The MAX Drone for Autonomous Indoor Exploration

Gustav Tolt
FOI[1], Sweden
gustav.tolt@foi.se

Joakim Rydell
FOI, Sweden
joakim.rydell@foi.se

Michael Tulldahl
FOI, Sweden
michael.tulldahl@foi.se

Max Holmberg
FOI, Sweden
max.holmberg@foi.se

Oskar Karlsson
FOI, Sweden
oskar.karlsson@foi.se

Fredrik Bissmarck
FOI, Sweden
fredrik.bissmarck@foi.se

ABSTRACT

This paper presents the concept and prototype implementation of a drone for Multi-purpose Autonomous eXploration of indoor environments – MAX. The purpose of MAX is to support first responders in the difficult task of assessing unknown and potentially dangerous or hostile situations in indoor or underground environments. The approach for addressing challenges associated with this task has been to construct a custom-designed drone based on requirements and conditions of first responder missions. This paper reports on the first phase of development of the MAX drone, aimed for experimentation with autonomy functionality in first responder contexts and for enabling further development of advanced higher-level planning functions. It describes the overall design of the MAX drone, its capabilities in terms of robust positioning and autonomous mission execution, along with the status of key enabling algorithms for exploration, such as target point selection and path planning.

Keywords

UAV, exploration, navigation, positioning, autonomy.

INTRODUCTION

Natural and human-made disasters cause the deaths of thousands of people every year. Earthquakes, floodings, fires and terrorist attacks lead to substantial damage to infrastructure and leave many buildings semi-collapsed or unstable. First responders face serious challenges in search and rescue (SaR) operations in the aftermath of such events. Particularly dangerous activities are those who take place inside damaged or otherwise dangerous buildings in which victims may be trapped. Uncertainty, inadequacy or absence of information increase the risk of making decisions that may have grave consequences, jeopardize the health of first responders, or delay the intervention needed to extract trapped victims or to save substantial material assets.

In order to save more lives, first responders need tools that can enable reliable assessment of the interior of buildings. Robotic systems have an important role to play in such situations, and the use of unmanned vehicles (UxV) in first responder operations is increasing rapidly. Unmanned vehicles can significantly enhance the capabilities to bring sensors into hazardous or risky areas, assess the status of buildings, detect victims or to provide important clues about how a scenario is developing.

Manual piloting of drones in unknown buildings is challenging. The risk of collision is high, line of sight to the drone is frequently blocked and wireless communication with the drone may become more unreliable as the drone penetrates the building. As a response, many research initiatives focus on using of unmanned aerial vehicles

[1] FOI – Swedish Defence Research Agency

(UAV) for SaR in complex environments and address issues relating to control automation, to reduce the dependence of human operators for piloting drones, and of the availability of stable communication links.

Several approaches (e.g. Shen et al., 2012; Ravankar et al., 2018; Martínez Novo et al., 2022) show that efficient exploration and accurate mapping of unknown indoor environments is becoming feasible even with low-cost vehicles and sensors, albeit featuring less computationally demanding planning algorithms. One especially strong factor contributing to recent developments is the Robot Operating System (ROS) that provides a range of key functions for simplified sensor and system integration as well as for tackling the exploration problem (Papachristos et al., 2018). Many approaches to the exploration problem exploit the concept of frontiers, as proposed by Yamauchi (1997). Recent frontier-based approaches have been proposed by Hardouin et al. (2020), Brunel et al. (2021) and Krátký et al. (2021). The latter also provides a comparison of several planning approaches for UAVs. Williams et al. (2020) propose a method for 3D frontier-based exploration that do not use an explicit volumetric representation as an intermediate map and demonstrate the technique in a subterranean environment. Other recent exploration approaches demonstrated in underground environments include Alexis (2019) and Lindqvist et al., (2022).

How to best explore a building depends on the current situation. Time constraints, risk levels, a priori information, e.g. regarding the presence of humans inside the building and other factors may vary greatly from one situation to another. The problem of combining several aspects to decide the robot's optimal movement falls within the family of Multi-Criteria Decision-Making (MCDM). Zagradjanin et al. (2022) describe and compare different approaches related to the problem of finding next best view points for efficient exploration. Yu and Wang (2019) use a robot with depth imaging sensors and combine a partially observable Markov decision process (POMDP) with partial map simulation to plan the motion of the robot. Wang and Englot (2020) suggest taking the expected uncertainty of unknown navigation landmarks into account in the multi-criteria utility function. To target the difficulties of hand-crafting efficient and robust rules for exploration, approaches leveraging advances in machine learning techniques have been proposed. Chen et al. (2020) suggest a method based on Graph Neural Networks and Reinforcement Learning to give the robot the capability to learn an efficient exploration behavior.

Common for all approaches aimed for exploration of GNSS-denied and unknown environments is the need for simultaneous localization and mapping (SLAM) algorithms using data from one or several sensors onboard the platform. Recent work targeting the problem of achieving accurate and real-time SLAM with cameras onboard small drones includes Karam et al. (2022). Gupta and Fernando (2022) give a comprehensive overview of SLAM for drones. Chen et al. (2017) show that a map created through SLAM using a single monocular camera on a ground robot may produce valuable information for first responders.

Research activities aimed at advancing the use of drones in SaR applications are pursued in several EU-funded projects, which together address a broad spectrum of aspects related to drones, including their purpose, design, sensor payload and integration with other tools. One such project is INGENIOUS that aims at developing a toolkit of integrated devices and services that augment first responders' situational awareness in all stages of an operation. In INGENIOUS, the *Multi-purpose Autonomous eXploration* (MAX) concept is proposed, that combines *autonomous exploration* functionalities with *multi-sensor data acquisition* as a means to relieve first responders of the difficult task to assess complex indoor environments. In INGENIOUS, interpretation of explored scenes is enhanced through AI-based image understanding techniques applied to data acquired with the MAX drone (Zhang et al., 2022a; Zhang et al., 2022b). The ambition with MAX is to embed key capabilities – SLAM, exploration and autonomous navigation – in a UAV equipped with a multitude of sensors to provide first responders with better means of quickly and reliably assessing situations in future operations. This Work in Progress paper focuses on the autonomous platform aspects of the MAX drone rather than on the use of its sensor payload for specific applications.

OBJECTIVE AND CHALLENGES

The main objective of the MAX concept is to *support first responders in assessing complex indoor spaces, in semi-collapsed, dangerous and GNSS-denied buildings*. Fulfilling this objective implies addressing a number of challenges:

- **Positioning infrastructure not present.** Satellite-based positioning systems work poorly indoors, and pre-installed reference markers or navigations beacons in the building cannot be expected. Hence, MAX needs to be able to estimate its position relative to objects in the environment.
- **Many buildings are unknown.** The existence of building maps cannot be assumed. Therefore, MAX needs to create its own necessary representation of the scene in which it is deployed.
- **Environments are complex and diverse.** A variety in dimensions and layouts of buildings pose different conditions regarding the possibility for drones to estimate their position and to plan a collision-

free path through the environment. Therefore, MAX needs functions that help it navigate safely in a variety of geometries.

- **Unstable communications.** Uninterrupted communication cannot always be guaranteed and to fulfil its exploration mission even in the case of intermittent communication dropouts with a ground station outside the building, MAX must carry everything it needs to navigate autonomously through the building.
- **Varying information needs.** The information needed to perform building assessment may vary greatly between situations, from estimating the stability of supporting columns to detecting victims or sensing the presence of toxic gases. Responsibilities and expertise vary between first responders and their attention is focused on different phases of the event. Therefore, the drone should support different tasks and be able to provide a multitude of sensor data.

Table 1 summarizes how the MAX concept addresses the above challenges from a technical point of view.

Table 1. Challenges addressed by the MAX concept.

Challenge	MAX approach
Positioning infrastructure not present.	Real-time SLAM (simultaneous localization and mapping).
Many buildings are unknown.	Update a 3D occupancy grid map in real-time.
Environments are complex and diverse.	Dual SLAM systems to increase positioning robustness, obstacle detection and avoidance capabilities.
Unstable communications.	On-board computations of everything it needs to operate safely.
Varying information needs.	Equip the drone with several sensors.

THE MAX DRONE

Here we describe the approach taken to address the challenges discussed above and arrive at a first implementation of a MAX drone prototype (Figure 1). We describe the concept, the physical design and key characteristics, as well as the main components installed.

The purpose of MAX is to act as an extra teammate that first responders can send in to scout a potentially dangerous building. The scenario for which MAX is primarily intended is the assessment of an unstable, semi-collapsed building likely to have sustained enough strain to severely affect its stability. The building is assumed to offer at least one still usable entry point (door, window).

Figure 1. The MAX drone 2023.

The MAX drone is designed to be flexible and modular in terms of payload and components, to leave room for future extensions and component upgrades. The design is custom-made, based around a set of sensors needed to achieve exploration and mapping and utility sensors for scene assessment. The drone body has integrated propeller guards to increase safety both for surroundings and for the drone itself. The body consists of strong and lightweight carbon fiber-reinforced sandwich plate. A rotating, round-looking 3D-lidar sits on top of the body. The battery and utility sensors, such as visual and thermal cameras and environment sensors (temperature, gas indicators), are located in enclosures beneath the body. The body plate has milled-out housings for core components such as the motor speed controllers, controller for the lidar, onboard computers, stereo camera, and proximity sensing sonars.

To obtain a small drone footprint, the propellers are small and rotate fast. Despite the large amount of sensors and parts, the total length and width measure only 43 cm × 43 cm, enabling operation in relatively narrow indoor passages and spaces. With the current equipment and because of the small propellers the flight time is limited to about five minutes.

MAX can operate in three modes: *manual*, *automatic* and *autonomous exploration* mode. In automatic mode, MAX either follows a set of predefined waypoints, or – if only one waypoint is set – hovers to maintain this position. This is useful for a patrolling or gazing behavior. In autonomous mode, MAX explores a region of interest (ROI), the coordinates of which are specified relative to MAX's starting position. The mode is selected by the operator by flipping a switch on the radio control unit. Operator involvement is needed for safety reasons, and also increases the range of environments and operating conditions where MAX can be practically used. MAX currently requires manual control at take-off and landing but planned future enhancements include the ability to automatize these steps as well.

To enable exploration in unknown structures, MAX has a set of *navigation sensors* (lidar, depth cameras, IMU and sonars) along with algorithms that give the ability to navigate safely and efficiently through an unknown environment. All algorithms are embedded in an on-board computer system that allows continued operation even in the case of an interrupted communication link. It also has *utility sensors* (visual camera, thermal camera, temperature sensor, indicator of dangerous gases) that send data that can analyzed by first responders or by computerized analysis tools (e.g. object recognition). The anticipated relevance of the MAX concept in first responder missions is embodied by the drone's wide range of sensors allowing configurations enabling planning and executing its motion according to mission-specific criteria. Currently, only one basic type of exploration is implemented, namely exploration in the sense of 3D mapping with the laser scanner.

MAX has two on-board computers – a Raspberry Pi4 and an Nvidia Jetson Nano – that share data through an Ethernet connection. The software integration is based around ROS. The MAX drone communicates with a service laptop that shows status data from the drone, stores and visualizes sensor data and flight parameters. The laptop also serves as a communication node at which georeferenced sensor data, 3D maps and status data can be packaged or adapted and sent to other nodes, depending on the application.

ROBUST POSITIONING

A key component of an autonomous vehicle is the ability to accurately estimate its position, orientation and velocity. In order to do so robustly, the MAX drone has two separate SLAM-based positioning systems with independent sensors – one visual SLAM system based on the Intel RealSense D435i stereo camera and one lidar-SLAM system using an Ouster OS1-16 lidar. The two independent position estimates are fused together to get a single, more robust estimate. The two systems complement each other in that the visual SLAM system excels in small or cluttered areas but lacks the range in larger environments, whereas the lidar system works better in larger environments but struggles at very short ranges and in very cluttered environments.

Visual SLAM

The visual SLAM system detects SIFT interest points (Lowe, 2004) in the left and right camera views in the stereo camera. Points that are successfully matched between the left and right image are tracked through a sequence of images from the stereo camera. By tracking a number of such points, commonly referred to as landmarks, the motion of the camera relative to the surrounding environment can be estimated. Using information from both the left and right camera views to track landmarks enables estimation of their distances from the camera and removes the scale ambiguity which is present in monocular visual SLAM.

In order to facilitate the tracking of landmarks during rapid motion, while also rejecting landmarks in non-stationary parts of the images (e.g. people), measurements from an IMU (inertial measurement unit, consisting of a 3-axis accelerometer and a 3-axis gyroscope) are used to predict where tracked landmarks are expected to appear in images from the camera. For stationary landmarks, this makes association between observed and tracked landmarks easier. For non-stationary landmarks, the prediction is generally far from where the landmarks actually appear. This makes it easy to discard the non-stationary landmarks.

The fusion of inertial data and depth images is performed in an extended Kalman filter (EKF). More detailed descriptions of the visual SLAM system can be found in (Rydell and Bilock, 2015; Rydell et al., 2016). The visual SLAM system runs on the Nvidia Jetson Nano, since it enables interest point detection to run on the GPU. This considerably improves the computational performance of the implementation.

Lidar SLAM

The lidar SLAM used on the MAX drone is a feature-based algorithm combining geometrical features, inertial measurements and lidar odometry in a factor graph. The algorithm is developed for execution on a compact computer; on the current MAX drone prototype a Raspberry Pi 4 is used. The limited hardware imposes significant constrains on the complexity of the SLAM algorithm, which is why a feature-based system was chosen. The features used are planes and egdes and the feature extraction algorithm is based on the method presented in (Zhang and Singh, 2014). The factor graph method uses the GT-SAM graph library first presented in (Dellaert, 2012). A unique aspect of this lidar SLAM algorithm is the use of geometrical primitives directly in the factor graph. Similar algorithms often feed a separate lidar odometry estimate to the graph instead. Using the features as primitives directly in the graph is advantageous as all available information is incorporated in one global optimization instead of multiple separate optimizations that use only parts of the available information. The algorithm and its performance is described in more detail in (Holmberg et al., 2022).

SLAM Fusion

The two position estimates provided by the two SLAM systems need to be combined in order to control the MAX drone. This is a difficult task that requires a dynamic way of knowing when and to what degree the separate systems can be trusted. On the MAX drone this is achieved through the use of a factor graph based solution, the structure of which is shown in Figure 2. The pose increments of the individual position estimates are fed to the graph as factors between the "fusion poses". When available, the loop closing measurements are also included in the graph. Since the two SLAM systems are working at different rates, the incoming data need to be interpolated to a common rate. In Figure 3, the estimates from the two different SLAM systems and the fused trajectory are shown for two test environments. The experience so far from using the MAX drone is that the SLAM fusion increases the overall accuracy and robustness, but no quantitative evaluation has yet been performed.

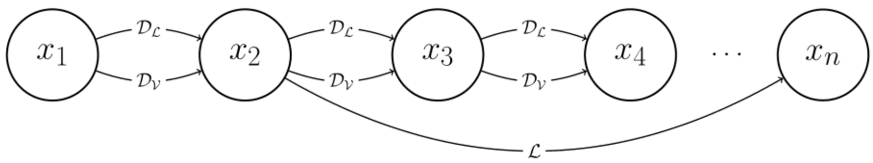

Figure 2. Example of the graph-structure. x_i are pose values, D_L are the pose delta measured by the lidar SLAM, D_V are the pose delta measured by the visual-SLAM and L are loop closing factors.

Loop closure

Without a stable reference, a SLAM system based on relative measurements will drift over time. One way of mitigating such drift is to revisit previously mapped areas and calculate an accurate transform between the previous pose and the current pose, so called *loop-closing*. Due to the graph based slam-fusion algorithm (described in the previous section) a loop-closing algorithm can seamlessly be integrated to improve the position estimate. The loop-closing algorithm consists of the following steps:

1. Identify possible loop-closing opportunities by considering the age and distance of all previous poses compared to our current pose.
2. Try to match the lidar-scan from the suggested loop-closing opportunity pose with the current pose lidar-scan using the Iterative Closest Point (ICP) algorithm for 3D point cloud registration.
3. If a sufficiently good match is found, pass the transform to the SLAM-fusion algorithm.

The loop closing is currently performed off-line as a post-processing step to provide as accurate 3D maps as possible to the first responders, but since it is computationally tractable the plan is to integrate it on the MAX platform. Figure 4 shows a map and trajectory with and without loop-closing. In this particular example, closing the loop brought the final error down to below 10 cm. The process is described in detail in Holmberg et al. (2022) along with some empirical results.

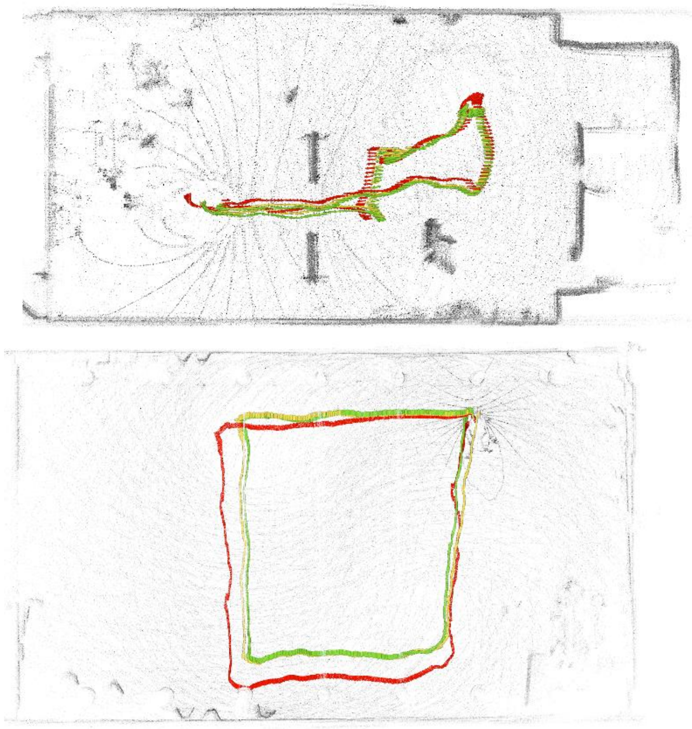

Figure 3. 2D visualizations of lidar point cloud map and trajectories estimated by the lidar SLAM (green), visual SLAM (red) and fuse pose estimate (yellow), respectively, in a smaller room, around 10 m x 20 m (top), and a larger hall, around 65 m x 40 m (bottom).

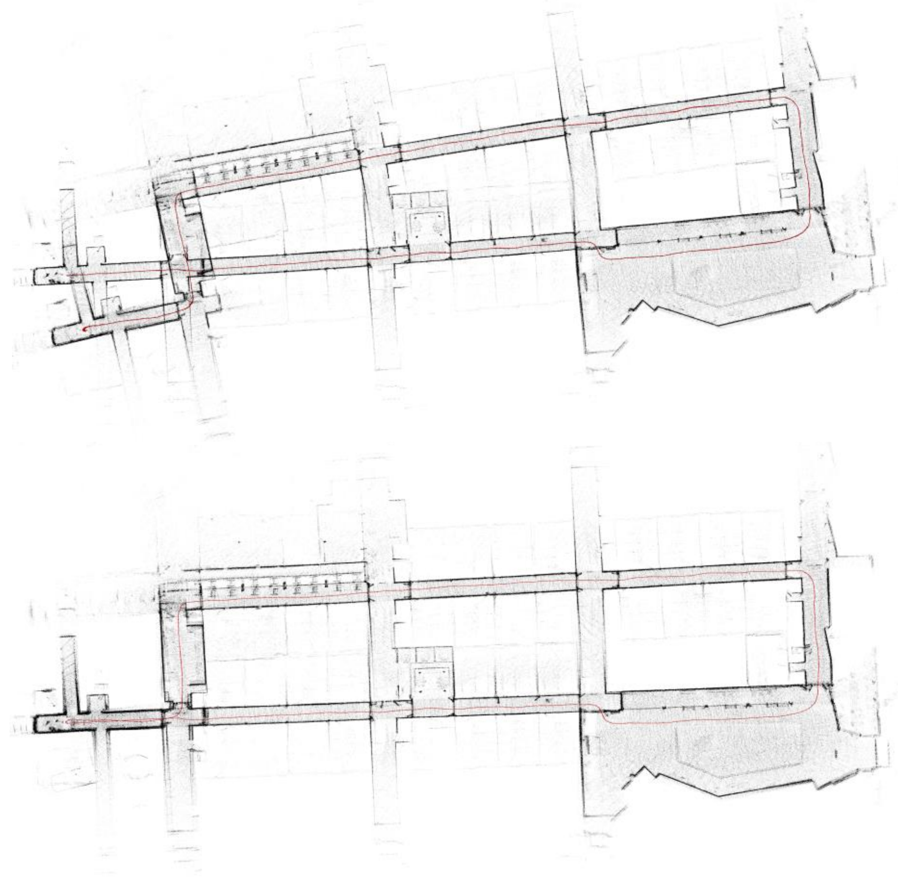

Figure 4. Top: Map and position trajectory(red line) generated by the lidar-SLAM, without loop closure. Bottom: Result after loop closure between the start and end of the trajectory.

EXPLORATION BEHAVIOR

The exploration uses the established concept of 3D Occupancy Grid Maps (OGM), implemented as Octomaps (Hornung et al., 2013) built and analyzed in real-time. The OGM allow describing every part of the environment as *occupied* (blocked), *free* (navigable) or *unknown* (not yet observed).

The target point selection mechanism currently implemented on the MAX drone is rather simple and inspired by the frontier-based approach suggested in (Yamauchi, 1997). Frontiers represent virtual boundary surfaces between known and unknown space. The concept is very useful for assessing and visualizing how the exploration progresses; If no more frontiers exist, exploration is complete. The process of choosing a target point and a path to get there is as follows:

1. Update the OGM using 3D lidar data and current platform pose (position and orientation)
2. Identify frontiers and remove too small/spurious frontiers.
3. Randomly select a frontier and estimate its local direction.
4. Query a 3D point P at a fixed distance perpendicular from the estimated frontier direction.
    a. If P lies in a free voxel and there is a collision-free path to P from the drone's current position, then accept P as next target point.
    b. If P does not qualify as target point, goto step 4 and repeat until a new target point is found or no more frontiers exist.

The path is computed using the Open Motion Planning Library (OMPL[2]).

One limitation of the current method is that chosen target points may not be ideal from a movement efficiency perspective. Further development will concern enhancements of this process by considering different quality factors such as the possibility to maintain accurate position estimates while moving, prioritize target points based on distances from already visited points, or to direct cameras towards certain parts of the scene. Learning-based exploration techniques will also be considered. Evaluation factors for exploration could include requirements on estimated quality of utility sensor data. One example would be to require visual image data acquired on all walls (occupied parts) at a suitable camera-to-wall distance interval to ensure a base-line quality level of image data and limit the exploration time of a large building.

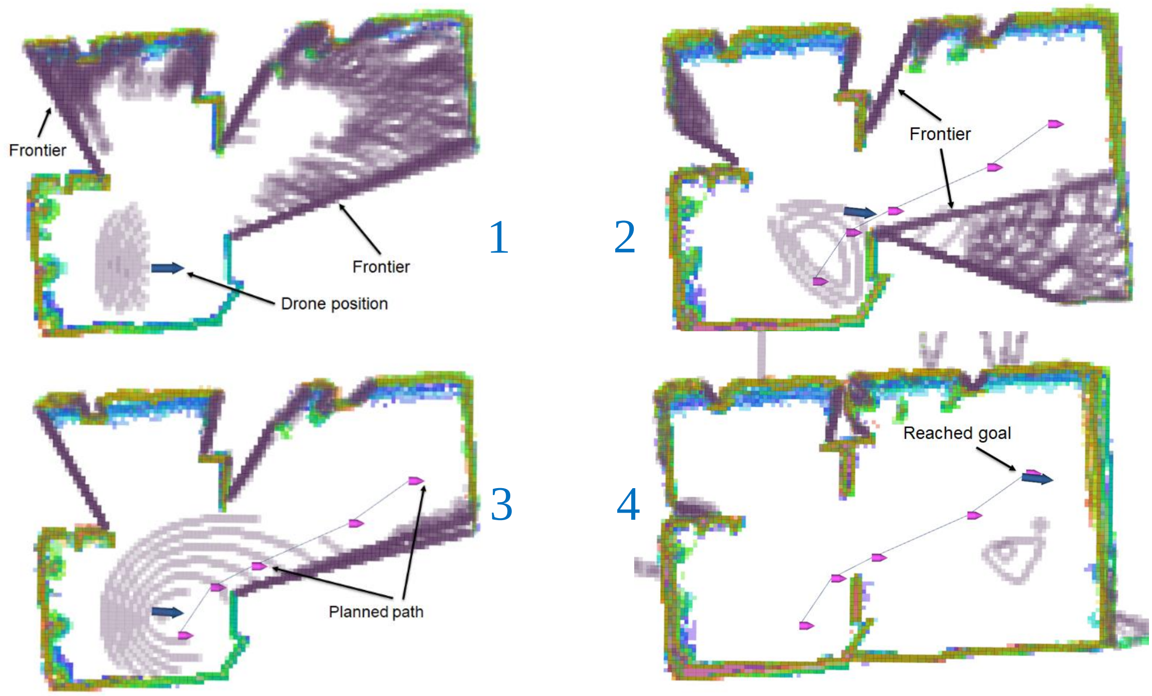

Figure 5. Illustrations of occupancy grid map (OGM) and the target point selection and path planning during a flight. In image 1, there are large frontiers indicating that large parts of the room are still unexplored. In image 4, MAX has followed a path to the target point and updated the OGM with sensor data, drastically reducing the frontier count and the still unexplored part of the room.

[2] ompl.kavrakilab.org

NAVIGATION

MAX navigates to a target point by following the path proposed by the planner. To do so, MAX tries to reach a successive list of waypoints and makes quick adjustments to the path in-flight to account for imperfections in the OGM used by the planner and to avoid suddenly appearing local obstacles.

Path Following

The lowest-level flight control onboard MAX is handled by the open source flight controller software ArduCopter, running on a Pixhawk Mini. This controller is responsible for keeping MAX stable in the air. ArduCopter is also capable of hovering in place, flying at a predefined velocity or even following a list of waypoints. However, that requires ArduCopter to know the actual position and velocity of the drone. This information is usually provided to the flight controller from an onboard GNSS receiver, but that is not possible when operating MAX indoors. Hence, MAX is controlled by sending tilt angles and vertical velocities to the flight controller. ArduCopter converts these into thrust settings for the four motors.

The desired tilt angles and vertical velocities are computed based on a target waypoint and the current position and velocity, as computed by the positioning systems described earlier in this paper. The target waypoint comes from the exploration module in autonomous exploration mode or from a list of pre-defined waypoints in automatic mode. A target velocity is computed based on the distance between the current and desired position. This velocity is compared to the current velocity, and tilt angels in the roll and pitch direction are computed based on the acceleration needed to achieve the desired velocity.

When MAX is sufficiently close to the target waypoint, its behavior varies depending on the current mode. In autonomous exploration mode, a new waypoint is requested from the exploration module. In automatic mode, the next predefined waypoint is loaded (if there are no more waypoints, MAX hovers in place). In INGENIOUS, the choice has been to give MAX a floor-wise exploration behavior as this how first responders often explore buildings. However, the 2D behavior is not inherent in the positioning, mapping or path planning algorithms but only a result of how waypoint candidates are selected in the very last step; a more or less horizontal movement is obtained by discarding waypoints close to frontiers that would imply moving or looking upwards. This makes the core algorithms applicable to other platforms too.

Obstacle avoidance

To avoid collisions with suddenly appearing objects in the planned path, MAX detects and localizes obstacles using its onboard lidar and modifies its route to try to go around them. Obstacles are detected as any objects observed by the lidar, at an altitude where MAX may collide with them and within two meters from MAX. Obstacles are avoided by modifying the desired velocity by minimizing a simple cost function c, which is a sum of two terms:

$$c(\varphi) = a\, o(\varphi) + b\, |\varphi - \varphi_0| \,,$$

where $o(\varphi)$ is based on of how close to an obstacle MAX will get by flying in the direction φ – φ| ,φ_0| is the angular difference between φ and the direction of the desired velocity φ_0, and a and b are weighting parameters controlling the trade-off between the two terms. The desired velocity is modified in such a way that MAX will travel in the direction of minimum cost. This is essentially the direction closest to φ_0 that does not risk causing a collision. If the updated direction of travel does not bring MAX closer to the target waypoint, a new waypoint is requested from the exploration module (in autonomous exploration mode).

Figure 6 shows this algorithm in action. MAX starts at position A, with a pre-defined list of waypoints (shown by the dashed line): First go to B, then to C, and finally back to A. The solid black curve shows the actual motion of MAX, while gray dots represent obstacles. The blue arrows show where MAX would go without collision avoidance, i.e., direct to the next waypoint. The red arrows show the result of the collision avoidance algorithm, i.e., the direction in which MAX should accelerate in order to avoid colliding with the obstacle.

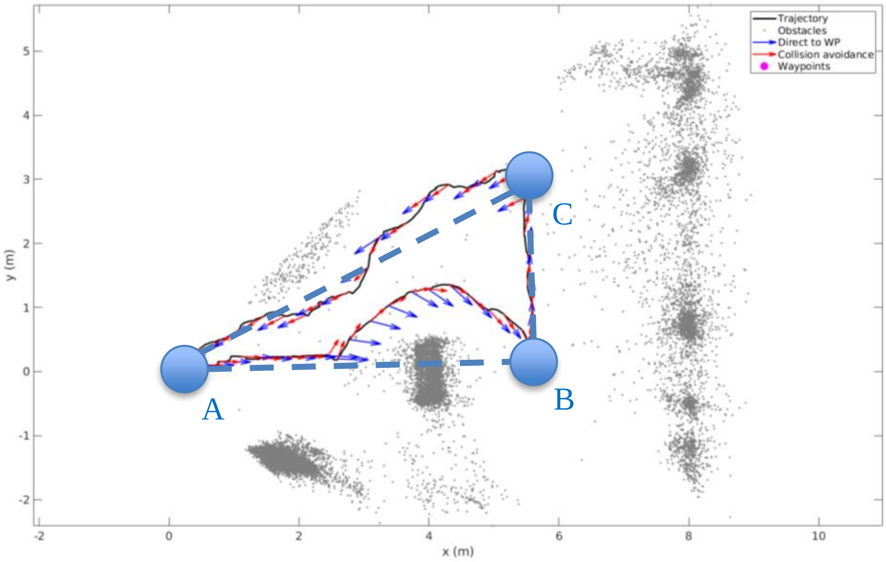

Figure 6. Collision avoidance. Position A marks the starting point, and positions B and C shows are waypoints. Note the obstacle detected between A and B that forces MAX to re-plan a path around it. The solid black line shows the route taken by MAX and the arrows show where the drone would go without collision avoidance (blue) and where it will go with collision avoidance (red).

Fly-through-opening algorithm

To pass safely through even narrow openings, MAX has a special fly-through-opening algorithm that is activated whenever there are nearby obstacles on both sides of the planned path. Sensors such as lidar and sonar often have a practical minimum operating distance. If the distance to a nearby object is closer than this distance, data from those sensors need to be ignored while the fly-through-opening routine is being executed. The routine consists of the following steps carried out in sequence (colors refer to Figure 7):

1. Detect and locate the opening including its orientation based on the obstacle map.
2. Compute a coarse starting point (red) in front of the opening, based on the current obstacle map
3. Fly to the starting position.
4. Find a refined position (green) at a predetermined distance from the opening (1.6 m).
5. Fly to the refined position and orient MAX with equal distances to perceived obstacles, perpendicular to the opening (so that a forward movement would take MAX through the opening)
6. Compute a goal position (blue) on the other side of the opening with clearance from obstacles.
7. Obtain a stable hover at the refined position by maintaining pose until the velocity is small (< 5 cm/s).
8. Fly to the goal position.

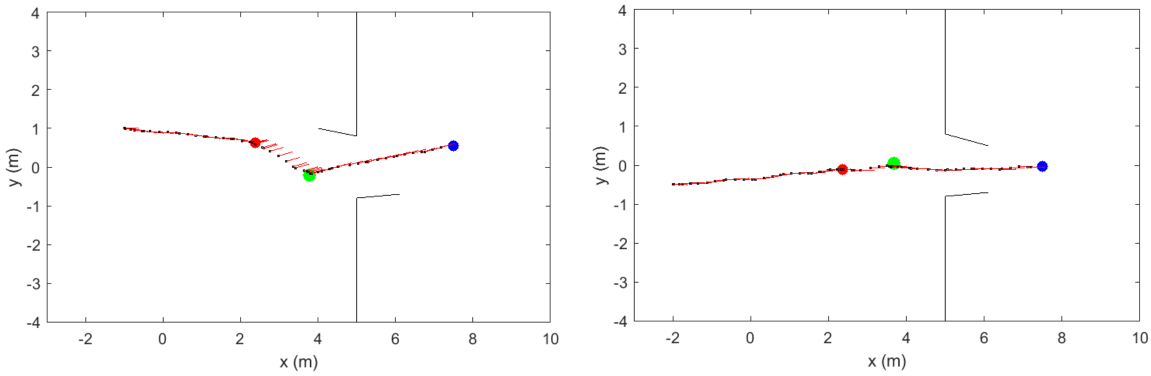

Figure 7. Illustrations of two simulated fly-through-opening situations. See text for details.

CONCLUSIONS AND FUTURE WORK

MAX is a custom-built drone to serve as a modifiable platform for development, test and evaluation of novel support tools for first responders. This paper has presented the overall MAX concept, the current state of the MAX drone prototype and its capabilities. The goal up to now has been to integrate all necessary components and functions into a system that enables autonomous exploration of an unknown indoor environment, provides multi-sensor data for relevant first response operations, and forms a basis platform for continued research and development.

In terms of exploration behavior, the current capabilities of the MAX drone are limited to a rather basic, randomized frontier-based exploration. Making progress in that area will be a main ambition in the next phase, through development and integration of a more advanced and flexible target candidate point selection process. In addition, mechanism to adapt the behavior of the drone according to different conditions and information needs of first responders, e.g. regarding exploration speed or exploration completeness, preferred search strategies or the quality of acquired utility sensor data, is another area for further development.

ACKNOWLEDGMENTS

This work has been supported by the European Union's Horizon 2020 Research and Innovation Programme and the Korean Government under Grant Agreement No 833435 (INGENIOUS). Content reflects only the authors' view; The Research Executive Agency (REA) and the European Commission are not responsible for any use that may be made of the information it contains.

REFERENCES

Alexis, K. (2019) Resilient Autonomous Exploration and Mapping of Underground Mines using Aerial Robots, *19th International Conference on Advanced Robotics (ICAR)*, December 2019, 1-8.

Brunel, A., Bourki, A., Demonceaux, C. and Strauss, O. (2021) SplatPlanner: Efficient Autonomous Exploration via Permutohedral Frontier Filtering, *ICRA 2021 – 38th IEEE International Conference on Robotics and Automation*, Xi'an, China, May 2021, 608-615.

Chen, F., Martin, J., Huang, Y. and Wang, J. (2020) Autonomous Exploration Under Uncertainty via Deep Reinforcement Learning on Graphs, *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 6140-6147.

Chen, X., Zhang, H., Lu, H., Xiao, J., Qiu, Q., and Li, Y. (2017). Robust SLAM system based on monocular vision and LiDAR for robotic urban search and rescue. *2017 IEEE International Symposium on Safety, Security and Rescue Robotics (SSRR)*, 41–47.

Dellaert, F. (2012) Factor graphs and GTSAM: A hands-on introduction. *Tech. rep. Georgia Institute of Technology*.

Gupta, A. and Fernando, X. (2022) Simultaneous Localization and Mapping (SLAM) and Data Fusion in Unmanned Aerial Vehicles: Recent Advances and Challenges. *Drones*, 6, 85.

Hardouin, G., Morbidi, F., Moras, J., Marzat, J. and Mouaddib. E. M. (2020) Surface-driven Next-Best-View planning for exploration of large-scale 3D environments, IFAC-PapersOnLine, 53, 2, 15501-15507.

Holmberg, M., Karlsson, O. and Tulldahl, M. (2022) "Lidar Positioning for Indoor Precision Navigation". In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, 2022, pp. 359–368.

Hornung, A., Wurm, K. M., Bennewitz, M., Stachniss, C. and Burgard, W. (2013) OctoMap: An Efficient Probabilistic 3D Mapping Framework Based on Octrees, *Autonomous Robots*, 34, 189-206.

Karam, S., Nex, F., Chidura, B. T. and Kerle, N. (2022) Microdrone-Based Indoor Mapping with Graph SLAM, *Drones*, 6, 352.

Krátký, V., Petráček, P., Báča, T. and Saska, M. (2021) An autonomous unmanned aerial vehicle system for fast exploration of large complex indoor environments, *Journal of Field Robotics*, 38, 8, 1036-1058.

Lindqvist, B., Kanellakis, C., Mansouri, S. S., Agha-mohammadi, A. and Nikolakopoulos, G. (2022) COMPRA: A COMPact Reactive Autonomy Framework for Subterranean MAV Based Search-And-Rescue Operations, *Journal of Intelligent & Robotic Systems*, 105, 49.

Lowe, D. G. (2004) "Distinctive image features from scale-invariant keypoints," *International Journal of Computer Vision*, vol. 60, 91–110.

Martínez Novo, Á., Lu, L. and Campoy, P. (2022) FAST RRT* 3D-Sliced Planner for Autonomous Exploration Using MAVs, *Unmanned Systems*, 10, 2, 175-186.

Papachristos, C., Kamel, M. S., Popovic, M. and Khattak, S. (2019) Autonomous Exploration and Inspection Path Planning for Aerial Robots Using the Robot Operating System, *Robot Operating System (ROS)*, 67-111.

Ravankar, A., Ravankar, A., Kobayashi Y. and Emaru T. (2018) Autonomous Mapping and Exploration of UAV Using Low Cost Sensors, *5th International Electronic Conference on Sensors and Applications*, 4, 5753.

Rydell, J. and Bilock, E. (2015) "Panther Chameleon: Real-time Image-based Positioning and Mapping", *Proceedings of the 2015 International Technical Meeting of The Institute of Navigation*, Dana Point, California, January 2015, 768-777.

Rydell, J., Bilock, E. and Larsson, H. (2016) "Navigate into Danger", *Proceedings of IEEE/ION PLANS 2016*, Savannah, GA, 57-66.

Shen, S., Michael, N. and Kumar, V. (2012) Autonomous indoor 3D exploration with a micro-aerial vehicle, *2012 IEEE International Conference on Robotics and Automation*, May 2012.

Wang, J. and Englot, B. (2020) Autonomous Exploration with Expectation-Maximization, *Springer Proceedings in Advanced Robotics*, 10, 759–774.

Williams, J., Jiang, S., O'Brien, M., Wagner, G., Hernandez, E., Cox, M., Pitt, A., Arkin, R. and Hudson, N. (2020) Online 3D Frontier-Based UGV and UAV Exploration Using Direct Point Cloud Visibility, *IEEE International Conference on Multisensor Fusion and Integration for Intelligent Systems (MFI)*, September 2020, 263-270.

Yamauchi, B. (1997) A frontier-based approach for autonomous exploration, *Proceedings of the 1997 IEEE International Symposium on Computational Intelligence in Robotics and Automation CIRA'97*, 146-151.

Yu, N. and Wang, S. (2019) Enhanced Autonomous Exploration and Mapping of an Unknown Environment with the Fusion of Dual RGB-D Sensors, *Engineering*, 5, 1, 164-172.

Zagradjanin, N., Pamucar, D., Jovanovic, K., Knezevic, N. and Pavkovic, B. (2022) Autonomous Exploration Based on Multi-Criteria Decision-Making and Using D* Lite Algorithm, *Intelligent Automation & Soft Computing*, 32, 3, 1369–1386.

Zhang, J. and Singh, S. (2014) "LOAM: Lidar odometry and mapping in real-time." *In: Robotics: Science and Systems*. Vol. 2. 9. Berkeley, CA, 1–9

Zhang, N., Nex, F., Vosselman, G. and Kerle, N. (2022a) Training a Disaster Victim Detection Network for UAV Search and Rescue Using Harmonious Composite Images, *Remote Sensing*, 14, 13.

Zhang, N., Nex, F., Vosselman, G. and Kerle, N. (2022b) Unsupervised harmonious image composition for disaster victim detection, *The International Archives of Photogrammetry, Remote Sensing and Spatial Information Sciences*, 43, 1189--1196.

Zhang, Z., Shi, C., Zhu P., Zeng, Z. and Zhang, H. (2021) Autonomous Exploration of Mobile Robots via Deep Reinforcement Learning Based on Spatiotemporal Information on Graph, *Applied Sciences*, 11, 18.