018.
- [14] Yong Li, Lihui Wang, Yuan Ren, Feipeng Chen, and Wenxing Zhu. Fiimap: Fast incremental inflate mapping for autonomous mav navigation. *Electronics*, 12(3):534, 2023.
- [15] H. Moravec and A. Elfes. High resolution maps from wide angle sonar. In *Proceedings. 1985 IEEE International Conference on Robotics and Automation*, volume 2, pages 116–121, 1985.
- [16] Manuel Yguel, Olivier Aycard, and Christian Laugier. Update policy of dense maps: Efficient algorithms and sparse representation. In *Field and service robotics*, volume 42, pages 23–33. Springer, 2008.
- [17] Péter Fankhauser and Marco Hutter. A Universal Grid Map Library: Implementation and Use Case for Rough Terrain Navigation. In Anis Koubaa, editor, *Robot Operating System (ROS) – The Complete Reference (Volume 1)*, chapter 5. Springer, 2016.
- [18] John Amanatides, Andrew Woo, et al. A fast voxel traversal algorithm for ray tracing. In *Eurographics*, volume 87, pages 3–10, 1987.
- [19] Jiarong Lin and Fu Zhang. R³live++: A robust, real-time, radiance reconstruction package with a tightly-coupled lidar-inertial-visual state estimator. *arXiv preprint arXiv:2209.03666*, 2022.
- [20] Jiarong Lin, Chongjian Yuan, Yixi Cai, Haotian Li, Yunfan Ren, Yuying Zou, Xiaoping Hong, and Fu Zhang. Immesh: An immediate lidar localization and meshing framework. *IEEE Transactions on Robotics*, 2023.
- [21] Fanze Kong, Xiyuan Liu, Benxu Tang, Jiarong Lin, Yunfan Ren, Yixi Cai, Fangcheng Zhu, Nan Chen, and Fu Zhang. Marsim: A light-weight point-realistic simulator for lidar-based uavs. *IEEE Robotics and Automation Letters*, 8(5):2954–2961, 2023.
- [22] Fangcheng Zhu, Yunfan Ren, Fanze Kong, Huajie Wu, Siqi Liang, Nan Chen, Wei Xu, and Fu Zhang. Decentralized lidar-inertial swarm odometry. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2023.
- [23] Wei Xu, Yixi Cai, Dongjiao He, Jiarong Lin, and Fu Zhang. Fast-lid2: Fast direct lidar-inertial odometry. *IEEE Transactions on Robotics*, 2022.
- [24] Fangcheng Zhu, Yunfan Ren, and Fu Zhang. Robust real-time lidar-inertial initialization. In *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 3948–3955. IEEE, 2022.
- [25] Guozheng Lu, Wei Xu, and Fu Zhang. On-manifold model predictive control for trajectory tracking on robotic systems. *IEEE Transactions on Industrial Electronics*, 2022.
- [26] Yunfan Ren, Siqi Liang, Fangcheng Zhu, Guozheng Lu, and Fu Zhang. Online whole-body motion planning for quadrotor using multi-resolution search. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2023.