

Scalable and Elastic LiDAR Reconstruction in Complex Environments Through Spatial Analysis

Yiduo Wang¹, Milad Ramezani¹, Matias Mattamala¹ and Maurice Fallon¹

Abstract—This paper presents novel strategies for spawning and fusing submaps within an elastic dense 3D reconstruction system. The proposed system uses spatial understanding of the scanned environment to control memory usage growth by fusing overlapping submaps in different ways. This allows the number of submaps and memory consumption to scale with the size of the environment rather than the duration of exploration. By analysing spatial overlap, our system segments distinct spaces, such as rooms and stairwells on the fly during exploration. Additionally, we present a new mathematical formulation of relative uncertainty between poses to improve the global consistency of the reconstruction. Performance is demonstrated using a multi-floor multi-room indoor experiment, a large-scale outdoor experiment and a simulated dataset. Relative to our baseline, the presented approach demonstrates improved scalability and accuracy.

I. INTRODUCTION

3D reconstruction is a common part of applications such as active mapping [1], [2], collision avoidance [3] and the inspection of industrial assets [4], [5]. Building Information Models (BIM) are commonly available for modern buildings, but there are scenarios where these models no longer represent the real situation, e.g. after renovations or disasters. Although systems have been developed to reconstruct these models offline using point clouds from laser scanners, autonomous exploration and reconstruction in multi-storey environments is still an open challenge in mobile robotics. This has been the motivation of international competitions such as the DARPA SubT challenge [6], [7].

In this paper, we propose a dense 3D reconstruction system which uses inputs from a 3D LiDAR on a mobile robot to map complex environments, such as the three-storey building demonstrated in Fig. 1. Our system creates and maintains local occupancy submaps instead of a global map to account for loop closure corrections and improve the global consistency of the reconstruction. We design strategies for spawning and fusing submaps based on geometric understanding of known spaces, enabling on-the-fly segmentation of areas that are isolated from one another, e.g. individual rooms indoors.

The proposed system also leverages explicitly known free space in its occupancy representation to analyse spatial overlap among submaps, and to fuse submaps together to reduce redundant reconstruction. In addition, by computing

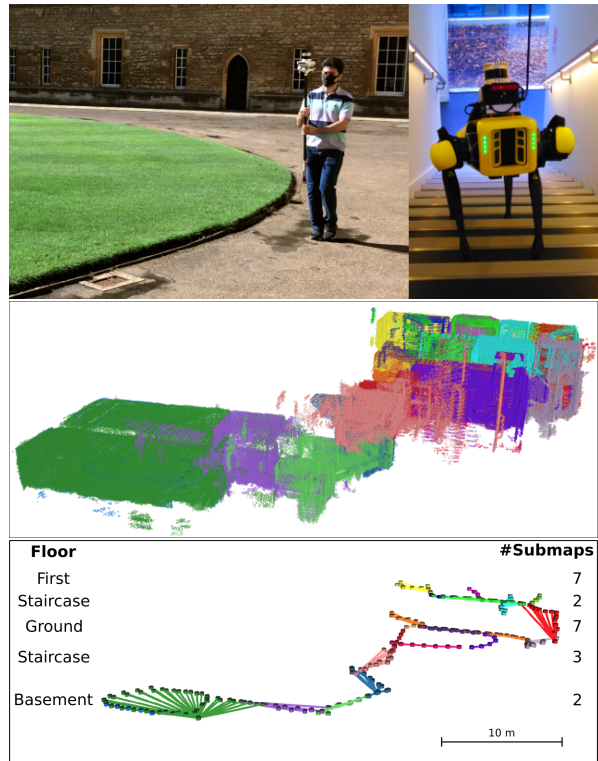


Fig. 1: Our proposed system has been tested in a large-scale outdoor and a multi-floor multi-room indoor environments. **Top**: A handheld LiDAR system exploring New College, Oxford, UK (left) and the quadruped robot Spot climbing stairs in Oxford Robotics Institute (ORI) (right); **Middle**: The reconstruction result of ORI, in which each room is segmented by a unique colour; **Bottom**: The clustered pose graph nodes using spatial overlap analysis.

the relative uncertainty among poses in a SLAM system, our system rejects unreliable submap fusion and improves reconstruction accuracy. We demonstrate the performance of the proposed system using both simulation and real-world experiments, especially its improvement in reconstruction accuracy, system scalability and the capability of room segmentation.

The contributions of our research are the following:

- New strategies for spawning and fusing submaps using probabilistic and spatial understanding.
- A new formulation for relative uncertainty derived from the work of Mangelson *et al.* [8] and GTSAM [9], and a formal treatment of uncertainty in submap fusion.
- Improved reconstruction accuracy and scalability in both outdoor exploration experiments and a multi-

This research is supported by the ESPRC ORCA Robotics Hub (EP/R026173/1). M. Fallon is supported by a Royal Society University Research Fellowship.

¹ These authors are with the Oxford Robotics Institute, University of Oxford, UK. {ywang, milad, matias, mfallon}@robots.ox.ac.uk

storey/multi-room exploration with a legged robot.

The overarching goal of the work is to achieve scalability with the size of the environment instead of the exploration length by controlling the growth of the submap number and memory consumption¹.

The remainder of this paper is organised as follows. In Section II, we review the existing work related to 3D reconstruction and room segmentation. Section III explains our new strategies of spatial overlap analysis and uncertainty formulation. Experimental results are presented in Section IV, and Section V discusses conclusions and future improvement.

II. RELATED WORK

While the literature on dense 3D reconstruction is substantial, in this paper we focus on systems incorporating submaps and methods designed for room segmentation because these aspects are the most relevant to our work.

A. 3D Reconstruction via Submaps

Submapping is a common technique in SLAM systems such as the *Atlas* framework by Bosse *et al.* [10] and DenseSLAM by Nieto *et al.* [11]. These SLAM systems reuse existing submaps for localisation when revisiting a known area rather than remapping that location. Maintaining a collection of submaps keeps memory and computation bounded as opposed to creating a single global map.

Submaps can also enable elasticity when a reconstruction requires correction at the event of a loop closure. Leveraging OctoMap-based [12] submaps, Ho *et al.* [13] and Sodhi *et al.* [14] designed elastic occupancy maps. Voxgraph [15] uses the same strategy to create a globally consistent reconstruction instead represented by a series of Signed Distance Function (SDF) submaps. In our previous work [16], we developed submap-based elastic reconstruction pipelines which supported both occupancy and SDF representations.

These approaches spawn new submaps after either a temporal interval [13]–[15] or distance travelled [16] to bound the local odometry drift within each submap. Extending upon [16], this paper proposes submap spawning and fusion based on a deeper spatial analysis to uniquely represent confined areas such as rooms. This allows local planning to use only a submap in challenging exploration tasks such as [17].

B. 3D Room Segmentation

Separately, several works have developed methods to segment LiDAR reconstructions or floor plans into individual enclosed spaces. Turner and Zakhor [18] designed a room-segmentation pipeline that partitioned 2D point cloud maps into 2.5D building models via triangulation. To achieve 3D reconstruction, the approach assumed that interior walls were vertical and flat.

More sophisticated methods were developed to parse 3D point clouds into rooms, such as detecting void spaces

between walls using point density histograms [19], and extracting planar features before partitioning separate rooms via a multi-label energy minimisation formulation [20], [21]. These methods were limited to single-storey reconstructions.

Ochmann *et al.* [22] and Nikoohemat *et al.* [23] proposed methods that handle unstructured 3D point clouds for multi-storey room segmentation. Ochmann *et al.* introduced a versatile integer linear programming method to incorporate hard constraints, e.g. wall connectivity, to ensure a plausible reconstruction. Nikoohemat *et al.* employed a mobile LiDAR SLAM system and separated building levels by assuming that sloped trajectory segments represent staircase traversals. They further segmented rooms using an adjacency graph of planer segments.

While the effect of partitioning confined spaces is similar to our proposed system, we focus on online segmentation running onboard a mobile robot instead of segmenting a complete reconstruction offline. Our proposed system aims to create segmented 3D reconstructions of confined spaces such as individual rooms and staircases on the fly in the context of multi-floor exploration.

III. METHODS

Our system expands upon the framework of [16] by adding principled strategies which use spatial understanding and pose graph relative uncertainty, to improve the spawning and fusion of submaps in *Graph Clustering*. Fig. 2 provides an overview of the proposed system.

An input to the reconstruction system is a pose graph published by a SLAM system [24], made up of $Q + 1$ nodes $X_k, k \in \{0, \dots, Q\}$ in a *Registered Cloud List*. Each node describes the estimated pose of a LiDAR sensor frame $\{\mathcal{L}_k\}$ with respect to a fixed map frame $\{\mathcal{M}\}$, denoted as ${}^{\mathcal{M}}\mathbf{T}_{\mathcal{L}_k} \in \mathbf{SE}(3)$. Each node is also associated with a raw point cloud C_k expressed in the LiDAR frame $\{\mathcal{L}_k\}$.

The output of the reconstruction system consists of $N + 1$ submaps. Each submap $\mathcal{S}_i, i \in \{0 \dots N\}$ contains an accumulated submap cloud $C_{\mathcal{S}_i}$ in $\{\mathcal{M}\}$, a volumetric occupancy reconstruction in the submap frame $\{\mathcal{S}_i\}$ and its root pose ${}^{\mathcal{M}}\mathbf{T}_{\mathcal{S}_i}$.

We compute and maintain Axis-Aligned Bounding Boxes (AABB) for each submap in $\{\mathcal{M}\}$ using the reconstruction. Submap AABB can be significantly affected by the orientation of $\{\mathcal{M}\}$, so we use AABBs not to accurately estimate scanned space, but as a lightweight method to determine non-overlapping submaps (Sec. III-B).

In the proposed system, we introduce the following set of measures which we use to decide when to spawn or fuse submaps:

- *Cloud Overlap Estimate*: A criteria to adjust the submap spawning decisions made by Local Mapping.
- *Submap Overlap Estimate*: A criteria to propose submap fusion for Global Mapping.
- *Relative Uncertainty*: A criteria to reject unreliable submap fusions.

These measures will be explained in the following sections.

¹Sparsification of the underlying SLAM pose graph is related research topic which we do not explore in this work.




Fig. 2: An overview of the proposed system. *Principled Clustering* contains the new strategies based on cloud and submap overlap estimation as well as relative pose uncertainty in the SLAM graph. This module interacts with *Local Mapping* to spawn new submaps (Sec. III-A), and proposes or rejects submap fusion for *Global Mapping* (Sec. III-B and Sec. III-C).

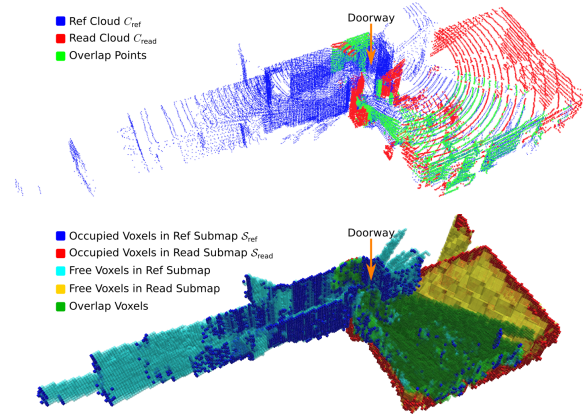


Fig. 3: An example of Cloud Overlap Estimate and Submap Overlap Estimate. **Top**: Limited overlap between the reference cloud C_{ref} and the read cloud C_{read} when entering a room through a narrow doorway. **Bottom**: The voxel overlap between the volumetric occupancy reconstructions of the reference submap S_{ref} and the read submap S_{read} .

A. Cloud Overlap Estimate

Many systems such as [15] and [13] spawn submaps at fixed frequencies to bound the size of each submap as well as robot odometry drift within the submap. Our previous system [16] used a travel distance threshold to achieve the same purpose, under the assumption that the odometry drift is proportional to the distance travelled.

Cloud Overlap Estimate adds another trigger to spawn submaps based on point cloud overlap. Because our odometry system is based on Iterative Closest Point (ICP) [25], performance is affected by major changes in overlap such as when entering a new room [26]. Hence when traversing between two disconnected spaces via a narrow passage, i.e. the scenario presented in Fig. 3 (top), it is beneficial to spawn a new submap and create an elastic connection. In a room network, our proposed system will spawn a new submap when going through a doorway and segment rooms online, as demonstrated in Fig. 1.

Alg. 1 presents how the overlap is measured between the point cloud of a new scan C_{read} and the accumulated submap cloud $C_{S_{ref}}$ of a reference submap S_{ref} . If C_{read} shares sufficient overlap with the accumulated submap cloud $C_{S_{ref}}$, the new scan is integrated into a submap S_{ref} , and the accumulated submap cloud $C_{S_{ref}}$ grows by adding C_{read} .

Algorithm 1: Cloud Overlap Estimate.

input: New LiDAR cloud C_{read} and submap cloud $C_{S_{ref}}$,
output: Cloud overlap ratio $R_{point,read}$

```

begin
  Voxel filter  $C_{read}$  and  $C_{S_{ref}}$  to resolution  $r_{filter}$ 
  for Point  $P_i \in C_{read}$  do
    Search for  $P_{neighbour} \in C_{S_{ref}}$  that is the closest to  $P_i$ 
    if  $\|P_i, P_{neighbour}\| < \sqrt{3} \times r_{filter}$  then
       $N_{point,overlap} = N_{point,overlap} + 1$ 
    end
  end
   $R_{point,read} = N_{point,overlap} / N_{point,read}$ 
  return  $R_{point,read}$ 
end
```

For instance, in Fig. 4 (a) and (b), node \mathcal{L}_{17} is integrated into submap S_5 , and LiDAR scan C_{17} is accumulated into the submap cloud C_{S_5} . However, we constrain point cloud accumulation of each submap by not combining accumulated submap clouds together during submap fusion. In Fig. 4, the volumetric occupancy submaps S_0 and S_5 in (c) are fused together into one submap S'_0 in (d), but submap cloud C_{S_5} is not combined with C_{S_0} . This stops these submap clouds from growing indefinitely as exploration continues.

B. Submap Overlap Estimate

Submap fusion merges existing submaps, and reduces the memory usage of the overall system by fusing repeated reconstructions of the same physical space together. Our previous pipeline [16] triggered submap fusion using the loop closures detected in the SLAM system. Hence it only merged submaps that were created when the robot travelled very close to a previous pose.

In a large-scale (outdoor) environment, a long range (≈ 60 m) LiDAR sensor can repeatedly scan the same space from poses that are far away from one another, resulting in significant redundancy between submaps that loop closure fusion cannot address. Therefore we introduce an additional strategy for submap fusion based on the overlapping scanned spaces, utilising the explicit representation of free space in our volumetric occupancy map. This improves the reconstruction scalability when revisiting explored areas (although the SLAM pose graph continues to grow linearly).

Submap overlap is computed by comparing the volumetric reconstruction as well as the occupancy information stored in




Fig. 4: An example of scan integration and submap fusion. (a) – There are 17 LiDAR scans $\mathcal{L}_{0:16}$ in the existing reconstruction, clustered into 6 submaps $\mathcal{S}_{0:5}$. Each rectangle represents the AABB of each submap. \mathcal{L}_{17} is the latest new LiDAR scan, and there have been no loop closures. (b) – *Graph Clustering* allocates \mathcal{L}_{17} to \mathcal{S}_5 , and the new scan passes the *Cloud Overlap Estimate* criteria described in Sec. III-A. Hence \mathcal{L}_{17} is integrated into \mathcal{S}_5 , expanding the AABB. (c) – There is a new loop closure edge (black line) given by the SLAM pose graph between \mathcal{L}_0 and \mathcal{L}_{17} . The head and tail of the loop closure connection defines the submap overlap search range from \mathcal{S}_0 to \mathcal{S}_5 . (d) – *Graph Clustering* proposes the fusion of submaps \mathcal{S}_5 and \mathcal{S}_0 into \mathcal{S}'_0 , and *Submap Overlap Estimate* (Sec. III-B) proposes the fusion between submaps \mathcal{S}_3 and \mathcal{S}_2 in \mathcal{S}'_2 . Both fusion proposals pass the *Relative Uncertainty* criteria (Sec. III-C) and are executed. All submap indices are updated accordingly. The AABBs of \mathcal{S}_0 and \mathcal{S}_2 are therefore expanded, but their accumulated submap clouds are not, as explained in Sec. III-A.

Algorithm 2: Submap Overlap Estimate.

```

input: Pair of occupancy submaps  $\mathcal{S}_{\text{read}}$  and  $\mathcal{S}_{\text{ref}}$ ,
output: Submap overlap ratios  $R_{\text{voxel,read}}$  and  $R_{\text{voxel,ref}}$ 
begin
  for Voxel  $V_{\text{read}} \subset \mathcal{S}_{\text{read}}$  do
    Find  $V_{\text{ref}} \subset \mathcal{S}_{\text{ref}}$  at the same coordinates as  $V_{\text{read}}$ 
    if  $V_{\text{ref}}$  is not unknown then
      if Both  $V_{\text{ref}}$  and  $V_{\text{read}}$  are free or occupied then
         $N_{\text{voxel,overlap}} = N_{\text{voxel,overlap}} + 1$ 
      end
    end
  end
   $R_{\text{voxel,read}} = N_{\text{voxel,overlap}} / N_{\text{voxel,read}}$ 
   $R_{\text{voxel,ref}} = N_{\text{voxel,overlap}} / N_{\text{voxel,ref}}$ 
  return  $R_{\text{voxel,read}}$  and  $R_{\text{voxel,ref}}$ 
end

```

each individual submap. Fig. 3 (bottom) demonstrates such a case. Alg. 2 describes the estimation of voxel overlap between a pair of submaps $\mathcal{S}_{\text{read}}$ and \mathcal{S}_{ref} . This pair of submaps are fused together if either the ratio $R_{\text{voxel,ref}}$ or $R_{\text{voxel,read}}$ exceed a configurable threshold λ_{fusion} .

To ensure that the root poses of submaps are corrected by loop closure before fusion, we define a submap overlap search range using the head and tail of each loop closure. For example, in Fig. 4 (c), SLAM loop closure is between \mathcal{L}_0 and \mathcal{L}_{17} , and they belong to \mathcal{S}_0 and \mathcal{S}_5 , respectively. Hence the search range for Submap Overlap Estimate is $\mathcal{S}_{0:5}$, and \mathcal{S}_3 and \mathcal{S}_2 are fused together due to significant submap overlap.

Iterating through all voxels is a computationally intense process. Therefore, we add a conservative but efficient preliminary heuristic based on the AABB of each submap before computing submap voxel overlap. Using the AABBs of \mathcal{S}_{ref} and $\mathcal{S}_{\text{read}}$, we compute the volumetric overlap percentages $\{R_{\text{aabb,ref}}, R_{\text{aabb,read}}\}$ and compare them with λ_{fusion} . If both AABB overlaps are smaller than the threshold, such as \mathcal{S}'_1 and \mathcal{S}'_4 in Fig. 4 (c), the proposed system skips computing voxel overlap.

C. Relative Uncertainty

In our current system design, after fusing a pair of submaps the individual submaps are discarded for memory efficiency. Each submap is internally rigid, so local consistency is essential. To improve submap fusion reliability and to retain global consistency, we propose a strategy that uses relative uncertainty between root poses of the submaps which are fused.

The measurement of relative uncertainty allows us to detect and reject uncertain fusions. Fig. 4 presents the case of fusing two pairs of submaps, namely \mathcal{S}_0 and \mathcal{S}_5 , and \mathcal{S}_2 and \mathcal{S}_3 . In the example of fusing \mathcal{S}_0 and \mathcal{S}_5 , we first compute the relative uncertainty between the root poses of \mathcal{S}_0 and \mathcal{S}_5 , which are \mathcal{L}_0 and \mathcal{L}_{15} .

The proposed mathematical model for relative uncertainty between two poses is derived from the same notation as the GTSAM library [9], because we use it as the back-end of our pose graph SLAM system [24]. See Appendix I for the detailed derivation of the computation of relative uncertainty.

We compute the eigenvalues of the relative uncertainty matrix as a quality metric of the fusion. For $\text{SE}(3)$ transformations ${}^{\mathcal{M}}\mathbf{T}_{\mathcal{S}_0}, {}^{\mathcal{M}}\mathbf{T}_{\mathcal{S}_5} \in \mathbb{R}^6$, the relative uncertainty $\Sigma_{\mathcal{S}_0\mathcal{S}_5}$ is a 6×6 matrix, and there are 6 eigenvalues — 3 for translation and 3 for rotation. We compare the 3 eigenvalues of translation against a configurable threshold $\lambda_{\text{uncertainty}}$. If one of the eigenvalues exceeds the threshold, the fusion between \mathcal{S}_0 and \mathcal{S}_5 is rejected.

IV. EXPERIMENTS AND EVALUATION

In this section we evaluate the improvement in system scalability and global consistency brought about by the proposed submap spawning and fusion strategies, compared with our previous reconstruction pipeline [16] as the baseline. The proposed system has been assessed using a large-scale outdoor experiment with a handheld device in the Newer College Dataset (NCD Long experiment) [27] and a multi-storey multi-room exploration experiment with a Boston Dynamics Spot robot (ORI experiment), as shown in Fig. 1. Lastly, we include experiments in a small and a large room network (Fig. 6) using the Gazebo simulator to demonstrate




Fig. 5: The proposed spatial overlap analysis improves the global consistency in the reconstruction. Figure (a) and (b) presents the mesh and volumetric representations of NCD Long experiment, created by the baseline and the proposed system, respectively. The volumetric representations are shown in grey overlaid with submap clusters. The mesh representations are created from the green submaps. Utilising spatial understanding of the environment leads to more reliable submap fusion and therefore better alignment. The mesh created by the previous pipeline has misaligned double surfaces while the proposed system improves the consistency in the mesh.

Experiment	Model	LiDAR properties		
		Vertical FoV	Horizontal FoV	Max range (m)
NCD Long (handheld)	Ouster OS1-64	33.2°	360°	120
ORI (with Spot)	Ouster OS0-64	90°	360°	50
Simulation (with UGV)	Ouster OS0-64	90°	360°	50

TABLE I: LiDAR sensors used in the experiments and their properties. FoV: Field of View

the performance of the proposed system. The exploration in the small room network took on a loopy trajectory to evaluate the effect for long term operation. The large room network was designed to assess the system’s room segmentation capability.

Table I gives details of the different LiDAR sensors used in these experiments. The LiDAR sensors produce organised point cloud scans of 64×1024 points at 10 Hz. The SLAM system created a node in its pose graph every 2m travelled when exploring. The proposed system and the baseline integrated the LiDAR scans at each SLAM node using the *MultiresOFusion* mode in [16] for volumetric occupancy reconstruction. The voxel resolution used in these experiments was 6.5 cm and we integrated LiDAR ranges between 0.5 m and 60 m. These settings give high resolution and long range while retaining 3 Hz integration.

Surface mesh representations of the reconstruction, for example Fig. 5, were created by applying the Marching Cubes algorithm [28] on the zero-crossings of the occupancy.

A. Large-scale Outdoor Experiment

In this section we present the performance of the proposed system and the baseline when tested using NCD Long. This dataset consists of a 2.2 km exploration over 44 min in a $135 \times 122 \text{ m}^2$ environment. The principled clustering strategies in the proposed system improve the consistency of the global reconstruction, as shown in Fig. 5. The proposed system also demonstrates improved scalability in memory usage compared with the baseline.

1) *Reconstruction Accuracy*: Fig. 5 (a) shows the mesh and volumetric reconstructions created by the baseline. In






Fig. 6: The Gazebo environments of a small (top) and a large (bottom) room network for experiments in simulation.

particular the green submap in the Quad area was created as a rigid fusion of several earlier submaps after multiple loop closures were established. As shown in the birds-eye view next to the mesh representation, this submap contains scans taken both in the Quad and the Tunnel. These scans have limited overlap with one another and registration between them is unreliable as a result. This led to a duplicate reconstruction of the indicated wall of the Quad.


By introducing our proposed submap fusion algorithms, the consistency of reconstruction is improved as shown in Fig. 5 (b). The Cloud Overlap Estimate strategy makes better decisions when the handheld device travels between the Quad and the Tunnel — maintaining elastic connections between these two spaces in the global reconstruction. The measurement of relative uncertainty also rejects unreliable submap



(a) The performance in the NCD Long experiment



(b) The performance in the small room network



(c) The performance in the large room network

Fig. 7: The memory usage and submap counters of the proposed system and the baseline system.

fusions. The global volumetric map and the submap mesh both demonstrate improved accuracy in surface alignment.

2) *Memory Consumption*: In the proposed system, submaps that have significantly overlapping scan volumes are fused according to their Submap Overlap Estimate. This is in addition to submap fusion based on loop closures, and properly allows the map to scale with the size of the environment rather than the length of the exploration. Fig. 7 (a) presents the memory usage and submap counter of the proposed system compared to the baseline in NCD Long. We compute the memory usage by summing the size of allocated memory for each submap’s octree in RAM.

For the baseline method the submaps can only be merged when loop closures occur which causes memory usage to grow over time. In contrast, our proposed approach can merge spatially overlapping volumes such that the number of submaps can plateau. For the NCD Long experiment, submap count stabilised at about 30 submaps when the entire environment has been explored at scan 650. Memory usage actually decreased after scan 650 to ~ 18 GB while maintaining the 6.5 cm resolution reconstruction. By the end of the experiment, there was a 65% reduction in memory usage compared to the baseline.





Fig. 8: The volumetric submap reconstructions of each floor (**Top**: first; **Middle**: ground; **Bottom**: basement) in ORI, with each room segmented into unique submaps by the proposed system on the fly during exploration.


B. Multi-storey Multi-room Indoor Exploration

In the ORI experiment a quadruped robot Spot explored three floors of a typical university research lab (Fig. 1); we show the reconstruction of every floor in the building in Fig. 8. New submaps were spawned when the robot entered or exited rooms because the proposed system detected a decrease in cloud overlap. Spatial overlap analysis then merged overlapping submaps in each room, creating a unique reconstruction for each enclosed space. It further ensured that these submaps remained independent allowing future SLAM loop closures to re-position the room submaps as needed. In the supplementary video we demonstrate the full experiment and the incremental mapping of the building.

Segmenting rooms on the fly allows real-time applications



(a) Submap reconstruction in the small room network



(b) Submap reconstruction in the large room network

Fig. 9: The proposed system segments individual rooms and fuses redundant submaps during simulation experiments.

such as path planning and obstacle avoidance to consider only the minimal submaps rather than the entire global reconstruction. The proposed system can thus improve the scalability of other applications.

C. Room Networks in Simulation

To test the performance for long term operation missions, we carried out two experiments in Gazebo simulation with a wheel robot carrying a simulated 3D LiDAR in the environments in Fig. 6.

In the small room network, the mission looped around the environment three times in different directions. As shown by the clustered poses in Fig. 9 (a) there is a clear division between submaps at each doorway. Each room was constructed with either one or two submaps — even after multiple revisits. Fig. 9 (b) further demonstrated the room segmentation capability in the large room network.

Fig. 7 (b) presents the memory usage in the small room network experiment, where the submap count and memory usage both plateaued as the environment was repeatedly scanned. In the large room network, there were many small rooms in this environment, and the proposed system created individual submaps for each of them, leading to a higher submap count than the baseline as shown in Fig. 7 (c). However, the proposed system required about 30% less memory than the baseline to fully map the environment, because submaps in large areas that shared significant overlaps were merged based on spatial analysis, slowing down

memory usage growth compared to the baseline. Meanwhile, submaps of small enclosed rooms occupy limited space in memory and share almost no overlap with other areas. These submaps contributed little to redundant mapping, therefore the proposed system required less memory though created more submaps in the large room network experiment.

V. CONCLUSION AND FUTURE WORK

This paper introduced a set of principled strategies that leverage spatial overlap analysis to improve the spawning and fusion of 3D reconstruction submaps. The proposed system can merge together submaps that scan the same space, allowing the memory consumption of our reconstruction to scale with the volume of the explored area rather than the duration of operation. Additionally, we improved submap fusion reliability by directly computing relative uncertainty between pairs of poses in the SLAM pose graph. These strategies delay submap fusion in 3D reconstruction until there is sufficient confidence to improve global consistency of the overall map.

Furthermore, our system can segment individual rooms on the fly as a robot explores a room network by using spatial overlap. This feature can be beneficial to planners exploring room networks such as in the DARPA SubT challenge.

In future work, we would like to introduce semantic information into the reconstruction, further segmenting objects of interest from the scene for tasks such as industrial inspections. We also aim to design a system that can autonomously infer floor plans from a multi-storey reconstruction, which enables high-level path planning between floors.

APPENDIX I

COMPUTATION OF RELATIVE UNCERTAINTY IN POSE GRAPH SLAM

To compute the relative uncertainty between any pair of nodes in our pose graph, we derived the following formula based on the conventions of GTSAM with reference to Mangelson *et al.* [8].

First we define the probability distribution of the relative transformation from submap \mathcal{S}_i to node \mathcal{S}_j as:

$${}_{\mathcal{S}_i}\mathbf{T}_{\mathcal{S}_j} = {}^{\mathcal{M}}\mathbf{T}_{\mathcal{S}_i}^{-1} {}^{\mathcal{M}}\mathbf{T}_{\mathcal{S}_j} \quad (1)$$

where the poses ${}^{\mathcal{M}}\mathbf{T}_{\mathcal{S}_i}$ and ${}^{\mathcal{M}}\mathbf{T}_{\mathcal{S}_j}$ indicate probability distributions on $\mathbf{SE}(3)$ following a right-hand composition:

$$\mathbf{T} = \bar{\mathbf{T}} \text{Exp}(\xi) \quad (2)$$

$\bar{\mathbf{T}}$ is the mean transformation of the distribution, and ξ is a perturbation that follows a Gaussian distribution. In Eq. (1) we consider that the poses have covariances $\Sigma_{\mathcal{M}\mathcal{S}_i}$ and $\Sigma_{\mathcal{M}\mathcal{S}_j}$, respectively.

In order to derive the expressions for the relative uncertainty, we need the adjoint action of $\bar{\mathbf{T}}$ on ξ , denoted as $\text{Ad}_{\bar{\mathbf{T}}}(\xi)$, defined as follows:

$$\begin{aligned} \text{Ad}_{\bar{\mathbf{T}}}(\xi) &:= \text{Ad}_{\bar{\mathbf{T}}}\xi = \text{Log}(\bar{\mathbf{T}} \text{Exp}(\xi)\bar{\mathbf{T}}^{-1}) \\ \text{Exp}(\text{Ad}_{\bar{\mathbf{T}}}\xi) &= \bar{\mathbf{T}} \text{Exp}(\xi)\bar{\mathbf{T}}^{-1} \\ \bar{\mathbf{T}}^{-1}\text{Exp}(\text{Ad}_{\bar{\mathbf{T}}}\xi) &= \text{Exp}(\xi)\bar{\mathbf{T}}^{-1} \\ \text{Exp}(\xi)\bar{\mathbf{T}} &= \bar{\mathbf{T}} \text{Exp}(\text{Ad}_{\bar{\mathbf{T}}^{-1}}\xi) \end{aligned} \quad (3)$$

Expanding Eq. (1) using Eq. (2) and Eq. (3):

$$\begin{aligned}
& \bar{\mathbf{T}}_{S_j} \text{Exp}({}^{S_i} \xi_{S_j}) \\
&= \text{Exp}(-{}^{\mathcal{M}} \xi_{S_i}) {}^{\mathcal{M}} \bar{\mathbf{T}}_{S_i}^{-1} {}^{\mathcal{M}} \bar{\mathbf{T}}_{S_j} \text{Exp}({}^{\mathcal{M}} \xi_{S_j}) \\
&= {}^{\mathcal{M}} \bar{\mathbf{T}}_{S_i}^{-1} \text{Exp}(-\text{Ad}_{\mathcal{M} \bar{\mathbf{T}}_{S_i}} {}^{\mathcal{M}} \xi_{S_i}) {}^{\mathcal{M}} \bar{\mathbf{T}}_{S_j} \text{Exp}({}^{\mathcal{M}} \xi_{S_j}) \\
&= {}^{\mathcal{M}} \bar{\mathbf{T}}_{S_i}^{-1} {}^{\mathcal{M}} \bar{\mathbf{T}}_{S_j} \text{Exp}(-\text{Ad}_{\mathcal{M} \bar{\mathbf{T}}_{S_j}^{-1}} \text{Ad}_{\mathcal{M} \bar{\mathbf{T}}_{S_i}} {}^{\mathcal{M}} \xi_{S_i}) \text{Exp}({}^{\mathcal{M}} \xi_{S_j})
\end{aligned} \tag{4}$$

Let ${}^{S_i} \bar{\mathbf{T}}_{S_j} \triangleq {}^{\mathcal{M}} \bar{\mathbf{T}}_{S_i}^{-1} {}^{\mathcal{M}} \bar{\mathbf{T}}_{S_j}$, then we can establish the following equivalence:

$$\text{Exp}({}^{S_i} \xi_{S_j}) = \text{Exp}(-\text{Ad}_{\mathcal{M} \bar{\mathbf{T}}_{S_j}^{-1}} \text{Ad}_{\mathcal{M} \bar{\mathbf{T}}_{S_i}} {}^{\mathcal{M}} \xi_{S_i}) \text{Exp}({}^{\mathcal{M}} \xi_{S_j}) \tag{5}$$

The covariance of the perturbation on the left should be equal to the one on the right. However, we cannot compute the covariance directly because of the properties of the exponential map. Instead, we define ${}^{S_i} \xi'_{S_j} = -\text{Ad}_{\mathcal{M} \bar{\mathbf{T}}_{S_j}^{-1}} \text{Ad}_{\mathcal{M} \bar{\mathbf{T}}_{S_i}} {}^{\mathcal{M}} \xi_{S_i}$, and use the Baker-Campbell-Hausdorff (BCH) formula [29] up to first order:

$$\begin{aligned}
E[{}^{S_i} \xi_{S_j} {}^{S_i} \xi'_{S_j} T] &\approx E[{}^{\mathcal{M}} \xi'_{S_i} {}^{\mathcal{M}} \xi'_{S_j} T] + E[{}^{\mathcal{M}} \xi_{S_j} {}^{\mathcal{M}} \xi'_{S_j} T] \\
&\quad + E[{}^{\mathcal{M}} \xi'_{S_i} {}^{\mathcal{M}} \xi_{S_j} T] + E[{}^{\mathcal{M}} \xi_{S_j} {}^{\mathcal{M}} \xi_{S_i} T]
\end{aligned} \tag{6}$$

which, after computing the covariance terms, provides an approximation for the covariance of the relative transformation:

$$\begin{aligned}
\Sigma_{S_i S_j} &\approx (\text{Ad}_{\mathcal{M} \bar{\mathbf{T}}_{S_j}^{-1}} \text{Ad}_{\mathcal{M} \bar{\mathbf{T}}_{S_i}}) \Sigma_{\mathcal{M} S_i} (\text{Ad}_{\mathcal{M} \bar{\mathbf{T}}_{S_j}^{-1}} \text{Ad}_{\mathcal{M} \bar{\mathbf{T}}_{S_i}})^T \\
&\quad + \Sigma_{\mathcal{M} S_j} \\
&\quad - (\text{Ad}_{\mathcal{M} \bar{\mathbf{T}}_{S_j}^{-1}} \text{Ad}_{\mathcal{M} \bar{\mathbf{T}}_{S_i}}) \Sigma_{\mathcal{M} S_i, \mathcal{M} S_j} \\
&\quad - \Sigma_{\mathcal{M} S_i, \mathcal{M} S_j}^T (\text{Ad}_{\mathcal{M} \bar{\mathbf{T}}_{S_j}^{-1}} \text{Ad}_{\mathcal{M} \bar{\mathbf{T}}_{S_i}})^T
\end{aligned} \tag{7}$$

REFERENCES

- [1] A. Bircher, M. Kamel, K. Alexis, H. Oleynikova, and R. Siegwart, "Receding horizon path planning for 3D exploration and surface inspection," *Autonomous Robots*, vol. 42, no. 2, pp. 291–306, 2018.
- [2] A. Dai, S. Papatheodorou, N. Funk, D. Tzoumanikas, and S. Leutenegger, "Fast frontier-based information-driven autonomous exploration with an MAV," in *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, Paris, France, June 2020.
- [3] H. Oleynikova, Z. Taylor, M. Fehr, R. Siegwart, and J. Nieto, "Voxblox: Incremental 3D Euclidean signed distance fields for on-board MAV planning," in *IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, 2017.
- [4] G. A. Hollinger, B. Englot, F. S. Hover, U. Mitra, and G. S. Sukhatme, "Active planning for underwater inspection and the benefit of adaptivity," *Intl. J. of Robotics Research*, vol. 32, no. 1, pp. 3–18, 2013.
- [5] F. S. Hover, R. M. Eustice, A. Kim, B. Englot, H. Johannsson, M. Kaess, and J. J. Leonard, "Advanced perception, navigation and planning for autonomous in-water ship hull inspection," *Intl. J. of Robotics Research*, vol. 31, pp. 1445–1464, 2016.
- [6] A. Bouman, M. F. Ginting, N. Alatur, M. Palieri, D. D. Fan, T. Touma, T. Pailevanian, S.-K. Kim, K. Otsu, J. Burdick, *et al.*, "Autonomous Spot: Long-range autonomous exploration of extreme environments with legged locomotion," *arXiv preprint arXiv:2010.09259*, 2020.
- [7] K. Ebadati, Y. Chang, M. Palieri, A. Stephens, A. Hatteland, E. Heiden, A. Thakur, N. Funabiki, B. Morrell, S. Wood, *et al.*, "Lamp: Large-scale autonomous mapping and positioning for exploration of perceptually-degraded subterranean environments," in *2020 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2020, pp. 80–86.
- [8] J. G. Mangelson, M. Ghaffari, R. Vasudevan, and R. M. Eustice, "Characterizing the uncertainty of jointly distributed poses in the Lie algebra," *IEEE Transactions on Robotics*, vol. 36, no. 5, pp. 1371–1388, 2020.
- [9] F. Dellaert and C. Beall, "GTSAM 4.0," URL: <https://bitbucket.org/gtborg/gtsam>, 2017.
- [10] M. Bosse, P. Newman, J. Leonard, M. Soika, W. Feiten, and S. Teller, "An Atlas framework for scalable mapping," in *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, vol. 2. IEEE, 2003, pp. 1899–1906.
- [11] J. Nieto, J. Guivant, and E. Nebot, "Denseslam: Simultaneous localization and dense mapping," *Intl. J. of Robotics Research*, vol. 25, no. 8, pp. 711–744, 2006.
- [12] 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.
- [13] B.-J. Ho, P. Sodhi, P. Teixeira, M. Hsiao, T. Kusunur, and M. Kaess, "Virtual occupancy grid map for submap-based pose graph slam and planning in 3d environments," in *IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*. IEEE, 2018, pp. 2175–2182.
- [14] P. Sodhi, B.-J. Ho, and M. Kaess, "Online and consistent occupancy grid mapping for planning in unknown environments," in *IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, 2019, pp. 7879–7886.
- [15] V. Reijgwart, A. Millane, H. Oleynikova, R. Siegwart, C. Cadena, and J. Nieto, "Voxgraph: Globally consistent, volumetric mapping using signed distance function submaps," *IEEE Robotics and Automation Letters*, 2020.
- [16] Y. Wang, N. Funk, M. Ramezani, S. Papatheodorou, M. Popovic, M. Camurri, S. Leutenegger, and M. Fallon, "Elastic and efficient lidar reconstruction for large-scale exploration tasks," *IEEE International Conference on Robotics and Automation (ICRA)*, 2021.
- [17] D. Bellicoso, M. Bjelonic, L. Wellhausen, K. Holtmann, F. Guenther, M. Tranzatto, P. Fankhauser, and M. Hutter, "Advances in real-world applications for legged robots," *J. of Field Robotics*, vol. 35, pp. 1311–1326, 2018.
- [18] E. Turner and A. Zakhor, "Floor plan generation and room labeling of indoor environments from laser range data," in *IEEE Intl. Conf. on Computer Graphics Theory and Applications (GRAPP)*, 2014, pp. 1–12.
- [19] I. Armeni, O. Sener, A. R. Zamir, H. Jiang, I. Brilakis, M. Fischer, and S. Savarese, "3d semantic parsing of large-scale indoor spaces," in *Proc. IEEE Int. Conf. Computer Vision and Pattern Recognition*, 2016, pp. 1534–1543.
- [20] C. Mura, O. Mattausch, A. J. Villanueva, E. Gobbetti, and R. Pajarola, "Automatic room detection and reconstruction in cluttered indoor environments with complex room layouts," *Computers & Graphics*, vol. 44, pp. 20–32, 2014.
- [21] C. Mura, O. Mattausch, and R. Pajarola, "Piecewise-planar reconstruction of multi-room interiors with arbitrary wall arrangements," in *Computer Graphics Forum*, vol. 35, no. 7. Wiley Online Library, 2016, pp. 179–188.
- [22] S. Ochmann, R. Vock, and R. Klein, "Automatic reconstruction of fully volumetric 3D building models from oriented point clouds," *ISPRS J. of Photogrammetry and Remote Sensing*, vol. 151, pp. 251–262, 2019.
- [23] S. Nikoohemat, A. A. Diakit , S. Zlatanova, and G. Vosselman, "Indoor 3D reconstruction from point clouds for optimal routing in complex buildings to support disaster management," *Automation in Construction*, vol. 113, p. 103109, 2020.
- [24] M. Ramezani, G. Tinchev, E. Iuganov, and M. Fallon, "Online LiDAR-SLAM for legged robots with robust registration and deep-learned loop closure," in *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, 2020, pp. 4158–4164.
- [25] F. Pomerleau, F. Colas, R. Siegwart, and S. Magnenat, "Comparing ICP Variants on Real-World Data Sets," *Autonomous Robots*, vol. 34, no. 3, pp. 133–148, Feb. 2013.
- [26] S. Nobili, G. Tinchev, and M. Fallon, "Predicting alignment risk to prevent localization failure," in *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, 2018, pp. 1003–1010.
- [27] M. Ramezani, Y. Wang, M. Camurri, D. Wisth, M. Mattamala, and M. Fallon, "The Newer College Dataset: Handheld LiDAR, inertial and vision with ground truth," in *IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, 2020.
- [28] W. E. Lorensen and H. E. Cline, "Marching Cubes: A high resolution 3D surface construction algorithm," *SIGGRAPH Comput. Graph.*, vol. 21, no. 4, p. 163–169, Aug. 1987.
- [29] S. Klarsfeld and J. Oteo, "The Baker-Campbell-Hausdorff formula and the convergence of the Magnus expansion," *Journal of physics A: mathematical and general*, vol. 22, no. 21, p. 4565, 1989.