Manhole Detection and Traversal for Exploration of Ballast Water Tanks using Micro Aerial Vehicles

Mihir Dharmadhikari, Paolo De Petris, Huan Nguyen, Mihir Kulkarni, Nikhil Khedekar, and Kostas Alexis

***Abstract*— This paper presents a method for the autonomous exploration of multiple compartments of a Ballast Water Tank inside a vessel using Micro Aerial Vehicles. Navigation across the compartments of ballast tanks often requires the robot to pass through narrow cross-section "manholes" (e.g., 0.8×0.6m). Hence, this work presents an algorithm to explicitly detect and localize such manholes using 3D LiDAR data and a strategy to reliably navigate through them to enable autonomous exploration of multiple compartments of the tank. Two ablation studies are presented analyzing the effective 3D space with respect to the manhole in which reliable detection takes place. Furthermore, the method is evaluated onboard a collision-tolerant aerial robot in two autonomous exploration experiments in relevant mock-up scenarios.**

I. INTRODUCTION

In recent years robotic systems have been deployed in a variety of industrial and natural environments [1–12]. Building upon this, several researchers along with companies are investigating methods to automate the inspection and surveillance process across industries. Among many, a highly challenging and demanding application is the inspection of Ballast Water Tank (BWTs) inside vessels. These facilities need to be inspected periodically to identify, assess and fix possible corrosion, cracks, and deformations. However, such environments can present varied geometry and scale. Specifically, BWTs can present highly confined geometries, are completely dark, hard to access, and might contain pockets of low/no oxygen and harmful gasses. Furthermore, BWTs consist of several compartments often connected by narrow openings called "Manholes", navigation through which is necessary for accessing all the compartments. Currently, the inspection is mostly done by humans who need to enter these challenging settings posing risk to them, while some teleoperated use of drones and other robots has been used by the industry but without advanced autonomy [13–15]. Simultaneously, apart from the risk to the inspectors' life, the vessel owners suffer heavy economic losses due to the down time during the inspection process. In response, several research groups develop robotic systems and algorithms for the inspection of vessels and tanks inside them [16–22]. These include aerial, crawling and swimming systems, path planning for exploration and inspection, defect detection, and methods for localization inside ballast tanks.

Motivated by the above, in this paper, we present a method for autonomous exploration and mapping of compartments of the ballast tanks using a Micro Aerial Vehicle (MAV). Traditional path planning methods utilize some form of map representation, with volumetric maps like [23, 24] being employed. However, the small size of the manholes (e.g., height $\times$ width $= 0.8\,\mathrm{m} \times 0.6\,\mathrm{m}$) make finding a path through them using volumetric maps of real-time tractable resolution significantly harder. Hence, in this work, we take the approach of explicitly detecting the manholes and localizing their pose which enables the MAV to traverse them. Both for Manhole Detection and Localization (MDL), our method uses the instantaneous point cloud and corresponding depth image data from a 3D LiDAR. In combination with our previous path planning work in [7, 8], the method aims to explore one BWT compartment at a time while detecting and traversing the manholes to navigate between compartments.

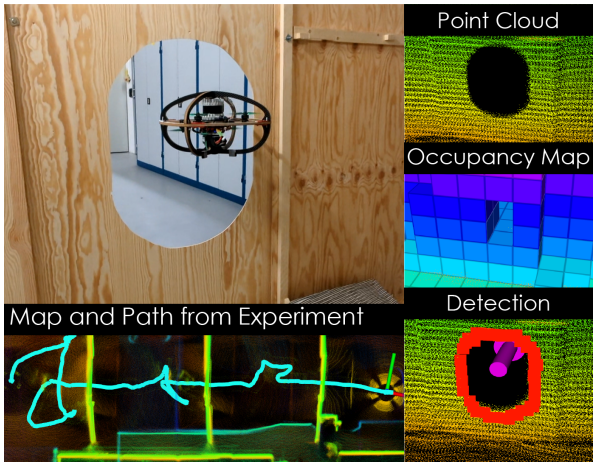

Fig. 1. Instance in one of the experiments showing the robot passing through the manhole along with the final map built and the path traversed by the robot from the mission. On the right, the manhole is shown in the point cloud data and occupancy map, along with its detected pose and boundary.

To verify the MDL method, we present a study analyzing the effective 3D space, with respect to the manhole, in which reliable detection can happen. We study the effect of the viewing angle of the LiDAR relative to the manhole using a) a handheld sensing setup with a LiDAR, and b) with a real robot. Moreover, we deploy our method on a collision-tolerant aerial robot [25] in experiments for the exploration of mock-ups of multiple horizontally-connected compartments as in BWTs.

In the remaining paper, Section II presents related work, followed by the system overview in III, the approach in Section IV, evaluation studies in Section V and conclusions in Section VI.

This material was supported by the Research Council of Norway under project SENTIENT (NO-321435).

The authors are with the Autonomous Robots Lab, Norwegian University of Science and Technology (NTNU), O. S. Bragstads Plass 2D, 7034, Trondheim, Norway mihir.dharmadhikari@ntnu.no

II. RELATED WORK

A niche community of researchers has investigated detection and navigation across manholes or similar narrow openings. The authors in [26] present a deep learning-based method for manhole detection using depth data. The method uses a convolutional neural network to detect the manhole in the depth image and temporal filtering to remove false positives. [27, 28] presents a system for navigation in confined environments including manhole-sized corridors. The method in [27] detects manholes in a stereo-rectified monochromatic image from a fish-eye camera by first extracting regions below a darkness threshold followed by performing a set of morphological operations to remove noise and small dark regions. The manhole is assumed to be rectangular and hence rectangular closed contours among these are selected. In [29], a window detection method using LiDAR as well as a depth camera is proposed. The LiDAR detection method identifies candidate window edge points and associates them with existing edges or uses prior knowledge about the window dimensions to classify a window. The work in [30] presents a system that can navigate sewer pipes following high-level human operator commands. A set of works have investigated flying through narrow gaps using active vision. [31] presents an aggressive trajectory planning method for navigating across narrow gaps that is independent of the underlying gap detection method used. [32] proposes a novel method for flying through arbitrarily shaped narrow openings using purely monocular vision.

III. SYSTEM OVERVIEW

The overall aim of the system is to explore and map a ballast water tank using a MAV. This task typically requires the robot to navigate and explore multiple compartments of the tank connected by manholes. These can be quite narrow (e.g., height $\times$ width $= 0.8\,\mathrm{m} \times 0.6\,\mathrm{m}$) therefore, they may be omitted or under-represented size-wise in occupancy mapping frameworks (Figure 1) and other representations with finite resolution. Thus, finding a path through them is statistically unlikely using a sampling-based path planner employing uniform random sampling. Hence, in this work, we focus on explicitly detecting the pose of the manhole to calculate a path to pass through it. This process is detailed in Sections IV-A and IV-B. The logical flow diagram of the overall method is shown in Figure 2. The robot starts in one compartment of the BWT and begins exploring that compartment. During exploration, it finds manholes that lead to further compartments and stores their poses. When the exploration is completed, the robot aligns itself in front of the closest not-yet traversed manhole, triggers re-detection to confirm that the manhole selected for traversal is not a false detection, and traverses through it by calculating a straight line path through the center of the manhole along the direction perpendicular to its plane. Subsequently, exploration of the next compartment takes place and the process continues iteratively. In this work, we use our previous open-sourced work on Graph-based exploration path planning (GBPlanner) [7, 8], summarized in Section IV-C as the exploration planner.

IV. APPROACH

This section presents details on the manhole detection and localization, traversal, and exploration path planning algorithms.

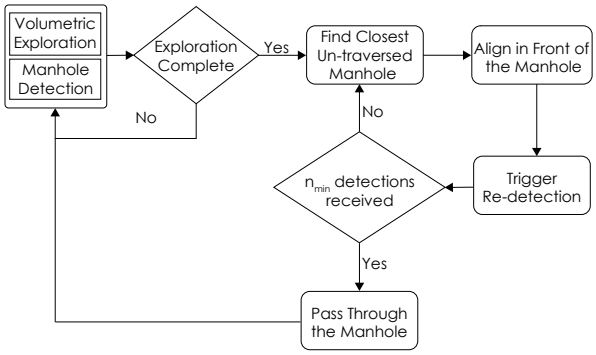

Fig. 2. This figure shows the flow diagram of the method. The robot starts in one compartment and begins the exploration of it. During this time, it detects and stores the poses of all the manholes observed. Upon exploration completion of the current compartment, the planner finds a path to align the robot in front of the closest, previously untraversed, manhole. Once executed, the MDL is re-triggered to verify that the manhole under consideration is not a false detection. If a minimum of n_{min} detections are received, a straight line path through the center of the manhole along its orientation is calculated and commanded to the robot. After traversing the manhole, the exploration planner is triggered and the process is repeated.

A. Manhole Detection and Localization (MDL)

As previously described, navigation across a manhole is statistically unlikely to be successful using a sampling-based planner employing uniform random sampling on an occupancy map in reasonable computation time on the MAV's onboard computer. Hence, our method aims to explicitly detect the 3D position $\mathbf{p} = [x, y, z]$ and orientation $\mathbf{q}$ of each of these manholes to find a path through them.

The detection takes place on the instantaneous point cloud measurements $\mathcal{P}_t$ and the corresponding depth image $\mathcal{I}_t$ obtained from the onboard LiDAR sensor $\mathcal{S}^D$ (the method can in principle be used with any depth sensor). The MDL algorithm is described in Algorithm 1, and Figure 3 shows the steps involved in the MDL algorithm on real data. It is noted that BWT manholes employ standard dimensions with the most typical being that of height $\times$ width $= 0.8\,\mathrm{m} \times 0.6\,\mathrm{m}$ [33, 34].

At a time instance t, the depth image $\mathcal{I}_t$ is processed to remove occlusions to the LiDAR caused by the robot frame. As these appear at fixed pixel locations in the depth image, they are removed simply through a fixed mask. After that, we find all the closed contours $\mathbb{C}_t$ (the subscript t is dropped hereafter for clarity) in the depth image using Canny Edge Detection [35] (Figure 3.1).

These closed contours need to be filtered out to remove contours that are not manholes. To do this, we use the following two properties of the manholes: a) they are planar, and b) they have standardized known dimensions. For each contour $c \in \mathbb{C}$, the points P_c in the point cloud $\mathcal{P}$ corresponding to

Algorithm 1 Manhole Detection and Localization

Inputs $\mathcal{P}_t, \mathcal{I}_t$

$\mathbb{D}_t \leftarrow \emptyset$
$\mathbb{C} \leftarrow$ **detectClosedContours**$(\mathcal{I}_t)$
for $\forall c \in \mathbb{C}$ **do**
  $\mathcal{P}_c \leftarrow$ **3DPoints**$(c, \mathcal{P}_t)$
  $\mathbf{k}_c \leftarrow$ **fitPlane**$(\mathcal{P}_c)$
  $\mathcal{P}_c \leftarrow$ **removeOutliers**$(\mathcal{P}_c, \mathbf{k}_c)$
  $\mathbf{b}_c \leftarrow$ **orientedBoundingBox**$(\mathcal{P}_c)$
  if $\exists k \in [1, m] : l^k_{min} \leq b^l_c \leq l^k_{max}, l \in x, y, z$ **then**
    $\mathbf{p}_c \leftarrow$ **mean**$(\mathcal{P}_c)$
    $\mathbf{q}_c \leftarrow$ **getOrientation**$(\mathbf{k}_c)$
    $\mathbf{p}^w_c, \mathbf{q}^w_c \leftarrow$ **transformToFixedFrame**$(\mathbf{p}_c, \mathbf{q}_c)$
    $\mathbb{D}_t \leftarrow \mathbb{D}_t \bigcup [\mathbf{p}^w_c, \mathbf{q}^w_c]$

the pixels in c are found. As illustrated in Figure 3.2 (image on the left), points corresponding to some pixels in c might lie on the surfaces seen through the manhole whereas some lie on the same surface as the manhole. To correct this error, for each pixel we check the pixels within a 3×3 window around it to find the corresponding point that is closest to the robot. The result of this is shown in the right side image of Figure 3.2.

Even after the above step, certain points can still lie on the surfaces seen through the manhole due to the noise in the depth image (Figure 3.3, left image). To remove them, we first find the best fitting plane to $\mathcal{P}_c$. Let the coefficients of this plane be $\mathbf{k}_c = [\alpha_c, \beta_c, \gamma_c, \delta_c]$. All points whose perpendicular distance from this plane is more than a threshold d_{thp} are removed from $\mathcal{P}_c$ and we get the result shown in the right side image of Figure 3.3.

Next, the point cloud sets $\mathcal{P}_c$ that do not satisfy the dimension criteria, which relates to the minimum oriented bounding box of $\mathcal{P}_c$ not belonging to any of the known manhole dimensions, are rejected (Figure 3.4). Specifically, a minimum oriented bounding box $\mathbf{b}_c = [b^x_c, b^y_c, b^z_c]$, ($b^x_c, b^y_c, b^z_c$ are the x, y, z dimensions of the bounding box respectively) of $\mathcal{P}_c$ is calculated. As the manholes have standard dimensions, the dimensions of $\mathbf{b}_c$ are checked against a set of possible known manhole dimensions. The contours c for which the following criteria do not hold true are rejected: $l^k_{min} \leq b^l_c \leq l^k_{max}, l \in \{x, y, z\}, \ k = 1...m$, where m is the number of possible manhole sizes against which $\mathbf{b}_c$ is checked. All contours c which satisfy the above criteria are selected as valid manhole detections.

The position and orientation of these manholes are then calculated as follows. The position $\mathbf{p}_c$ of the manhole is simply the geometric centroid of the points in $\mathcal{P}_c$. The orientation $\mathbf{q}_c$ of the manhole is set to be along the normal to the plane defined by $\mathbf{k}_c$. Note that both of these are in the sensor's frame and are transformed into the fixed frame using the estimates from the robot's onboard Simultaneous Localization And Mapping (SLAM) solution to obtain the position $\mathbf{p}^w_c$ and orientation $\mathbf{q}^w_c$ in the world frame. All such detections are added to the set $\mathbb{D}_t$.

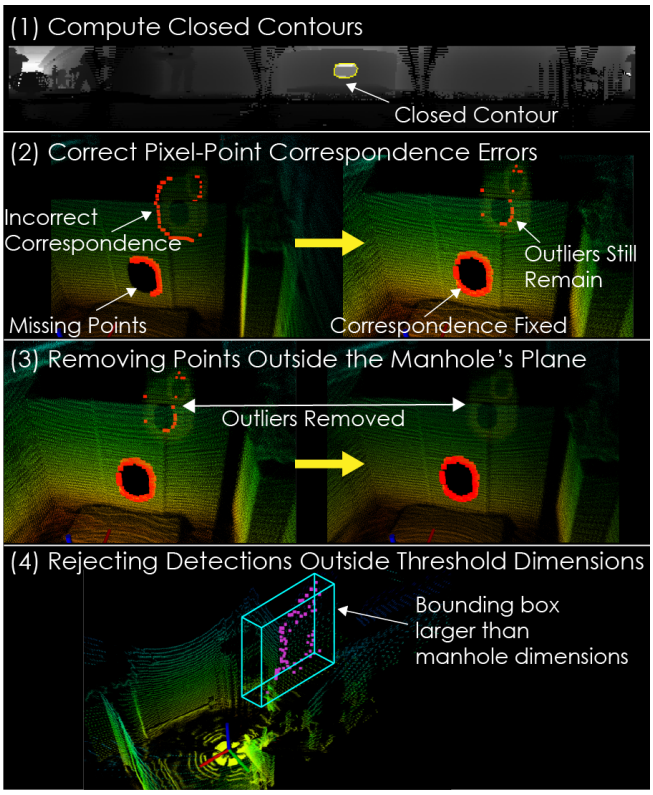

Fig. 3. Steps involved in MDL from an instantaneous point cloud and corresponding depth image. First, closed contours are detected in the depth image. The points in the point cloud corresponding to these contours are filtered to remove outliers as well as reject contours that do not correspond to a manhole using the following two properties of the manholes: a) they are planar, and b) they have standardized known dimensions.

Each detection $D_{c,t} \in \mathbb{D}_t$ is checked against the previously detected manholes. A set $\mathbb{H}$ is maintained to store these detections. Each element $H^i \in \mathbb{H}$ is a set of detections whose positions are within a distance λ_{th} from their mean position and whose orientations form an angle less than θ_{th} with the mean orientation. Every new detection $D_{c,t}$ is added to one of the sets in $\mathbb{H}$ if it satisfies the above criteria otherwise, a new set H^i is created for it. As the orientation is calculated using the normal to the plane of the manhole, the orientation of multiple detections of the same manhole might be flipped by $180°$. Hence, we transform the orientation of each new detection to lie within $-90°$ to $90°$ of the mean orientation of the set before comparing. Finally, only the manholes H^i having at least n_{min} detections are used for navigation to remove false detections.

B. Manhole Traversal

During the exploration step, the MDL module runs in the background to identify the manholes that exist in the compartment that the robot is exploring in. When the exploration of the current compartment is completed, the planner selects the closest not-yet traversed manhole to traverse to the next compartment. Two points $(\mathbf{p}_0, \mathbf{p}_1)$ on the line passing through the center of the manhole and along its orientation, both lying on the opposite sides of the manhole at a distance

ℓ from its center are selected. The planner then plans a path to the point closest among the two, and the robot follows that path to align itself in front of the manhole. The MDL step is then run again for a fixed time t_v. If within this time we get n_{min} detections, then a straight line path connecting the current location, the center of the manhole, and the selected point on the other side is commanded to the robot. After traversing this path, the corresponding manhole is marked as traversed and not used for further exploration. Then the exploration planner is triggered again and this process is repeated. If we get less than n_{min} detections, then that manhole is marked as a false manhole and the robot repeats the process for the next closest manhole. This re-detection step is done to ensure that the robot does not try to go through an area that is not a manhole. As seen from Section V-A, the highest probability of detection is when the robot is aligned in front of the manhole at close proximity. Hence, we trust the re-detection more than the original detection.

C. Exploration Planner

In this work, the exploration path planning functionality is implemented by interfacing with our previous open-sourced work on Graph-based exploration path planning (GBPlanner). This method is verified extensively in various subterranean and industrial environments and was used by Team CERBERUS in the DARPA Subterranean Challenge [5]. The method utilizes Voxblox [23] as the underlying volumetric mapping representation. It operates in a bifurcated local and global planning architecture. The local planner provides efficient paths for exploration around the robot whereas the global planner is responsible for repositioning the robot to unexplored areas of the map when the local planner cannot find an informative path, as well as for calculating a safe path for the robot to return to the home location. In this work, as the compartments of the ballast tank are relatively small, the local exploration mode is primarily used. In this mode, the planner first builds an undirected graph $\mathbb{G}_L$ along with its vertex and edge sets $\mathbb{V}, \mathbb{E}$ inside a local volume that is adaptively calculated based on surrounding environment geometry. An information gain Γ_L, called VolumetricGain, is calculated for each vertex $\nu \in \mathbb{V}$ and relates to the amount of unknown volume that will be mapped if the robot was to be at that vertex. Next, the planner calculates the shortest paths from the current robot location to all the vertices, aggregates the Γ_L of the vertices along the paths, and selects the path with the highest cumulative gain. This path is commanded to the robot, and upon execution, the above process is repeated until in one of the planning iterations, all the vertices in $\mathbb{G}_L$ have Γ_L below a threshold $\Gamma_{L,th}$. At this point, the planner declares that the local exploration is completed.

V. EVALUATION STUDIES

We verify our MDL algorithm through an ablation study to find out the effective 3D space with respect to the manhole in which reliable detection can take place. Furthermore, the entire exploration system is verified through two experiments in a mock ballast tank setup.

A. Ablation Study

The detection of the manhole primarily depends on the viewing angle of the LiDAR sensor with respect to the manhole. Hence, to evaluate the effective space in which reliable detections can happen, we conducted two studies analyzing the effect of the viewing angle on detection.

In the first study, a handheld sensing setup with an Ouster OS0-64 LiDAR was used to collect data for MDL. The sensing setup was moved in front of a stationary manhole with the viewing angle of the LiDAR with respect to the manhole varying in the range: azimuth angle: $\theta_{azi} \in [-70°, 0°]$, elevation angle: $\theta_{ele} \in [0°, 40°]$ and the distance from the manhole in the range $r \in [0.5\,\text{m}, 5.0\,\text{m}]$. Figure 4(a) shows the results of this study, whereas Figure 4(b) shows the arrangement of the LiDAR with respect to the manhole. The figure shows the percentage of successful detections as a function of the azimuth and elevation angles with respect to the manhole. The study shows that viewing angles close to zero provide the highest percentage of detections. It is highlighted that even though the data collected in this study only covers one of the octants with respect to the manhole, the symmetric nature of the manhole and the LiDAR implies that similar results will be found in the remaining octants. Furthermore, as the LiDAR is 360° its yaw is not considered in the study.

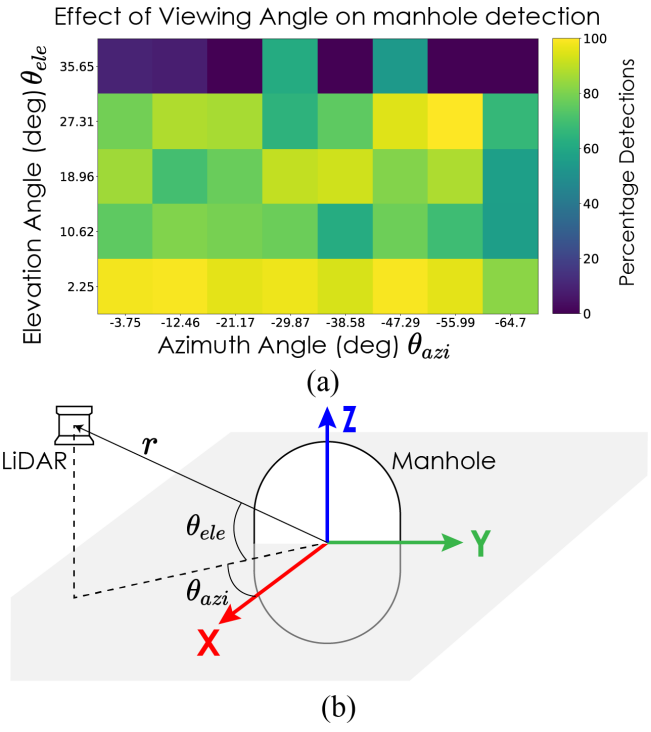

Fig. 4. Quantitative results of the first ablation study conducted using a handheld sensor setup. Sub-figure (a) shows the percentage of successful detections as a function of the azimuth and elevation viewing angles of the LiDAR sensor with respect to the manhole. It can be seen that the highest percentage of detection is at viewing angles close to zero. It is noted that as the LiDAR is 360° its yaw is not considered in the study. Sub-figure (b) shows an illustration of the arrangement of the LiDAR with respect to the manhole.

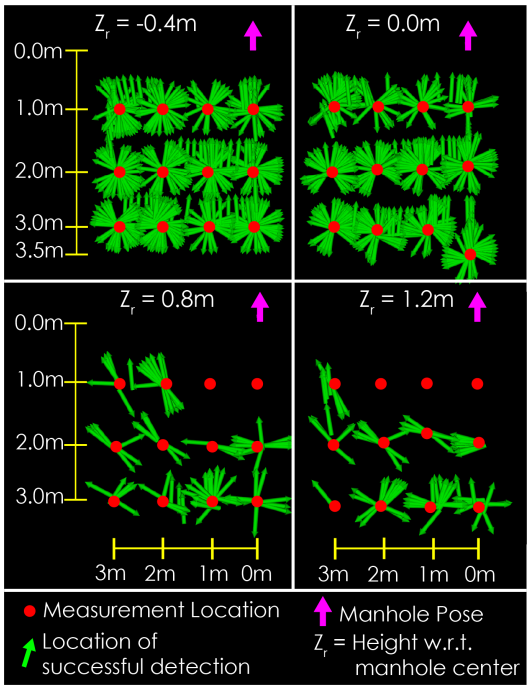

Fig. 5. Qualitative result of the second ablation study conducted using the LiDAR onboard the robot. The robot was placed at different locations in 3D space in front of the manhole and rotated around itself. The green arrows in the figure indicate the robot poses from which the manhole was successfully detected. This study helps understand the robot configuration space (relative to the manhole) in which the manhole can be detected reliably. It also shows the effects of occlusion from the robot frame to the LiDAR on MDL.

In the previous experiment, the LiDAR sensor had no occlusions in its Field of View (FoV). However, in reality, when the sensor is mounted on a robot, the frame of the robot can create occlusions to the sensor causing the detection to fail at certain robot headings and viewing angles with respect to the manhole. To study the effect of these occlusions, we conducted a study similar to the above but considering the yaw of the robot as well. The robot was placed at different points in the 3D space in front of the manhole and rotated 360° around itself. Figure 5 shows a qualitative result of the detections in this experiment. The green arrows indicate the pose of the robot where the manhole was successfully detected. This study shows the poses of the robot with respect to the manhole from which reliable detection can happen. It is important to conduct this study to select the parameters such as ℓ, of the MDL algorithm. Figure 6 shows the depth images captured on the robot during this study at the same x,y location but two different heights and robot headings to visualize what parts of the robot frame occlude the manhole at what robot poses relative to the manhole. The colored pixels in the depth images show the closed contours detected by the method. As shown in the second last image (Yaw = 0), the contour of the manhole is not detected due to occlusion.

B. Experimental Results

To verify the entire method on a real system we conducted two experimental evaluations in a lab setting creating two different configurations of compartments analogous in size to those found in common ballast tanks. Both configurations

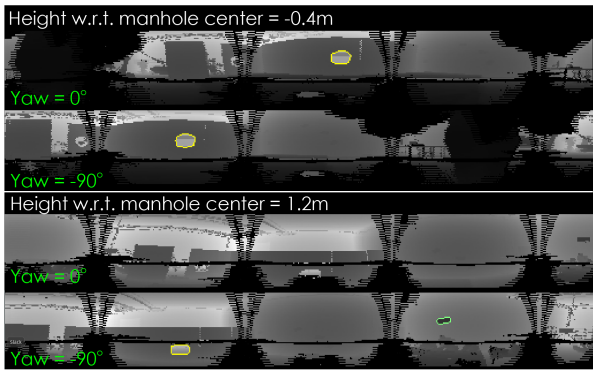

Fig. 6. Indicative depth images along with the detected contours for the same location but different robot heights and headings. The effect of occlusion from the robot frame can be clearly seen in this figure.

involved four compartments connected in different ways. The size of the manholes used in the experiments was height × width $= 0.8\,\mathrm{m} \times 0.6\,\mathrm{m}$. The results of these tests can be seen in Figure 7.

Both experiments were conducted using a custom-built collision-tolerant aerial robot called RMF-Owl [25] (with slight modifications from [25]). The robot integrates an Ouster OS0-64 3D LiDAR with 64 channels (FoV $[F_h, F_v] = [360, 90]^\circ$, maximum range $d_{\max} = 50\,\mathrm{m}$), a FLIR Blackfly S 0.4MP visual camera, and a Vectornav VN100 IMU. The path planning, SLAM, and control algorithms run onboard the robot on a Khadas VIM4 Single Board Computer incorporating $\times 4$ 2.2Ghz Cortex-A73 cores, paired with $\times 4$ 1.8Ghz Cortex-A53 cores. The methods in [36] and [37] were used for SLAM and position control respectively.

The first experiment involved the four compartments in a straight line and all the manholes falling on the same line. The dimensions of each compartment were length × width × height $= 3.0\,\mathrm{m} \times 3.6\,\mathrm{m} \times 2.5\,\mathrm{m}$. The robot started in a compartment at one end, explored each compartment, traversing through the manholes as described in Section IV-B, and finally landed in the last compartment. As seen in Figure 7, the robot was able to detect and traverse through all the manholes in the setup. For each compartment, autonomous exploration took place with the planner switching to this behavior for as long as necessary. Such a configuration of compartments and manholes is common in vessels and hence selected.

In the second experiment, the compartments were connected at right angles to each other Figure 7. The dimensions of each compartment were length × width × height $= 3.0\,\mathrm{m} \times 2.4\,\mathrm{m} \times 2.5\,\mathrm{m}$. The robot started in one of the two compartments that were not connected to each other via a manhole and explored each compartment till it reached the last one.

VI. CONCLUSIONS

In this paper a method for exploration of ballast water tanks through the explicit detection, localization and traversal of manholes using MAV is presented. The method utilizes

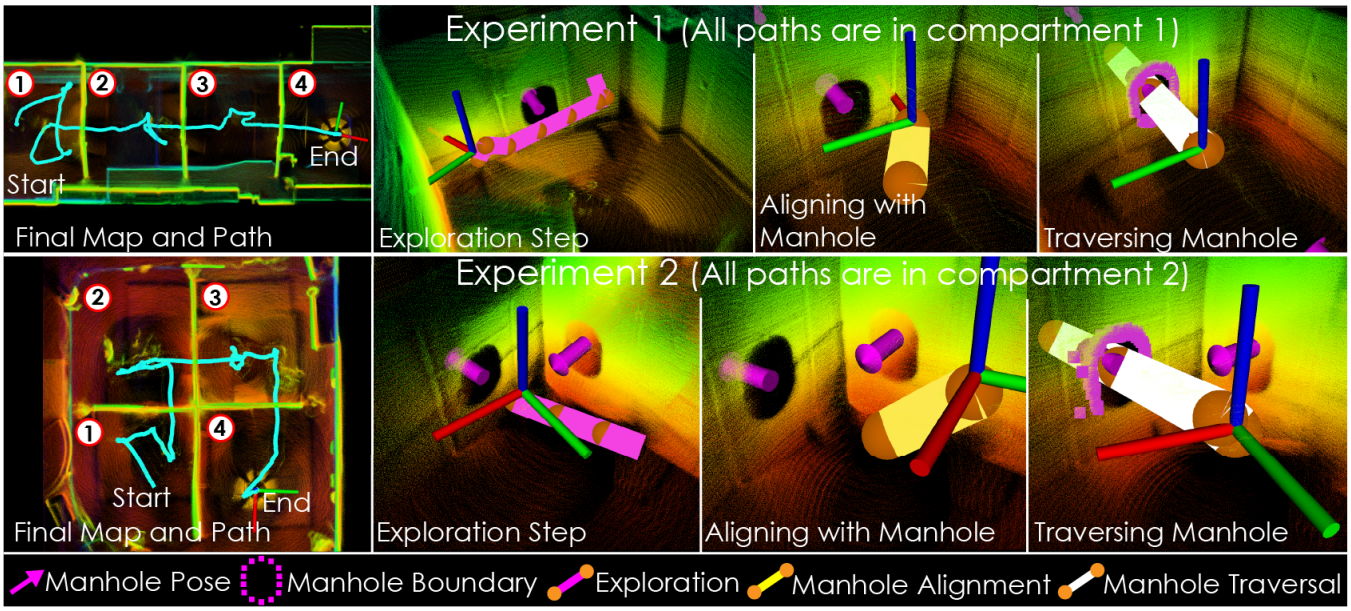

Fig. 7. Maps and paths traversed by the robot in the two experiments. The first experiment involved the four compartments in a straight line with all the manholes falling on the same line. The dimensions of each compartment were length$\times$width$\times$height $= 3.0\,\mathrm{m}\times3.6\,\mathrm{m}\times2.5\,\mathrm{m}$. In the second experiment, the compartments were connected at right angles to each other, with the dimensions of each compartment being length$\times$width$\times$height $= 3.0\,\mathrm{m}\times2.4\,\mathrm{m}\times2.5\,\mathrm{m}$. The subfigures on the left show the final map and the path traversed by the robot during the missions. The remaining subfigures show indicative paths during the exploration, manhole alignment, and manhole passing stages in one of the compartments.

a 3D LiDAR sensor for detection exploiting its instanteous point cloud. The manhole detection and localization algorithm is able to provide reliable results in a large set of robot poses relative to the manhole as verified in the presented ablation study. Furthermore, two experiments using a collision-tolerant aerial robot in mock-up scenarios demonstrate the ability of the method to run on real robots and enable them to perform exploration of multiple ballast tank compartments connected through manholes in diverse geometric configurations.

References

[1] I. Sa and P. Corke, "Vertical infrastructure inspection using a quadcopter and shared autonomy control," in *Field and service robotics*. Springer, 2014, pp. 219–232.

[2] B. Chan, H. Guan, J. Jo, and M. Blumenstein, "Towards uav-based bridge inspection systems: A review and an application perspective," *Structural Monitoring and Maintenance*, vol. 2, no. 3, pp. 283–300, 2015.

[3] A. Bircher, K. Alexis, M. Burri, P. Oettershagen, S. Omari, T. Mantel and R. Siegwart, "Structural inspection path planning via iterative viewpoint resampling with application to aerial robotics," in *IEEE International Conference on Robotics and Automation (ICRA)*, May 2015, pp. 6423–6430. [Online]. Available: https://github.com/ethz-asl/StructuralInspectionPlanner

[4] G. A. Hollinger, B. Englot, F. Hover, U. Mitra, and G. S. Sukhatme, "Uncertainty-driven view planning for underwater inspection," in *Robotics and Automation (ICRA), 2012 IEEE International Conference on*. IEEE, 2012, pp. 4884–4891.

[5] M. Tranzatto, T. Miki, M. Dharmadhikari, L. Bernreiter, M. Kulkarni, F. Mascarich, O. Andersson, S. Khattak, M. Hutter, R. Siegwart, and K. Alexis, "Cerberus in the darpa subterranean challenge," *Science Robotics*, vol. 7, no. 66, p. eabp9742, 2022.

[6] M. Tranzatto, M. Dharmadhikari, L. Bernreiter, M. Camurri, S. Khattak, F. Mascarich, P. Pfreundschuh, D. Wisth, S. Zimmermann, M. Kulkarni *et al.*, "Team cerberus wins the darpa subterranean challenge: Technical overview and lessons learned," *arXiv preprint arXiv:2207.04914*, 2022.

[7] M. Kulkarni, M. Dharmadhikari, M. Tranzatto, S. Zimmermann, V. Reijgwart, P. De Petris, H. Nguyen, N. Khedekar, C. Papachristos, L. Ott, R. Siegwart, M. Hutter, and K. Alexis, "Autonomous teamed exploration of subterranean environments using legged and aerial robots," in *2022 International Conference on Robotics and Automation (ICRA)*, 2022, pp. 3306–3313.

[8] T. Dang, M. Tranzatto, S. Khattak, F. Mascarich, K. Alexis, and M. Hutter, "Graph-based subterranean exploration path planning using aerial and legged robots," *Journal of Field Robotics*, vol. 37, no. 8, pp. 1363–1388, 2020.

[9] N. Hudson, F. Talbot, M. Cox, J. Williams, T. Hines, A. Pitt, B. Wood, D. Frousheger, K. L. Surdo, T. Molnar *et al.*, "Heterogeneous ground and air platforms, homogeneous sensing: Team csiro data61's approach to the darpa subterranean challenge," *Field Robotics*, vol. 2, pp. 595–636, 2022.

[10] A. Agha, K. Otsu, B. Morrell, D. D. Fan, R. Thakker, A. Santamaria-Navarro, S.-K. Kim, A. Bouman, X. Lei, J. Edlund *et al.*, "Nebula: Quest for robotic autonomy in challenging environments; team costar at the darpa subterranean challenge," *Field Robotics*, vol. 2, pp. 1432–1506, 2022.

[11] T. Rouček, M. Pecka, P. Čížek, T. Petříček, J. Bayer, V. Šalanský, D. Heřt, M. Petrlík, T. Báča, V. Spurný *et al.*, "Darpa subterranean challenge: Multi-robotic exploration of underground environments," in *International Conference on Modelling and Simulation for Autonomous Systems*. Springer, Cham, 2019, pp. 274–290.

[12] S. Scherer, V. Agrawal, G. Best, C. Cao, K. Cujic, R. Darnley *et al.*, "Resilient and modular subterranean exploration with a team of roving and flying robots," *Field Robotics*, 2021.

[13] GeoOceans, "Rov marine class inspection." [Online]. Available: https://www.geooceans.com/services/rov-class-inspection/

[14] F. AS, "Aerial platform for irregular confined spaces." [Online]. Available: https://www.robins-project.eu/fly-drone/

[15] ScoutDI, "Maritime vessel inspection using drones." [Online]. Available: https://www.scoutdi.com/applications/#maritime

[16] R. Y. Brogaard, R. E. Andersen, L. Kovac, M. Zajaczkowski, and E. Boukas, "Towards an autonomous, visual inspection-aware 3d exploration and mapping system for water ballast tanks of marine vessels," in *2021 IEEE International Conference on Imaging Systems and Techniques (IST)*, 2021, pp. 1–6.

[17] L. Christensen, "Cost-effective autonomous robots for ballast water tank inspection," *Journal of ship production and design*, vol. 27, 2011.

[18] F. Bonnin-Pascual and A. Ortiz, "On the use of robots and vision technologies for the inspection of vessels: A survey on

recent advances," *Ocean Engineering*, vol. 190, p. 106420, 2019. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0029801819305682

[19] R. Andersen, L. Nalpantidis, O. Ravn, and E. Boukas, "Investigating deep learning architectures towards autonomous inspection for marine classification," in *2020 IEEE International Symposium on Safety, Security, and Rescue Robotics (SSRR)*, 2020, pp. 197–204.

[20] K. Ishizu, N. Sakagami, K. Ishimaru, M. Shibata, H. Onishi, S. Murakami, and S. Kawamura, "Ship hull inspection using a small underwater robot with a mechanical contact mechanism," in *2012 Oceans - Yeosu*, 2012, pp. 1–6.

[21] *Application of Mini-ROV Technology for FSO Ballast Tank Inspection and Thickness Gauging*, ser. SPE Asia Pacific Oil and Gas Conference and Exhibition, 10 2015, sPE-176049-MS. [Online]. Available: https://doi.org/10.2118/176049-MS

[22] F. Bonnin-Pascual, A. Ortiz, E. Garcia-Fidalgo, and J. P. Company-Corcoles, "A reconfigurable framework to turn a MAV into an effective tool for vessel inspection," *Robotics and Computer-Integrated Manufacturing*, vol. 56, no. April, pp. 191–211, Apr. 2019. [Online]. Available: https://doi.org/10.5281/zenodo.4408239

[23] 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 International Conference on Intelligent Robots and Systems (IROS)*, 2017.

[24] A. Hornung, K. M. Wurm, M. Bennewitz, C. Stachniss, and W. Burgard, "OctoMap: An efficient probabilistic 3D mapping framework based on octrees," *Autonomous Robots*, 2013.

[25] P. D. Petris, H. Nguyen, M. Dharmadhikari, M. Kulkarni, N. Khedekar, F. Mascarich, and K. Alexis, "Rmf-owl: A collision-tolerant flying robot for autonomous subterranean exploration," in *2022 International Conference on Unmanned Aircraft Systems (ICUAS)*, 2022, pp. 536–543.

[26] R. E. Andersen, M. Zajaczkowski, H. Jaiswal, J. Xu, W. Fan, and E. Boukas, "Depth-based deep learning for manhole detection in uav navigation," in *2022 IEEE International Conference on Imaging Systems and Techniques (IST)*, 2022, pp. 1–6.

[27] P. D. Petris, H. Nguyen, M. Kulkarni, F. Mascarich, and K. Alexis, "Resilient collision-tolerant navigation in confined environments," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*, 2021, pp. 2286–2292.

[28] P. De Petris, H. Nguyen, T. Dang, F. Mascarich, and K. Alexis, "Collision-tolerant autonomous navigation through manhole-sized confined environments," in *2020 IEEE International Symposium on Safety, Security, and Rescue Robotics (SSRR)*, 2020, pp. 84–89.

[29] V. Pritzl, P. Stepan, and M. Saska, "Autonomous flying into buildings in a firefighting scenario," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*, 2021, pp. 239–245.

[30] A. R. Castaño, H. Romero, J. Capitán, J. L. Andrade, and A. Ollero, "Development of a semi-autonomous aerial vehicle for sewerage inspection," in *Robot 2019: Fourth Iberian Robotics Conference*, M. F. Silva, J. Luís Lima, L. P. Reis, A. Sanfeliu, and D. Tardioli, Eds. Cham: Springer International Publishing, 2020, pp. 75–86.

[31] D. Falanga, E. Mueggler, M. Faessler, and D. Scaramuzza, "Aggressive quadrotor flight through narrow gaps with onboard sensing and computing using active vision," in *2017 IEEE International Conference on Robotics and Automation (ICRA)*, 2017, pp. 5774–5781.

[32] N. J. Sanket, C. D. Singh, K. Ganguly, C. Fermüller, and Y. Aloimonos, "Gapflyt: Active vision based minimalist structure-less gap detection for quadrotor flight," *IEEE Robotics and Automation Letters*, vol. 3, no. 4, pp. 2799–2806, 2018.

[33] ISO, "Ships and marine technology - manholes with bolted covers," 2018. [Online]. Available: https://www.sis.se/api/document/preview/80008224

[34] ——, "Ships and marine technology - manholes with bolted covers," 1999. [Online]. Available: https://cdn.standards.iteh.ai/samples/26573/687db42f4e9a415a9cf8c88c75770cd1/ISO-5894-1999.pdf

[35] J. Canny, "A computational approach to edge detection," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. PAMI-8, no. 6, pp. 679–698, 1986.

[36] S. Khattak, H. Nguyen, F. Mascarich, T. Dang, and K. Alexis, "Complementary multi–modal sensor fusion for resilient robot pose estimation in subterranean environments," in *2020 International Conference on Unmanned Aircraft Systems (ICUAS)*, 2020, pp. 1024–1029.

[37] M. Kamel, T. Stastny, K. Alexis, and R. Siegwart, "Model predictive control for trajectory tracking of unmanned aerial vehicles using ros," *Springer Book on Robot Operating System (ROS)*, pp. 3–39, 2017.