










# An aerial/ground robot team for autonomous firefighting in urban GNSS-denied scenarios

Simón Martínez-Rozas<sup>1,\*</sup> , Rafael Rey<sup>1,\*</sup> , David Alejo<sup>1,\*</sup> , Domingo Acedo<sup>1,\*</sup> ,  
 José Antonio Cobano<sup>1</sup> , Alejandro Rodríguez-Ramos<sup>2</sup> , Pascual Campoy<sup>2</sup> ,  
 Luis Merino<sup>1</sup>  and Fernando Caballero<sup>1</sup> 

<sup>1</sup>Service Robotics Laboratory, Universidad Pablo de Olavide, Crta. Utrera, km. 1, 41013, Sevilla, Spain

<sup>2</sup>Computer Vision and Aerial Robotics Group, Centre for Automation and Robotics (C.A.R.), Universidad Politécnica de Madrid (UPM-CSIC), Calle Jose Gutierrez Abascal 2, 28006 Madrid, Spain

**Abstract:** This paper presents a framework for urban firefighting with a heterogeneous aerial/ground robot team. The system was developed to address Challenge 3 of the MBZIRC'20. The challenge required autonomously detecting, locating, and extinguishing fires on multiple interior floors and on exterior facade and ground surfaces. Our multi-robot system consists of a robot team of up to three Unmanned Aerial Vehicles (UAV) and one Unmanned Ground Vehicle (UGV). We describe the main hardware and software components for UAV and UGV platforms and present the system's main algorithmic components: a 3D LIDAR-based Mapping and Localization module able to work in GPS-denied scenarios, a Global Planner and a fast local re-planning system for robot navigation, Infrared-Based Perception and robot actuation control for fire extinguishing, and Mission Executive and coordination module based on Behavior Trees. Finally we describe our results from the competition, where the system worked fully autonomously and scored in all the trials performed. The presented system ended in 7<sup>th</sup> place (out of 20 teams) in Challenge 3 and 5<sup>th</sup> place (out of 17 teams) in the Grand Challenge's Challenge 3 portion.

**Keywords:** Multi-robot, GPS-denied, Autonomous navigation, Fire-extinguishing

## 1. Introduction

The application of robotic technologies in disaster management and rescue operations has increased in recent years, due to advances in platform design and capabilities for mapping, localization, perception, planning, and coordination capabilities (Delmerico et al., 2019). The goals continue to be: Decrease the risks that humans face in such activities and, at the same time, increase operational effectiveness.

Robot competitions foster advancements in robotics research for these applications and allow proper benchmarking for different approaches. Some examples related to search-and-rescue and

---

\*Equal contribution.

Received: 30 September 2020; revised: 22 February 2021; accepted: 16 May 2021; published: 21 March 2022.

**Correspondence:** Simón Martínez-Rozas, Service Robotics Laboratory, Universidad Pablo de Olavide, Crta. Utrera, km. 1, 41013, Sevilla, Spain, Email: [simon.martinez@uantof.cl](mailto:simon.martinez@uantof.cl)

This is an open-access article distributed under the terms of the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

Copyright © 2022 Martínez-Rozas, Rey, Alejo, Acedo, Cobano, Rodríguez-Ramos, Campoy, Merino, Caballero

disaster-management scenarios are: Defense Advanced Research Projects Agency (DARPA) Robotics Challenge (Pratt and Manzo, 2013), RoboCup Rescue competition (Sheh et al., 2016), euRathlon and ERL Emergency (Winfield et al., 2017), DARPA Fast Lightweight Autonomy program (Mohta et al., 2018), Trinity College Robot Competition Schedule Changes (TrinityCollege, 2020) and the Mohamed Bin Zayed International Robotics Challenge (MBZIRC, 2020).

Among these missions and challenges, firefighting constitutes a relevant application of robot systems, given the dangerous conditions in which human firefighters must operate. Human agents are at a constant risk of being burned, becoming trapped, or inhaling smoke. The risks of losing lives could be greatly reduced by using autonomous robots capable of searching, detecting and extinguishing fires (Chen et al., 2019). Current automated systems, however, are generally fixed in place. For instance, automatic fire sprinklers and alarms are used in heavily populated and hazardous areas for rapidly extinguishing any fire. Mobile robotic firefighting systems, such as ground and aerial robots, on the other hand, could traverse unsafe and/or inaccessible areas to collect information through onboard sensors (visual and infrared cameras, LIDAR, etc) (Pecho et al., 2019) and even to perform extinguishing operations (Innocente and Grasso, 2019). Such technology would offer critical new capabilities on real-time situation awareness and remote extinguishing to firefighters. Achieving this vision still requires, though, advancing the state of the art and improving robustness, both in software and hardware. We also need progress in other, specific areas (Delmerico et al., 2019), such as performance in indoor and outdoor environments, clear-fire image detection (Aydin et al., 2019), and autonomous navigation (Imdoukh et al., 2017).

The MBZIRC 2020 competition addressed this opportunity directly in its Challenge 3: “Team of Robots to Fight Fire in High Rise Building”. The mission required a collaborative robot team to detect, localize, and extinguish fires in a simulated high-rise building. The tasks involved putting out fires located inside and on the walls of a building at different heights and floors, and outside the building, by using water and fire blankets. This challenging mission requires extending the state-of-the-art in several areas: deploying and coordinating multiple, heterogeneous robots (UAVs and UGV) in the same urban workspace; navigating autonomously in cluttered, indoor/outdoor environments; autonomously seeking and locating of fires; adapting hardware and control systems for fire extinguishing.

This paper details a fully autonomous, multi-robot system for extinguishing fires in such urban high rise building firefighting scenario. Our robot team comprises three aerial and one ground vehicles. Of 20 teams competing in Challenge 3, our system achieved seventh place. In the Challenge 3 portion of the Grand Challenge, we finished fifth out of 17 teams.

Our main contributions relate directly to the technical advances required to address MBZIRC Challenge 3 with a fleet of heterogeneous robot in fully autonomous mode. More specifically, we developed:

- A multi-modal robust localization module, able to operate in GPS-denied scenarios
- A flexible mission composition and execution system based on Behaviour Trees (BT) (Colledanchise and Ogren, 2017) in charge of combining all modules to define the behaviors of each robot and that of the whole team to accomplish missions
- Specialized navigation, control and perception modules, and hardware employed in the competition.

While the competition allowed the use of GPS and even DGPS (with a penalization in this latter case), the ability to operate without GPS was crucial to navigate autonomously and score in all trials, given the intermittent reception due to satellite occlusion and other factors. The versatility of BTs was as well fundamental to adapting quickly to competition conditions.

The paper is organized into eight sections. Section 2 describes the hardware designs for the UAV and UGV platforms. Section 3 presents the mapping and localization system implemented. The autonomous navigation system of each platform is described in Section 4. Section 5 shows the approach proposed to detect and extinguish the fires and Section 6 presents the framework

used to carry out the mission. Experiments are presented in Section 7. Finally, lessons learned and conclusions are described in Section 8.

### 1.1. State of the art

The idea of using robots for firefighting has been present in the community for many years. Early designs can be traced back to (Bradshaw, 1991), which describes the functional and mechanical design of a ground robot for fire detection and first intervention in indoor environments. In (Amano et al., 2001), the authors show the design of a climbing robot for helping fire activities in buildings. All in all, robots are considered a promising tool to fight fires because of their ability to reach inaccessible zones, assist in dangerous environments, and to carry sensors to monitor and control the fires.

Actually, robotic technology is being often deployed as part of firefighting operations in the last years. Current examples can be found in the Notre Dame fire in 2019 with the firefighting ground robot called Colossus<sup>1</sup>. It is a remote-controlled firefighting automaton designed and built by Shark Robotics. TC800-FF is another remotely operated ground robot with some autonomous navigation capabilities designed to assist fire-fighter during operations<sup>2</sup>. Two of the most popular firefighting ground robots are the Thermite RS1-T3 and RS2-T2<sup>3</sup>. The aim of all these robots is to replace fighters and to keep them safe and free from the heavy work that distracts them and takes time away from solving problems quickly and effectively. Most of these systems are mainly teleoperated. While teleoperation can be enough for certain scenarios, the inclusion of autonomous capabilities expands further the utility of such systems, and can be very relevant when fleets of several robots are involved, like in the MBZIRC challenge.

Furthermore, in the case of high-rise buildings and in other scenarios, the use of aerial robots can be beneficial. In the last 15 years, Unmanned Aerial Vehicles (UAVs) have been more and more used in wildfire fighting (e.g. (Ambrosia et al., 2003; Casbeer et al., 2006; Merino et al., 2006; Skeele and Hollinger, 2016; Bailon-Ruiz et al., 2018)). These systems are mainly used as “eyes in the skies”, to provide situational awareness to first response teams, given their capability of carrying sensors and position in vantage points. Some of these works (Casbeer et al., 2006; Merino et al., 2006; Bailon-Ruiz et al., 2018) have shown the advantages of using teams of cooperating UAVs for firefighting. In particular, we have also shown in our former works (Merino et al., 2006; Merino et al., 2012) the added value of cooperating heterogeneous UAVs with different capabilities and that can carry complementary sensors in order to localize precisely fire spots and discard false alarms. Currently, drones are used as part of wildfire operations, as for instance during the Australian bushfires where drones were used to search and rescue koalas affected by the fires (Witt et al., 2020).

Those works do not consider the inclusion of fire suppression mechanisms, which for UAVs is also a challenge. A novel hose type robot, which can fly directly into the fire source via a water-jet, has been proposed in (Ando et al., 2018). In (Yamada and Nakamura, 2016) the concept of attaching a drone to the tip of a fire hose and remotely extinguishing a fire is presented. These studies have not been tested in a real fire extinguishing task. Problems such as the payload and influence of the tension of the hose itself on the flight must be addressed to bring the UAV’s hose closer to the fire source.

When considering the use of autonomous robots to fight urban fires in buildings as in MBZIRC, additional challenges are present, like the navigation in cluttered and GPS-denied environments or with limited GPS coverage, as is typically the case in urban canyons and/or close to and inside buildings, and the operation in low-visibility conditions (Schneider and Wildermuth, 2017). The last years have seen important advances in the navigation of UAVs in GPS-denied scenarios (for instance (Shen et al., 2014; Mohta et al., 2018; Perez-Grau et al., 2018; Usenko et al., 2020)). However, there

<sup>1</sup> <https://www.shark-robotics.com/colossus>

<sup>2</sup> <https://www.robotpompier.com/en/>

<sup>3</sup> <http://www.roboticfirefighters.com/>



**Figure 1.** Arena of Challenge 3. There are simulated fires on the facades of the building (top-left), and inside the building (down, right), located on the ground, first and second floor (to be suppressed with water). There are also fires located on the ground outside (down, left) that need to be suppressed with a blanket.

are still issues that need to be addressed to navigate and precisely position the UAVs to attack fires in urban scenarios like the one posed by Challenge 3 of MBZIRC.

The cooperation between UAVs and UGVs presents additional challenges. In (Michael et al., 2014) the synergies between a UGV and UAVs are exploited to map buildings affected by earthquakes. But there are not many works in the literature that consider fully-autonomous ground and aerial robots for firefighting in urban scenarios. In (Maza et al., 2011), a system of multiple UAVs cooperating with a static sensor network for a fire scenario in a building is presented. UAVs are used to deploy new sensors and provide situation awareness, which is then used to attack the fire by firefighters. The work shows again how cooperation can be very valuable. However, the system assumes GPS coverage, no ground robots are considered, and no extinguishing is performed by the robots themselves.

In sum, addressing a fire extinguishing scenario in a high-rise building by a team of robots, as the one posed by MBZIRC, still requires advances in localization, navigation, perception, and mission coordination and execution of the system as a whole. The paper presents in the next sections the proposed advances in those lines, discussing further literature pertinent for each case.

## 1.2. Problem statement

This section summarizes the main aspects of the rules, conditions, and scoring of the MBZIRC Challenge 3. The challenge requires a team of UAVs (maximum three) and UGV (maximum one) collaborating in order to autonomously extinguish a series of simulated fires in an urban high rise building firefighting scenario. Figure 1 shows the arena for Challenge 3, with a size of approximately that of half a football pitch. A 15-meter high structure simulates the building.

The robots have to extinguish a set of six fires located in the three floors of the building using water. For each floor, there is one fire inside the building and one fire on the facade (see Figure 1). Furthermore, there are two fires located outside the building at ground level that must be extinguished using fire blankets carried by the robots. The locations of the fires are changed randomly between trials.

Each fire is scored between 0 and 1. For the fires in the building the maximum score is obtained if the amount of water properly thrown into the fire is 1 liter. Simulated fire spots on the ground score when they are covered with a blanket (with a maximum when they are totally covered). The final score is obtained by a weighted sum for all eight fires. Different fires receive different weights according to their difficulty.

The mission can be attempted in either fully autonomous or manual mode. The mission is considered in manual mode from the moment there is an intervention by a human (the only interventions allowed are to replenish water and change batteries and blankets). Any score in manual mode is considered below a score in autonomous mode, or to disambiguate in case of a draw for



**Figure 2.** The ground platform used for Challenge 3.(a) Showing performance in the scenario of MBZIRC (b) Deploy of onboard sensors and Water Extinguishing System

scores in autonomous mode. GPS is allowed. RTK/DGPS can be used although a penalty of 25% will be applied if it is used. The duration of Challenge 3 is 15 minutes, with a 5-minute preparation slot.

## 2. Hardware

The robot team to address the challenge consists of one UGV and up to three UAVs. They are adequately equipped to perform navigation and to detect and extinguish the fire for this firefighting scenario.


### 2.1. Unmanned ground vehicle

The ground platform called SIAR (Alejo et al., 2020) is a 6-wheeled robot with independent traction system and a pan&tilt mechanism on top. The robot is able to navigate at a maximum speed of 0.7 m/s and it has battery autonomy for more than 4 hours of operation (motion and sensing/computation). The robot integrates an i7 computer, local sensing and a 3D LIDAR. The main sensors and mechanisms used in the ground platform are (see Figure 2):

- GPS sensor UBlox 8 and ArduIMUv3 Inertial Measurement Unit (IMU).
- LIDARs: A 3D OS-1-16 LIDAR Ouster, located on top of the robot. The LIDAR is able to provide 320k 3D points per second, ranging from 0.2 to 120 meters. The information gathered by this sensor is used for localization and mapping. Furthermore, a Hokuyo UTM-30LX 2D sensor is installed. The information of this sensor is used for local obstacle detection/avoidance.
- Cameras: A SeekThermal IR camera to detect the sources of heat and an Orbbec's ASTRA camera that provides the operator with images to operate the platform if needed.
- Pan&Tilt unit: It moves the thermal camera for searching and centering the fire, and at the same time, it holds and guides the hose of the extinguishing system to aim the fire.
- Extinguishing System: It consists in a series connection of two water pumps (5V, 4.8 W, 300cm Hmax, 300L/H, no valve), one MOSFET module 5V, hoses, and a 3-liter water tank (within the limits of the competition). To improve the intake of water, the tank is placed above the water pumps. Alike, the output hose is placed above the pump to prevent the water from going out, as the pumps do not have a valve. The system ejects 2 meters in a straight line water.

### 2.2. Unmanned aerial vehicles

Regarding to the aerial platforms, the models used for the competition are adaptations of the DJI Matrice 210 V2 (M210) and the DJI Matrice 600 Pro (M600) (see Figure 3). We selected these platforms because of the robustness of the hardware and software, support that the provider




**Figure 3.** (a) Adapted DJI Matrice 210v2 for facade fire extinguishing, with LIDAR and PC placed to avoid blocking the down-looking sensors. (b) Adapted DJI Matrice 600 for blanket deployment and facade fire extinguishing.

offers, and their facility to communicate the PC on-board with the Robot Operating System (ROS) (Quigley et al., 2009). The drone’s structure was adapted in order to perform the tasks and to integrate new hardware. For that, we take into account the maximum weight that can carry each drone and also the dimension constraints of the competition. Each UAV robot integrates an i7 NUC computer and local sensing. The main sensors and mechanisms added to the UAVs are (see Figure 3):

- DJI Matrice onboard sensors. This includes single GPS, Inertial Measurement Unit (IMU) and barometric altimeter. DJI Matrice 600 integrates 3 GPS receivers and 3 IMUs for redundancy and fault detection purposes, but they behave as single devices.
- 3D LIDAR. OS-1-16 LIDAR Ouster. The purpose of this sensor is threefold: environment mapping, localization, and odometry. It is the basis for the generation of the 3D map of the environment, and to compute an accurate robot odometry for the UAVs. It is also used as main sensor for robot localization in all robots, and to detect obstacles by the navigation module.
- Infrared camera: SeekThermal used for fire detection. In M210 carries one camera looking forward, while the M600 carries an additional one in a nadir configuration for the ground fires.
- Water Extinguishing System: we use the same pumps as the UGV platform, but with less capacity due to payload constraints. We use two 0.5-liter water tank, positioned on each leg of the drone to evenly distribute their weight. The tanks’ outlets are joined to carry the water to the pumps. An aluminum tube is used to guide the outlet hose from the pumps to a safe area in which the air turbulence generated by the drone itself does not disperse the water when it is ejected. A flow restrictor is placed at the outlet of the hose to increase the pressure and range.
- Blanket Extinguishing system (M600 only): we use six electromagnets to carry and drop the blanket. Electromagnets are located on a bar placed between the back two legs of the drone. The blanket is transported rolled up to prevent the air turbulence generated by the drone from producing a sail effect. The electromagnets are distributed into two groups of three electromagnets each. In each group, there is one electromagnet for rolling the blanket and two for dropping it. When desired, the rolling electromagnets are deactivated and the blanket is unrolled, and when the blanket is to be thrown, the drop electromagnets are deactivated.

### 3. 3D Mapping and Localization

The application considered requires that the robots are localized seamlessly and accurately both in GPS and GPS-denied environments (including buildings interiors, urban canyons and, in general, areas with high dilution of precision), with a smooth transition among them. Our solution consists of a map-based localization system that fuses data from different sensors.



**Figure 4.** Two views of the map generated for Challenge 3 of the scenario shown in Figure 1 with data provided by the M600 UAV. (a) a map perspective in 3D space and (b) the top view of the map.

The proposed mapping and localization system consists of two main components : a LIDAR-based odometry and mapping and an online map-based global localization module. These systems create the global map during the mapping phase, and estimate the global position of the robot into the map during the mission respectively. While there are some minor differences in the implementations for the UAVs and UGV, the modules are based on the same principles.

### 3.1. 3D Mapping and multi-sensor odometry

Different approaches can be used for 3D mapping and odometry in aerial vehicles. Stereo cameras (Kitt et al., 2010; Geiger et al., 2011; Schmid et al., 2013) and RGBD (Endres et al., 2012; Kerl et al., 2013) are common choices indoors or for outdoor navigation close to obstacles. However RGBD and time-of-flight cameras cannot work under strong sunlight, and stereo cameras would need a large baseline in order to accurately estimate distances at 25 meters approximately (as the case of Challenge 3). Bundle adjustment approaches (Cucci et al., 2017; Pagliari et al., 2015) can be used offline in combination to visual cameras or short-baseline stereo to build accurate maps outdoors, but they cannot provide reliable odometry during the robot operation. On the other hand, 3D-LIDAR approaches work outdoors, range up to 120 meters and can perceive the 360° of the robot’s environment.

Thus, the base of the proposed mapping and localization systems is a LIDAR-based odometry system that estimates the relative motion of the robot. This module considers IMU and 3D LIDAR measurements to compute a reliable 6DoF robot odometry. It also provides short-term aligned scans that can be used for local mapping, and a map of the environment.

The LIDAR-based odometry module makes use of the LIDAR Odometry and Mapping (LOAM) system (Zhang and Singh, 2014; Zhang and Singh, 2017), which decomposes the original SLAM problem into scan registration and mapping. The system computes high-precision motion estimates and maps. We make use of an open software implementation of the original LOAM system<sup>4</sup>.

In the mapping phase, a drone equipped with LIDAR, GPS and on-board sensors (IMU, altimeter, etc) was flown manually over the Challenge 3 arena , mapping the environment at different altitudes, from 2 meters to the top of the building. The resulting 3D map is shown in Figure 4. It can be seen how the system accurately mapped the interior of the building and also the four scaffolding structures in the corners. Originally, the map also included the surrounding buildings, but they were removed to reduce its size because these areas were unreachable by the robots. During the mapping stage, the GPS data was used to compute the relative transformation between the map-based global frame and the GPS coordinate system for multi-sensor localization.

<sup>4</sup> <https://github.com/HKUST-Aerial-Robotics/A-LOAM>

### 3.2. Multi-sensor robot localization

The localization module provides a robust global position of the robot within the map built in Section 3.1. Map-based approaches for localization are usually better than SLAM approaches in terms of reliability and computational requirements. Thus, decoupling the robot localization from the map building process saves computation resources because the map is built offline once, and the complexity of the algorithm is usually smaller with the corresponding impact in reliability.

One of the approaches based on predefined models of the environment is commonly known as teach-replay (Gridseth and Barfoot, 2019; Royer et al., 2007), which is accomplished in two stages: first the robot is manually piloted along the desired path as in a teaching phase, and an accurate 3D map of the environment is built, along with the robot motion from this learning path; afterwards, this map is used to locate the robot when it repeatedly visits the same path. Although very effective, these approaches are not general enough since they only enables a robot to follow a predetermined trajectory.

Monte Carlo Localization (MCL) is another approach that makes use of a known map of the environment, and is commonly used for robot navigation in indoor environments (Thrun et al., 2001). It is a probabilistic localization algorithm that makes use of a particle filter to estimate the pose of the robot within the map, based on sensor measurements. An extension of this approach for 6D localization based on 2D laser scanner is presented in (Hornung et al., 2010), but it is meant for the 2D motion of humanoid robots in a 3D environment, which makes it not suitable for aerial robots. The use of GPUs can enhance the computing capabilities of the on-board system regarding MCL-based approaches, as proposed by other authors (Kanai et al., 2015; Fallon et al., 2012).

The approach used in this paper is based on Monte Carlo localization. It extends our previous work presented in (Perez-Grau et al., 2017b) with the following aspects:

- The particle prediction model has been adapted to absorb non-modelled errors in the odometric system. One of the major issues of MCL approaches are their dependence on the odometric model. A faulty odometry quickly leads to filter divergence due to a wrong dispersion of the particles into the solution space. We included new random terms into the particle prediction model in order to account for these non-modelled errors.
- New sensor updates: Compass, altimeter and GPS were integrated into the localization approach in order to build a multi-sensor approach with increased reliability. The MBZIRC competition requires that the robots are localized seamlessly and accurately both in GPS and GPS-denied environments, with a smooth transition among them. Integrating multiple sensors improves the reliability of the solution, reducing the dependence on a single sensor and also eases the transition from GPS-denied to full-GPS navigation.
- General improvements for faster and robust execution. The localization system is an essential element of the robot navigation, it is involved in almost every robot task, so it must be robust, and efficient because robot on-board computation is very limited.

The source code of the version used in the experiments is publicly available<sup>5</sup>.

The localization filter maintains  $N$  pose hypotheses (particles) defined as  $\mathbf{p}_i^t = [x_i^t, y_i^t, z_i^t, \Psi_i^t]$ , where  $x_i^t, y_i^t, z_i^t$  and  $\Psi_i^t$  refers to the robot position and yaw angle for particle  $i$  at time  $t$ . Each particle has an associated weight  $w_i^t$  such as  $\sum_i^N(w_i^t) = 1$ . Note how the robot roll and pitch angles are not included into the robot pose definition. The onboard IMU already provides their values with high-enough accuracy. This greatly reduces the computational complexity of the algorithm, allowing for real-time onboard computation.

The particles are initialized by setting the initial pose (at the take-off spot) and distributing them in the hypotheses space. The LIDAR-based odometry described in Section 3.1 provides increments

<sup>5</sup> <https://github.com/robotics-upo/mcl3d>



of the robot pose in the robot frame (after roll and pitch compensation)  $[\Delta x^t, \Delta y^t, \Delta z^t, \Delta \Psi^t]$ , which are used to propagate the distribution of the particles at each time step. Thus, the state of each particle  $\mathbf{p}_i^t = [x_i^t, y_i^t, z_i^t, \Psi_i^t]$  will evolve according to the following expressions:

$$\begin{aligned} x_i^{t+1} &= x_i^t + N_x + \Delta x^t \cdot \cos(\Psi_i^t) - \Delta y^t \cdot \sin(\Psi_i^t) \\ y_i^{t+1} &= y_i^t + N_y + \Delta x^t \cdot \sin(\Psi_i^t) + \Delta y^t \cdot \cos(\Psi_i^t) \\ z_i^{t+1} &= z_i^t + N_z + \Delta z^t \\ \Psi_i^{t+1} &= \Psi_i^t + N_\Psi + \Delta \Psi^t \end{aligned} \quad (1)$$

The values of  $\Delta x^t$ ,  $\Delta y^t$ ,  $\Delta z^t$ , and  $\Delta \Psi^t$  are drawn randomly following a normal distribution centered in their actual values and standard deviations proportional to each increment itself. The values of  $N_x$ ,  $N_y$ ,  $N_z$  and  $N_\Psi$  are design parameters and they are drawn randomly from a zero-centered normal distribution with fixed standard deviation. The main purpose of these parameters is to account for outliers in the odometry estimation. They force to randomly sample solutions not included into the odometry estimation in view of possible errors. This random noise increases the robustness of the solution by exploring particle states near the odometry prior.

The approach evaluates the sensors only when the robot moves above given thresholds in translation or rotation. In this case, it performs an update: using the previous equations to predict robot position according to odometry, updating particles based on sensors and re-sampling the hypotheses space when required. The following sensor readings are used to update the weight  $w_i^t$  associated to each particle's hypothesis  $\mathbf{p}_i^t$  (for the sake of clarity, the time index  $t$  is removed from the following equations):

- The point clouds by the 3D LIDAR sensor are transformed to each particle pose in order to find correspondences between the cloud and what the map should look like from that particle's pose. Since this is very expensive computationally, a 3D probability grid is computed offline as in (Hornung et al., 2010; Perez-Grau et al., 2017a). This grid stores a value of how likely is that a given position falls within an occupied point of the map. Thus, the probability of each cell  $\mathbf{c} = [x_c, y_c, z_c]$  is computed as a Gaussian distribution centered in the closest occupied point in the map  $\mathbf{m} = [x_m, y_m, z_m]$ , and whose variance  $\sigma^2$  depends on the LIDAR noise.

$$grid(\mathbf{c}) = \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\|\mathbf{c}-\mathbf{m}\|^2/2\sigma^2} \quad (2)$$

This probability grid only needs to be computed once for a given environment, and relieves from performing numerous distance computations between each cloud point for each particle and its closest occupied point in the map. Besides, each point cloud is first transformed according to the current roll and pitch provided by the on-board IMU. This transformation is done just once per update, reducing the computational requirements as well. Then, for every point of the transformed cloud, we access its corresponding value in the 3D probability grid. This value is an indicator of how likely is that the point is part of the map. Assuming that the point cloud is composed of  $M$  3D points  $\mathbf{v}_j$ , the weight  $w_i^{map}$  of each particle  $\mathbf{p}_i$  is computed by adding the associated probability grid values:

$$w_i^{map} = \frac{1}{M} \sum_{j=1}^M grid(\mathbf{p}_i(\mathbf{v}_j)) \quad (3)$$

where  $\mathbf{p}_i(\mathbf{v}_j)$  stands for the transformation of the point  $\mathbf{v}_j$  to the particle's state  $\mathbf{p}_i$ , and  $grid(\mathbf{p}_i(\mathbf{v}_j))$  is the evaluation of the probability grid in such transformed position.

Equation (3) can be also computed as the product of all  $grid(\mathbf{p}_i(\mathbf{v}_j))$ . However this was discarded for the following reasons: a) outlier points  $\mathbf{v}_j$  might have a significant impact in the weight computation, leading to almost zero in some cases, b) point clouds area easily composed by dozens of thousands of points, given that  $grid(\mathbf{p}_i(\mathbf{v}_j))$  ranges from 1 to 0, we can fall into numerical errors in the product.

- The 3D position measurements provided by the GPS are used to check how well each particle matches the sensor reading. First, the GPS position is transformed into the map by means of the known fixed transformation computed during the mapping task. Then, a  $w^{gps}$  is computed for each particle. Assuming a GPS measurement  $\mathbf{p}^{gps}$ , we compute the associated weight  $w_i^{gps}$  as follows:

$$w_i^{gps} = \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\|\mathbf{p}_i - \mathbf{p}^{gps}\|^2 / 2\sigma^2} \quad (4)$$

where the distance computation dismisses the altitude contribution of the GPS, which is subject to significant noise. Also, the standard deviation of the GPS measurement  $\sigma$  is computed empirically from the data gathered during the mapping stage and to account for GPS errors in single configuration.

- We integrate the height above the ground, provided by the altitude sensor, into the MCL through the re-sampling stage. This forces the hypotheses to be distributed around the altitude provided by the robot altimeter. This way, we can reduce the dispersion of the particles around Z axis, which is easily observable by means of sensors such as laser altimeter, barometer or GPS.
- We also integrate the IMU yaw angle through the re-sampling stage. The rationale behind this decision is also based on the computational optimization and the nature of the yaw angle. In the experiments area, the yaw angle was checked to be only slightly distorted by the environment, so that we can trust the estimation from the magnetometer. Thus, every time a new re-sampling is performed, the orientation of the particles are drawn from a Gaussian distribution centered in the latest yaw value provided by the IMU. A global frame to map calibration is used to transform the IMU yaw into the map yaw.

As mentioned, GPS and point-cloud updates are integrated into the filter. Given the distinct nature of the two technologies involved, we calculate separate weights for each sensing modality. The final weight of each particle is computed as a weighted average:

$$w_i = \alpha * w_i^{map} + (1 - \alpha) * w_i^{gps} \quad (5)$$

where  $\alpha$  is chosen depending on the particularities of the indoor environment where the robot is going to operate. If the map used in the MCL does not contain the full environment or its accuracy is not enough to trust the map matching,  $\alpha$  should be lower than 0.5. Whereas if GPS measurements are not accurate,  $\alpha$  should be higher.

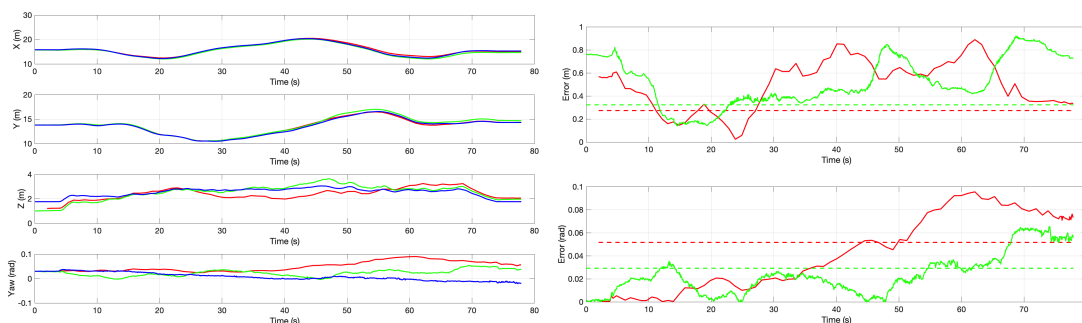
This approach is used for both robot systems, ground and aerial. However, the dispersion in Z axis is set to a very small value in the predictive model of the ground robot, as the UGV stays on the floor surface.

### 3.3. Experiments

A set of experiments has been conceived in order to validate and benchmark the presented approach. Thus, we first compared the estimation errors of our approach with respect to MCL3D (Perez-Grau et al., 2017b). We use a software implementation and datasets publicly available<sup>6</sup> to benchmark our approach. The dataset provides the LIDAR data gathered by an aerial robot flying indoors in a 20x20x5 meters volume. The dataset also provides the ground-truth robot position computed with a motion tracking system with millimetric accuracy.

Figure 5 shows the estimated robot position and orientation, as well as the computed errors with respect to the ground-truth. In this experiment only LIDAR information and the environment map was used, and no other sensor was considered. Both algorithms have been setup with the same general parameters and make use of 1000 particles. The LIDAR odometry presented in Section 3.1

<sup>6</sup> <https://github.com/fada-catec/amcl3d>



**Figure 5.** Robot localization estimation and errors. The red line stands for the MCL3D estimation, the green line denotes the estimation of our method and blue line stands for the ground-truth robot position and orientation (Left) Estimated robot position in X, Y, Z and Yaw angle. (Right) Computed instant error in position (meters) and yaw angle (radians) with respect to the ground-truth. Dashed lines shows the corresponding RMSE error.

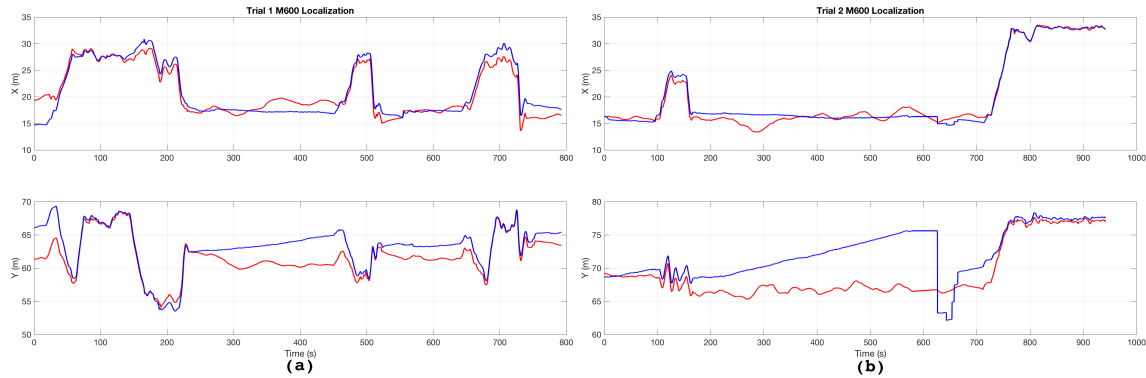
**Table 1.** Estimated RMSE localization error of the proposed approach

Approach	RMSE (m)	RMSE (rad)
MCL3D	0.2729	0.0517
Ours	0.3226	0.0292
Ours+Alt	0.2736	0.0245
Ours+Yaw	0.2555	0.0015
Ours+Alt+Yaw	0.2167	0.0016
Ours+Gps	0.3428	0.0302
Ours+Gps+Alt	0.2777	0.0241
Ours+Gps+Alt+Yaw	0.2205	0.0016
Ours+Gps+Yaw	0.2540	0.0016

was used to compute the robot odometry. It can be seen how the errors in position are approximately the same in both approaches. Both approaches follow the ground truth position with small errors in X and Y, while most of the errors are introduced in the Z estimation. On the other hand, the errors in the yaw estimation are clearly smaller in our approach, the RMSE error is roughly a 40% smaller. This improvement mostly comes from the new particle prediction model presented in (1). The possibility to introduce small prediction noise in the particle’s evolution helps to model small inaccuracies coming for the odometric system.

As previously commented, the errors in Z and yaw are predominant in our approach. This is related with the fact the 3D LIDAR mounted in the aerial robot mostly perceives the walls and structures, but not the floor. On the other hand, the yaw angle is also prone to drift due to error accumulation in the odometry. Section 3.2 proposes the integration of the localization approach with other sensors such as compass, GPS and altimeter to increase the accuracy and robustness of the system. Table 1 presents the computed RMSE error in the estimated robot’s position and orientation for different combinations of sensors.

Table 1 presents the localization error of our proposed system with noisy sensors. We generate these noisy readings by injecting Gaussian errors in measurements simulated from the ground truth. This is done in order to mimic the behaviour of a real sensor. Results show how the introduction of the yaw angle has a significant impact in the general estimation, not only in robot yaw (which is obvious) but also in position (see Ours+Yaw column). Similarly, the altimeter also improves the robot position estimation. However, introducing GPS measures does not provide a clear improvement in the estimation. This is probably produced by the relatively high sensor noise (0.5 m in X and Y position), which is in the order of the errors obtained in the map/point-cloud matching. Despite of this fact, having an alternative estimation of the robot position is beneficial from the point of view of redundancy, and this is the reason why we performed experiments also with GPS integration.



**Figure 6.** DJI M600 X-Y localization at the MBZIRC arena during execution of Trials 1 and 2. (a) Estimated XY position in Trial 1. (b) Estimated XY position in Trial 2. Red line denotes the MCL estimation and blue line denotes the GPS estimation.

The proposed approach integrating compass, altimeter and GPS was used and tested in experiments involving real robot operations during the MBZIRC trials. Figure 6 shows the estimated localization of our M600 during two experiments in the MBZIRC arena. The figure also shows the estimated GPS position of the M600 during the experiments (blue line in the graph), which indicates GPS estimation issues in both experiments. In this direction, Figure 6-b shows how from 175 to 625 seconds, approximately, the GPS estimation evolves at a constant velocity in Y axis, with a 15 meters jump in Y at second 625. During this period the drone was landed so the robot position should stay approximately constant.

Figure 6-a also shows constant velocity integration in axis Y from second 225 to 475 approximately. As previously noted, the GPS eventually recovers the right estimation, converging to the MCL estimation, which stays stable during the remaining of the experiment.

These GPS errors presented in Figure 6 were common during the MBZIRC 2020 execution. This behaviour was detected during the rehearsal days, and we finally decided to set the value of  $w^{GPS}$  in (5) to 0 to avoid integrating outliers into the estimation, leading to a complete GPS-free localization approach.

## 4. Autonomous robot navigation

The challenge requires the robots to reach arbitrary free poses in the map (outside and inside the building) to look for fires and to attack them once detected. We employ a two-stage autonomous navigation system, consisting of a global and a local planner, in order to deal with this environment. The Global Planner computes a global path in the map frame. The inputs are the goal to be reached, the static global map of the environment and the current robot localization given by the localization approach described in the previous section. The output is a feasible and safe path that should be followed in order to reach the goal. The Local Planner executes the global path and deals with dynamic and unmapped obstacles detected by the onboard 3D LIDAR. It takes as input the global path and a local map built by using the measurements of the onboard sensors. As output, it provides a free-collision local trajectory to the path tracker module. In the following, we describe the main techniques for each module.

### 4.1. Path planning

The application considered does not require precise dynamic motion planning for the vehicle. The main consideration is to provide safe paths within the workspace, handling the vehicle dynamics in the path tracker. At the same time, we need fast planning capabilities, as in the case of the local

planner a new local plan should be generated whenever a new LIDAR measure is obtained (and therefore a new local map is generated), at an approximate rate of 10Hz.

Different types of approaches for tackling 3D path planning have been proposed in the literature: discrete search optimal algorithms, sampling-based algorithms and bioinspired algorithms. Related to the first type, the A\* algorithm has been used in several works (Filippis et al., 2012; Sanchez-Lopez et al., 2019; Zhan et al., 2014). Another algorithm is the Lazy Theta\* algorithm, which is based on the A\* algorithm (Nash and Tovey, 2010). The Lazy Theta\* algorithm, an any-angle find-path algorithm, computes shorter paths than the A\* algorithm because it does not constrain paths to be formed by graph edges (to find short “any-angle” paths). Sampling-based algorithms such as RRT (Rapidly exploring Random Tree) have also been used in path planning (Lin and Saripalli, 2017). An comparison between A\* algorithm and RRT algorithm is done in (Zammit and van Kampen, 2020), both algorithms can be applied in real-time. The advantage of the graph-based algorithms with respect to sampling-based algorithms is that the computed path is closer to the shortest path, and the solution is deterministic. Bioinspired approaches imitate the behavior of humans or other animals, it is worth to mention Particle Swam Optimization (PSO) (Alejo et al., 2014), and also genetic algorithms (Conde et al., 2012). These algorithms generally require significant computational resources, and their performance may vary depending on the scenario considered. Considering the Challenge 3, we need to implement a safe and precise waypoint navigation system which rapidly plan paths. The Lazy Theta\* algorithm presents the most suitable characteristics and is chosen with respect to other planning algorithms due to its computational load and its high repeatability.

Thus, the path planners are based on the Lazy Theta\* algorithm (Nash and Tovey, 2010), a fast and reliable heuristic planner. However, we have modified the original Lazy Theta\* algorithm to foster safety while still maintaining its good computation time. Algorithm 1 shows the pseudo-code of the original algorithm and highlights the two modifications implemented (in red).

These two modifications of the Lazy Theta\* algorithm have the following effects:

- **Cost to reach a node.** We added a cost component to the Euclidean distance in the original algorithm (see lines 33 and 35 of the Algorithm 1). For each node in the grid, this cost is the distance to the nearest obstacle (`dist_obst` in lines 33 and 35). Then, the cost to reach a node is now a sum of the path length and the new cost component for the nodes that define the trajectory. The latter component is weighted by a factor  $C_w$  between 0 and 1. This results in paths with larger safe margins with respect to obstacles, depending on the value of the weight factor.
- **Line of sight.** The line of sight is limited in order to allow the algorithm to link two nodes only if they are not further than a given distance. Otherwise, the first modification would have no effect because the original Lazy-Theta\* algorithm joins nodes if there is a direct line of sight between them, regardless of the cost. Thus, a maximum line of sight distance is included, `max` in line 38 of the Algorithm 1, allowing the algorithm to link two nodes only if they are not further than this distance. If this distance is set to  $\infty$  then we recover the original Lazy-Theta\* algorithm. This will affect directly the number of expanded nodes.

These two modifications should provide paths with larger safety margins with respect to obstacles at the cost of exploring more nodes. The quality and the computation time of these paths will depend on the weight factor and maximum line of sight distance, so the values of these parameters in both planners should be tuned.

#### 4.1.1. UGV planner

The influence of the two modifications of the Lazy Theta\* algorithm has been analyzed in the 2D planner of the UGV. Figure 7 presents how the weight factor,  $C_w$ , and line of sight,  $LoS$ , influence the trajectories computed by the 2D planner. Figure 7-(a) shows the number of explored nodes for different weight factors and maximum line of sight distances. The number of explored nodes (and

```


1 Main()
2   open := closed := ∅;
3   g(sstart) := 0;
4   parent(sstart) := sstart;
5   /* g(s) is the length of the shortest path from sstart to s */
6   /* h(s) is the approximated distance from the goal to the vertex */
7   open.Insert(sstart, g(sstart) + h(sstart));
8   while open ≠ ∅ do
9     s := open.Pop();
10    SetVertex(s);
11    If s = sgoal then
12      | return "path found";
13    closed := closed ∪ s;
14    foreach s' ∈ nghbrvis(s) do
15      | If s' ∉ closed then
16        | If s' ∉ open then
17          | g(s') := ∞;
18          | parent(s') := NULL;
19          | UpdateVertex(s, s');
20      | return "no path found";
21  end
22 UpdateVertex(s, s')
23   gold := g(s');
24   ComputeCost(s, s');
25   If g(s') < gold then
26     | If s' ∈ open then
27       | open.Remove(s');
28     | open.Insert(s', g(s') + h(s'));
29  end
30 ComputeCost(s, s')
31   /* Path 2 */
32   /* c is the cost to reach a node */
33   If g(parent(s)) + c(parent(s), s') + Cw · distobst < g(s') then
34     | parent(s') := parent(s);
35     | g(s') := g(parent(s)) + c(parent(s), s') + Cw · distobst ;
36  end
37 SetVertex(s)
38   If NOT lineofsight(parent(s), s) & lineofsight(parent(s), s) < max then
39     | /* Path 1 */
40     | parent(s) := argmin (s' ∈ nghbrvis ∩ closed(g(s') + c(s, s')));
41     | g(s) := min (s' ∈ nghbrvis ∩ closed(g(s') + c(s, s')));
42  end

```


**Algorithm 1:** Lazy Theta\* with the introduced modifications in red.

thus, the computation time) augments as the weight factor increases for every value of the line of sight. This increase takes also place as the line of sight decreases. Figure 7-b shows trajectories computed with different weight factors and a maximum line of sight distance of 1.5m. The weight factors 0.15 and 0.3 generate a safer trajectory than the one obtained with the original algorithm (blue line). This can be appreciated when robot crosses the door.

In order to properly select the right values of  $C_w$  and  $LoS$ , we have to consider that the local planner must generate trajectories at an approximate rate of 10Hz to ensure fast planning



**Figure 7.** (a): Nodes explored depending on the weight factor and the line of sight. (b): 2D trajectories computed for the UGV coming into the building (2D map of navigation) by considering different values of the weight factor,  $C_w$ .




**Figure 8.** Influence of the line of sight and weight factor on the safety margin (top view). The distance cost is also shown, red area around the building is the area with a higher distance cost. (a) Trajectory computed without line of sight and the weight factor 0 (original Lazy-Theta\*), (b) Trajectory computed considering that line of sight is 2 and the weight factor 0.3. The base\_link frame is the initial position.

capabilities. Additionally, we must choose a good compromise solution between safety and trajectory length. The chosen values are  $LoS = 1.5$  and  $C_w = 0.15$ . It can be seen in Figure 7-(a) that the computation time considering  $C_w = 0.15$  is smaller than considering  $C_w = 0.3$ , while giving similar results. Although the number of explored nodes doubles those of the original algorithm, it still ensures the rate for fast planning capabilities.

#### 4.1.2. UAV planner

The same analysis has been performed in the 3D planner of the UAV. Figure 8 illustrates how the safety is fostered by showing a trajectory computed with the original Lazy Theta\* algorithm and another one considering the modified algorithm with  $LoS = 2$  and  $C_w = 0.3$ . The trajectory length in the first one is 14m and in the second one is 15.6m. This larger trajectory makes the safety margin with respect to the building to increase during the flight. The closest distance between the building and the path is 0.5m with the original algorithm. However, the trajectory of the modified algorithm departs from the building (up to 5.1m away) and the minimum distance, 2.5m, takes place when the UAV reaches the goal.



**Figure 9.** (a) Nodes explored to compute the trajectory with lines of sight from 1.0 to 4.0 and weight factors between 0 and 1 for each line of sight, (b) Influence of the line of sight and weight factor on the safety margin. The trajectories considered are the computed in the scenario shown in Figure 8 with the same initial point and goal.

Figure 9(a) shows the explored nodes with different values of the line of sight and weight factors. The explored nodes decrease as the line of sight increases. Instead, the number of nodes fluctuates a little with different values of the weight factor. Considering the original algorithm,  $LoS$  is not limited and  $C_w = 0$ , the number of explored nodes is much lower, 28,851. Therefore, it is noteworthy how the modified algorithm influences the number of explored nodes.

Figure 9(b) shows the maximum distance the paths moves away from the building considering different values of the line of sight and weight factor. In the case of the path computed with the original algorithm, this distance is always smaller than 2.5m (see Figure 8-a). Instead, the UAV is always at least 3.5m away as the modified algorithm is used, and even this distance is greater than 4m as  $LoS < 3$ . The highest distances are achieved with  $LoS = 1$ . Regarding the closest the path gets to the building, it is never shorter than 2.5m (distance between the goal and the building) as  $LoS < 2.5$ . This distance is 2.4m for  $LoS = 2.5$  and between 1.7m and 2.4m in the rest of cases,  $LoS = 3.0, 3.5, 4$ . Therefore, safer and smoother trajectories are computed thanks to the modifications performed. An example is the path shown in Figure 8-b.


Finally, we decided not to enter the building with UAVs in the trials. Although the advantages of the modified planning algorithm have been shown and the compromise between safety and planning capabilities can be reached, this decision guaranteed safety because the closest distance to the building would be the distance of the inspection points to the facade. For this reason, in the competition we did not limit the line of sight neither consider the distance cost,  $C_w = 0$ , in the 3D planner. This allowed us to maintain the fast planning capabilities of the 3D planner.

## 4.2. Path tracker


The Path Tracker module computes the velocity commands to navigate and follow the trajectory computed by the Local Planner in both platforms. The UGV makes use of a pure pursuit (PP) (Kelly, 1994; Park et al., 2014) controller that adapts the angular velocity of the robot to reach the closest point into the commanded trajectory. The main idea of PP is to calculate a real-time target waypoint (WP) which lays in the target path and is situated at a look-ahead (L) distance from the vehicle. When applied to non-holonomic robots, PP generates a commanded heading change  $\Delta\theta = \arctg(y/x)$  that is necessary to make the robot point towards WP, where  $x$  and  $y$  are the coordinates of WP in the reference frame of the robot.

On the other hand, the UAVs make use of the command velocity interface provided by the DJI software. This interface allows commanding linear velocities in X, Y and Z axes, and angular rates in yaw. A saturated proportional controller (see Figure 10) has been implemented to follow the commanded local trajectory in the UAVs. This control scheme gives us a trapezoidal velocity profile





**Figure 10.** (a) Basic control scheme of the path tracker implemented on the UAVs.



**Figure 11.** Top view (a) and 3D view (b) of the global planned trajectory (green line) and its associated executed trajectory, as estimated by our localization system (red line) during a flight of the DJI Matrice M210 in the second trial for Challenge 3.

in each axis in order to reach the current destination waypoint while ensuring convergence to the final destination as we have a first order system.


### 4.3. Experiments

Figure 11 presents the trajectory followed by the M210 UAV in one experiment, in red, alongside with the global planned trajectory in green. As we use a local planner in order to check for unmodeled obstacles on the environment, the tracking of the trajectory on the horizontal plane is not perfect and it tends to take shortcuts on the corners whenever possible. On the other hand, the system follows the altitude reference more closely.

The 3D planner obtained a remarkable performance that allowed us to perform all experiments during the MBZIRC competition without any collisions. The computation times of the global plan is of less than a second in its onboard NUC717 and using a single core. On the other hand, the local planner computes trajectories at 7Hz approximately which ensures fast planning capabilities of the 3D local planner.

## 5. Fire perception and extinguishing

Our approach for the fire extinguishing task in the MBZIRC - Challenge 3 consists of two modules: a perception module that detects fires and estimates their positions with a thermal camera; and an extinguishing module that attacks the fire by dropping a blanket or ejecting water.



**Figure 12.** (a) Simulated ground fire, colored temperature image and segmentation using the radiometric values (the circle shows the mean position of the fire). (b) Simulated facade fire, colored temperature image and segmentation using the radiometric values. The original temperature image gets saturated, while working with the radiometric field allows for a detailed segmentation.

### 5.1. Fire perception

The module of fire perception detects fires within the scene using as input the data from a thermal camera. As output, it indicates if a fire has been detected, as well as its 2D coordinates on the image plane and 3D coordinates in the scenario.

All fires in the competition were simulated using a thermal heated plate that reached temperatures up to  $110^{\circ}$ . The fires on the facades featured additionally a real fire using gas, while the remaining fires were also visually simulated by using red silk.

In our solution, we aimed for a fully infrared-based detection of the heat source. Thus, the main sensor for fire perception is the miniature infrared thermal camera from Seek Thermal mentioned above. The camera can detect temperatures between  $-40^{\circ}$  to  $330^{\circ}$  Celsius. We forked a ROS driver from ETHZ and made additional modifications<sup>7</sup>. The driver provides a grayscale temperature image, as well as a radiometric map.

Fire detection is carried out by analyzing the images provided by the infrared cameras. A calibration procedure transforms the raw radiometric data obtained by the cameras to temperature ranges. While obtaining absolute temperature values from radiometry is a very complex process, dependant on the material emissivity (for instance, a metal can radiate less IR energy for the same temperature than other materials with larger emissivity), atmospheric attenuation, and many other factors, here we are more interested on clear relative temperature differences, which simplifies the calibration procedure. The calibration is performed for the different kind of fires once. Then, given the clear difference of temperatures between the fire and the surroundings, a thresholding operation is performed to segment the fire pixels. As a result, the fire detection module provides coordinates on the image plane of fire spots. Figure 12 shows detection results obtained during the competition. During the competition, no false positives were generated by using the method.

### 5.2. Fire position estimation

The fire segmentation results are used for the visual servoing final approach to extinguish the fire (see Section 5.3). However, during the exploration of the scenario we need to estimate the 3D position of the fire spots to plan the approximation goals for fire extinguishing.

It is not possible to estimate the 3D position of the fires from the fire position in pixel coordinates, as there is no information about the range. In order to obtain the distance to the fires, the LIDAR information is considered. The LIDAR values are mapped to the image coordinates given the known static transformation between sensors and the camera internal calibration. This way we can associate the fire object to a LiDAR range measure, obtaining an estimation of the distance to the fire. This is not always possible given the resolution of the LIDAR (and depending on the actual

<sup>7</sup> [https://github.com/robotics-upo/seekthermal\\_ros](https://github.com/robotics-upo/seekthermal_ros)

distance to the fire). If no range information is available, the range is estimated according to the 3D map of the environment. For the fires located on the floor, the height of the UAV is considered to estimate the position of the fire.

Each 3D measurement is associated with a covariance matrix that reflects the uncertainty on the range, which is adapted according to the estimated range (when no range is available, a large uncertainty is associated to it), and other aspects like the position of the robot. Then, an Information Filter is used to fuse in time the measurements about the position of the fire from different view points in order to triangulate the fire position, as in our previous work (Dias et al., 2015).

### 5.3. Fire extinguishing

This module is in charge of confirming the fire presence, aiming adequately the extinguisher system, and releasing the water or the blanket. The module assumes that the robot is already located at an attacking distance to the fire. Three tasks should be accomplished separately in each robot: confirming, centering and extinguishing.

#### 5.3.1. Fire extinguishing with UGV using water

The confirming phase is done by scanning with the Pan&Tilt unit onboard the UGV to look for the precise position of the fire. To this end, we command a predefined trajectory to the pan and tilt unit using position commands. At this point, the Fire Detection module is activated.

Whenever the Fire Detection module confirms a fire, the centering phase use a velocity visual servoing proportional controller to command the Pan&Tilt unit joints. The error signal is the pixel difference between the detected fire position on the image plane and the center of the locked-on target zone. We have empirically defined this zone in such a way that the water ejecting mechanism hits on target. This zone is indicated in an area of the image and it depends on both the height and distance to the target.

Once the extinguisher has locked on the fire, the extinguishing phase begins, and the system ejects water. As a final note, it is difficult to perceive the trajectory of the water with the onboard sensors. Thus, to increase the chances of successfully hitting the target, the pan&tilt unit performs a cross-shaped movement while ejecting water during 60 seconds.

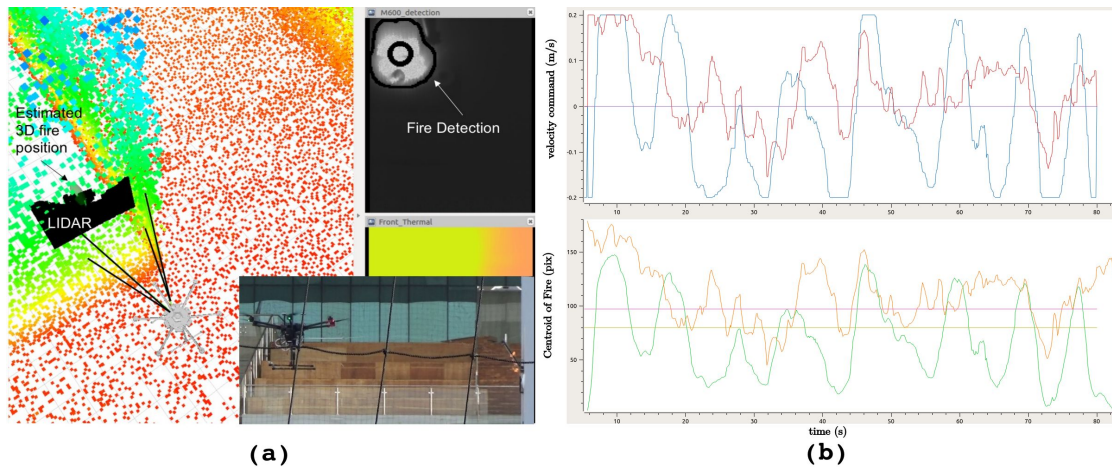
#### 5.3.2. Fire extinguishing with UAV using water

To extinguish the fires on the facade, once the UAV is positioned in the attack position, the UAV first performs a manoeuvre to confirm that there is a fire in the area (to compensate for imprecision on reaching the fire-attack point or on the estimated 3D fire position). This manoeuvre consists of performing a square motion parallel to the facade of the building.

As soon as a fire is detected the square sequence stops. We use a proportional velocity controller in the local Y and Z axes of the UAV so that the detected position of the fire on the image plane enters into the locked on target zone. As in the UGV case, we have empirically defined the locked on target area on the image assuming a distance of attack to the fire. Once the fire is on the locked on target zone, the water ejecting starts. As the UAV can be subjected to disturbances while flying which may cause the UAV move outside the lockup zone, the velocity controller keeps regulating the position during the ejection.

Figure 13 presents an example from the competition (during the Grand Challenge). It shows the output of our control system as a function of the coordinates of the centroid of the detected fire region on the thermal image. This control system sends velocity commands to the M600 autopilot in order to center the fire in the image. This particular fire was subjected to lateral wind gust of up to 8 m/s. It can be seen in Figure 13 how this provoked oscillations in the control in the lateral velocity, making it difficult to adequately point the water hose. The fire was never completely lost during the water ejection and the fire was considered partially extinguished.

For the fires inside the building, as before, first the drone starts a fire confirmation sequence. In this case, to prioritize the safety of the platform, the maneuver consists only of yaw rotations in



**Figure 13.** Facade fire extinction with an UAV using water. (a) The thermal camera is used to detect the fire. The 3D position of the fire is estimated by using the LIDAR information and tracked using an Information Filter. (b) Visual servoing for extinguishing. Top: commanded UAV speeds generated by the control system on  $y$  (horizontal UAV body axis, parallel to the wall, blue) and  $z$  (vertical UAV body axis, red). Bottom:  $u$  (columns, green) and  $v$  (rows, orange) image coordinates of the centroid of the fire detected on the image. The  $u$  and  $v$  target coordinates are plotted in pink and light green, respectively.

place. If a heat trace is found, then a proportional velocity controller in yaw is executed to center the fire horizontally on the image plane. When this is achieved the water ejecting starts, keeping the controller for correcting disturbances.

### 5.3.3. Fire extinguishing with UAV using blankets

In this case the UAV performs a confirming fire manoeuvre similar to the one done for facade wall fires, but this time the UAV moves in a plane parallel to the ground. As the blanket must cover the fire, the extinguishing action is more convoluted. Once the fire is confirmed on the image, the UAV is placed just above the fire. Then the drone descends to be more accurate when dropping the blanket. Afterwards, the UAV moves backwards until the fire is out of image and stop. Just in that moment, two electromagnets from the extinguishing system switch off to unroll the blanket. From there, the UAV moves forward passing through the fire, and so detecting the fire again with the thermal camera. The UAV will continue this movement, over passing the fire, and just when the thermal camera stop detecting the fire, the last four magnets will switch off and so the blanket drops covering the fire. From there, a fast move upward is performed in order to avoid any collision with the deployed blankets.

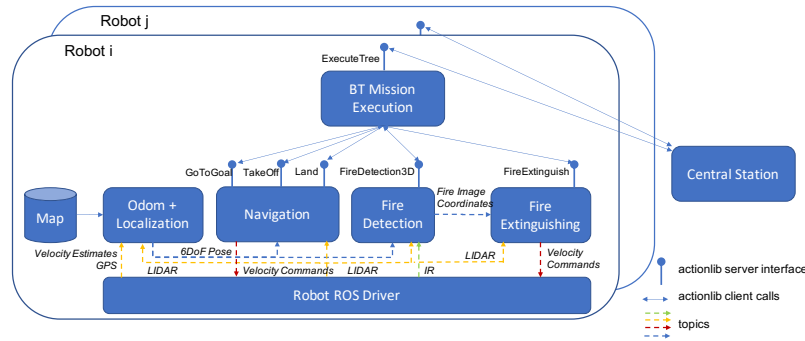
## 6. Mission Executive and Multi-Robot Coordination

### 6.1. System architecture and atomic tasks

All former components need to be integrated and work jointly to achieve the missions of the Challenge. This integration relies on the ROS framework, using the `ros-kinetic` distribution under Ubuntu 16.04. The algorithms described above have been implemented as ROS nodes.

The basic architecture of the system is shown in Figure 14. Continuous data flows between nodes are implemented through ROS topics. Some of these nodes also provide atomic tasks that will be composed to create the whole missions. As abstraction for such tasks, we employ the ROS' `actionlib` package<sup>8</sup>, which offers a simple interface to preemptable tasks. In particular, we use

<sup>8</sup> <http://wiki.ros.org/actionlib>



**Figure 14.** A depiction of the architecture. Each robot has operational autonomy to execute missions, and all processing takes place onboard, where all the enclosed modules enclosed run. A central Control Station can order the robots to execute missions through the BT Mission Execution node.

**Table 2.** Atomic tasks that can be carried out by each robot. Actionlib interface, parameters required, potential final status, and final result reported.

actionlib	Params	Final Status	Result
TakeOff	Height	CANCELLED if vehicle not ready SUCCEEDED if height reached ABORTED if height cannot be reached after takeoff	None
Land	None	CANCELLED if vehicle already landed SUCCEEDED landed ABORTED if the UAV cannot land	None
GoToGoal (Section 4)	Waypoint (WP)	CANCELLED if no path to WP SUCCEEDED if WP reached ABORTED if WP cannot be reached	None
FireDetection3D (Section 5.1)	Duration	SUCCEEDED if fire found within the given duration ABORTED if fire not found	3D position of fire
FireExtinguish (Section 5.3)	None	SUCCEEDED if fire can be locked on ABORTED if fire cannot be locked on	None

the task model offered by the SimpleActionServer implementation (actions can be ACTIVE when running; SUCCEEDED when successfully finished; CANCELED if they cannot be processed; ABORTED if they cannot be completed; finally, actions can be also PREEMPTED by the client). The list of atomic tasks considered by each robot are shown in Table 2. These tasks are the same for all robots, but of course they have specialized implementations depending on the robot (for instance, TakeOff is not applicable to the UGV, and the FireExtinguish action has the three variants described in Section 5.3).

### 6.2. Mission definition and execution

Given the atomic tasks that can be carried out by the modules described in the former sections, we need a framework to define whole missions as plans that combine those tasks, and a mechanism to execute and supervise the defined task plans. This mechanism is the mission executive, which is a fundamental component to allow the robots to have operational autonomy (Gancet and Lacroix, 2007; Molina et al., 2020). A popular example of executive is the ROS SMACH system (Bohren and Cousins, 2010), which uses nested Finite State Machines (FSM) to represent and execute high-level plans.

In our case, we employ Behaviour Trees (BTs) (Colledanchise and Ogren, 2017) as the framework for mission definition and execution for the robots of the team. They are an alternative to FSMs

that offer advantages in terms of modularity and reactivity. First originated as a tool for Non-Player Character (NPC) AI development in the video game industry, BTs are becoming widespread in robotics, mainly due to its modularity and simplicity (Colledanchise and Ögren, 2017; Paxton et al., 2017; Colledanchise and Ögren, 2018), including their use to define behaviours for UAVs (Ogren, 2012; Molina et al., 2020).

While a full description of the BT framework is out of the scope of this paper, here we summarize its main elements. BT model behaviors as hierarchical trees made up of nodes (an example can be seen in Figure 16). Trees are traversed from top to bottom at a given tick rate following a set of well-defined rules and executing the tasks/commands associated to the nodes that are encountered while doing so. Nodes' statuses are reported back up in the chain and the flow changes accordingly. A status can be either **SUCCESS**, **FAILURE** or **RUNNING**. According to their functionality, nodes can be classified as:

- **Composite**: it controls the flow through the tree itself and are similar to control structures in structured programming languages.
- **Decorator**: it processes or modifies the status it receives from its child.
- **Leaf**: this is where the actual task is performed, the atomic tasks that the robot can carry out, or other functionalities. As such, these nodes cannot have any children.

As it can be seen from the classification above, BT decouple logic from actual tasks in a natural way. When developing a tree, one only should care about the leaf nodes. In this case, these leaves correspond to the `actionlib` tasks described just above and another potential operations. The flow can later be defined and re-arranged constantly, creating new behaviors and expanding on what is already done. This *modularity* and *composability* (due to the hierarchical nature of the trees) of BTs with respect to alternatives like FSM is one of the advantages of the formulation (Colledanchise and Ögren, 2018), and was very relevant for defining the missions in the fleet and adapting to the lessons learned during the rehearsals. Missions could be re-defined very fast, and behaviours could be developed in parallel and easily integrated as sub-trees in more complex missions once well tested and validated.

The behavior of the individual robots of the team in the Challenge are designed using BTs. The mission executive that carries out and monitors the mission defined as a BT uses a BT engine. We have employed a publicly available BT implementation<sup>9</sup>. We have developed a ROS plugin for this engine. It ships with a library containing ROS subscribers, publishers, services and `actionlib` leaf nodes (mapping adequately the states of these actions to the corresponding states for the nodes BT) to be able to send and receive messages, and call services/actions to/from other nodes running in the network. And it includes a `BT Mission Execution` ROS node that offers an additional interface (as an `actionlib` as well) to load, start and stop user-defined behavior trees. This node (see Figure 14) runs onboard each robot, providing them with full operational autonomy to carry out complete missions.

### 6.3. Cooperation and Coordination

The approach to Challenge 3 takes advantage of the use of a heterogeneous robot team (mixing UGV and different types of UAVs). On the one hand, each robot specializes on one of the tasks for the Challenge. This means that we consider a loose cooperation between the robots. In particular, we devote the ground robot for the ground floor indoor fire, and each UAV to a different facade (with the larger M600 for the outside fires and the facade closest to the outside fires).

The Control Station module (see Figure 14) commands the different robots to execute out their missions, and those are carried out autonomously by each robot. Each robot runs all modules

<sup>9</sup> <https://www.behaviortree.dev>

**Table 3.** Results obtained during the MBZIRC 2020 competition. All tasks were achieved in autonomous mode. Maximum score for a task is by putting 1 litre or covering 100% of fire.

	<b>Outdoor Fires with blanket (Weights: 10 UAV, 5 UGV)</b>	<b>Facade Fires with UAVs (Weights: 14 ground floor, 8 first and second floor)</b>	<b>Indoor Fires (Weights: 24 second floor, 16 first floor, 10 ground floor)</b>	<b>Score</b>
TRIAL 1	1 fire, UAV. 50% covered	None	None	5
TRIAL 2	None	None	Ground floor, UGV. 350 ml	3.5
GRAND	N/A	Ground floor fire, under wind, UAV. 8 ml	Ground floor, UGV. 3 ml	0.13

onboard, including a different ROS master node. We use the software described in (Tardioli et al., 2019) to share between robots only vital information through the wireless network, including the position of each robot and the interfaces to control the robots from the station.

Thus, the main aspect to consider is the coordination in time and space to avoid collisions. The task allocation mitigates the risks, as the different vehicles will operate in different facades and, thus, different spaces. In any case, it is important to avoid close encounters between the vehicles. The takeoff zone is very narrow, so the Control Station regulates the order of initiating the missions for each vehicle so that no two UAVs take off at the same time. Furthermore, each vehicle has a landing zone that is separated from the rest of the team.

## 7. Experiments and Results

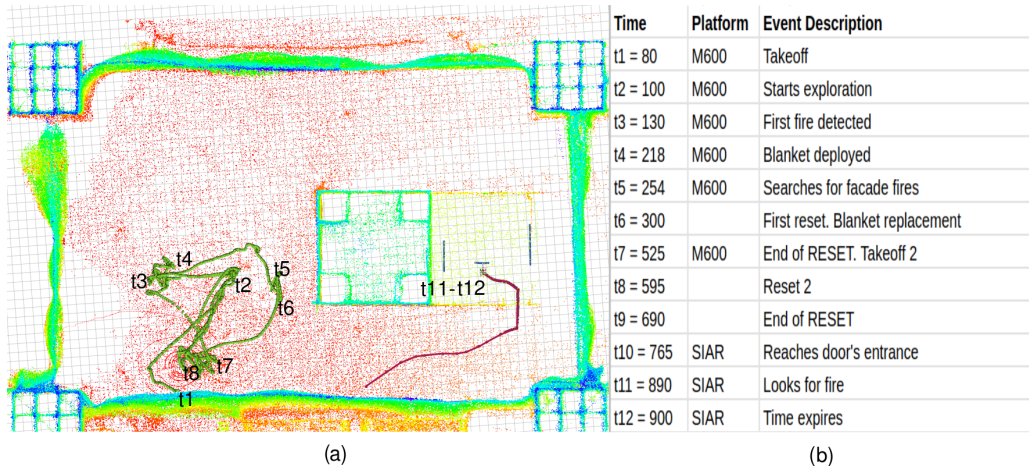
This section summarizes our performance during MBZIRC 2020. We focus on the results obtained by the integrated system as a whole summarizing high-level actions. For a detail on the performance of the individual systems, please refer to the previous sections.

The complete, integrated multi-robot system competed three times in Challenge 3. Twice during the main competition and once more for the final, Grand Challenge (please refer to Section 7.3). The result obtained in the competition are presented in Table 3. We scored in all trials (see Table 3), always in autonomous mode, and scored target hits on facade fires and both indoor and outdoor fires. We also managed to autonomously operate a team of three robots at the same time. With these results, we achieved the seventh place in Challenge 3 (out of 20 teams in the competition) and the fifth place in the Challenge 3 entry to the Grand Challenge (out of 17) in our first participation in the competition (contributing this way to an overall third place in the Grand Challenge for the team). The results reflect the complexity of the challenge. A video highlighting results from the MBZIRC competition can be accessed at <https://youtu.be/sx9R-6JrfQA>.

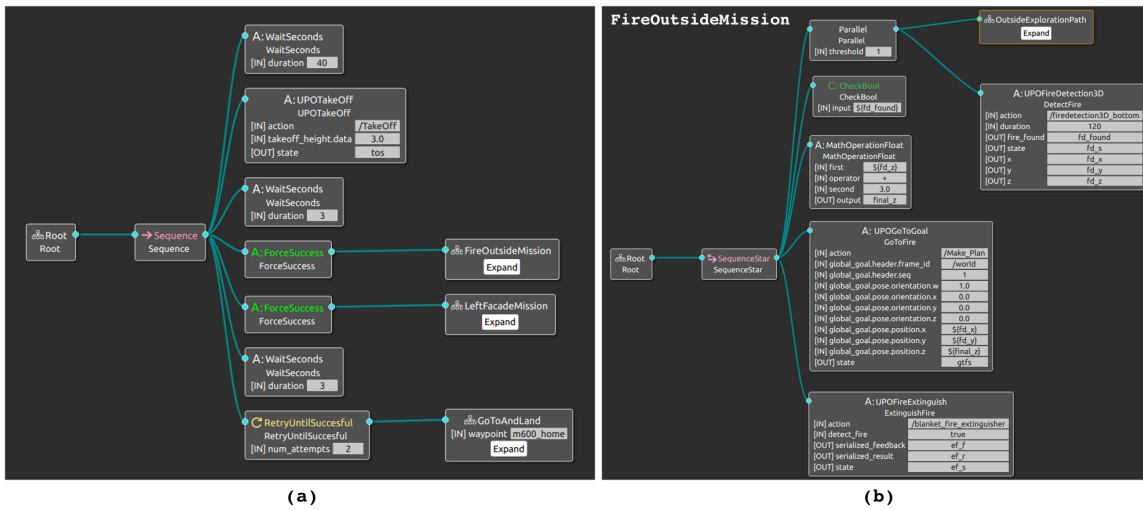
### 7.1. Trial 1

In this trial the ground robot SIAR was tasked to seek and extinguish an interior, ground-floor fire, while our M600 UAV was tasked to extinguish one outdoor fire with the blanket and then one facade fire with water. We encountered technical issues that prevented us from flying with the M210 platform. In contrast, our M600 platform did perform very well. It was able to successfully deploy a blanket which partially covered an outdoor fire. Figure 15 shows the timeline of the most important events during the first trial.

The mission definition for the M600 robot as a BT is presented in Figure 16. The UAV takes off (UPOTakeOff block) after a timeout (to wait for the ground robot to be at a safe distance;  $t_1$  in Figure 15). Then, the fire-blanket mission proceeds, represented by the FireOutsideMission sub-tree. Figure 16 (b), shows the elements of this sub-tree. First, the Parallel composite node executes in parallel a path to explore (OutsideExplorationPath, which uses the navigation actions of Section 4) and activates the UPOFireDetection3D action of Section 5.1. Once one of the two nodes returns SUCCESS (either because a fire is detected or the exploration path is finished), the Parallel composite returns SUCCESS (FAILURE otherwise).



**Figure 15.** (a) Trajectories of the 2 robots employed in Trial 1. The M600 UAV and the UGV SIAR trajectories are depicted in green and magenta, respectively. The labels indicate the timestamps when the main events occurred. (b) Task performed by each robot with the corresponding timestamp.

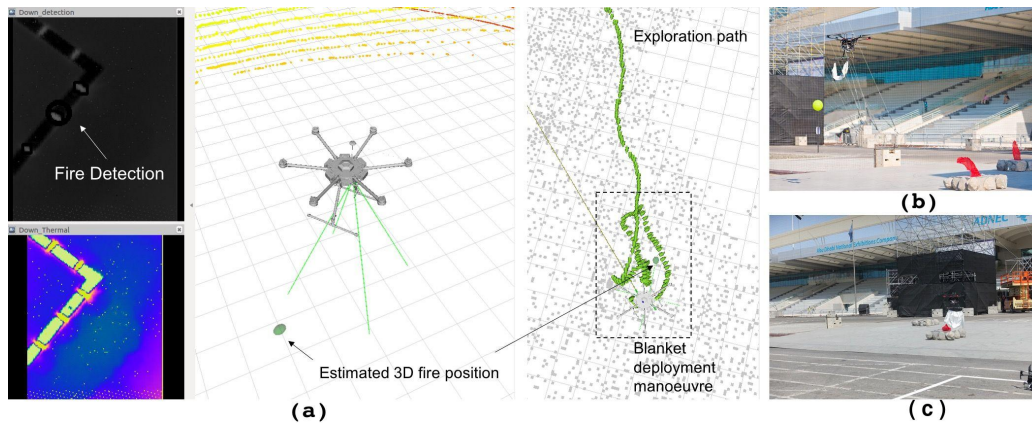


**Figure 16.** (a) General tree of a (simplified) BT for the M600 mission. The robot is tasked to put out an outside fire and then a fire on the left the facade of the building. The Sequence composite node executes the leaves from top to bottom. If a node returns FAILURE, the whole mission fails. Leaf nodes FireOutsideMission, LeftFacadeMission and GotoAndLand are sub-trees. (b) FireOutsideMission sub-tree.

In the M600 Trial 1 mission, the exploration path is followed until  $t_3$  in Figure 15. The leaf node CheckBool returns SUCCESS as the variable `fd_found` is TRUE. This means that a fire has been detected by the module, and its estimated coordinates are stored at  $(fd_x, fd_y, fd_z)$ . Then, the UAV is commanded to a fire-attack position 3 meters above the fire’s coordinates (computed using the MathOperationFloat leaf), by calling to the corresponding UPOGoToGoal (Section 4). Once there, the extinguishing procedure with the blanket is activated (UPOFireExtinguish, Section 5.3.3.3). The blanket is successfully deployed at time  $t_4$ , covering 50% of the fire. Figure 17 depicts the main steps of the fire blanket mission for the M600 UAV, as described in Figure 16.

After the blanket mission, and regardless of its outcome (determined through the ForceSuccess decorator), the M600 proceeds to the facade fire extinguishing sub-mission. This sub-mission





**Figure 17.** Data from one blanket deployment task of M600 UAV. (a) An exploration path is given and the methods of Section 5.1 are used to detect and to estimate the 3D position of the fire. Then, the blanket deployment manoeuvre is carried out. (b) Exploration phase. (c) Blanket deployment phase.

(LeftFacadeMission) is very similar to the FireOutsideMission described above (see Figure 16, b)), with different parameters. The main difference is that the exploration path comprises a set of waypoints that cover the entire facade, with the UAV oriented towards it. The second difference is that the frontal IR camera is used, and, when a fire is detected, a fire-attack waypoint is generated 2.5 meters in front of it. Finally, the UPOFireExtinguish task is called with the specialized version of Section 5.3.2.2.

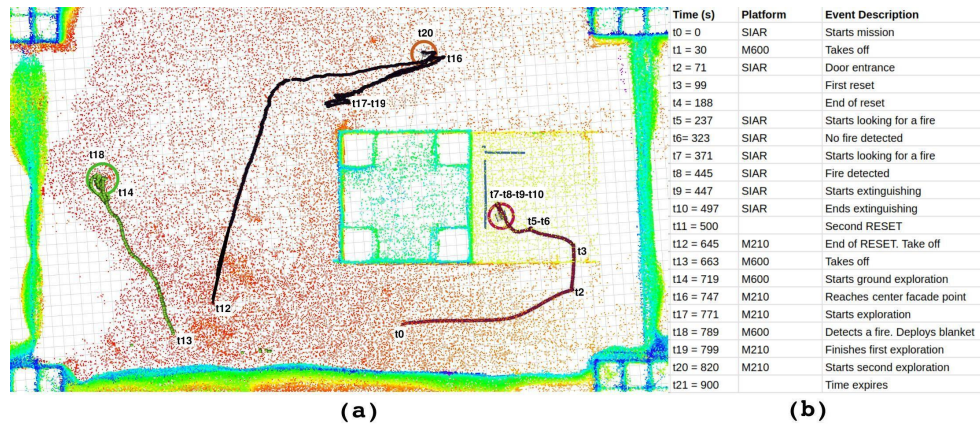
In normal conditions, after the exploration and regardless of the result of the facade mission (through the second ForceSuccess), the robot will go back to the home position. In this case, however, the mission was interrupted at  $t_6$ , when we asked for a reset to replace the deployed blanket and to start the UGV, instead of waiting for the full exploration of the facade.

After the reset (at  $t_7$ ), we assigned the M600 the same mission with a different exploration path to find the second fire outside, but the blanket was released accidentally just after take off, as the magnetic gripper failed, so we requested a second reset. The UGV SIAR proceeds to its mission (described in detail in Trial 2), but it can only explore one point inside the building ( $t_{11}$ ) before the trial time runs out.

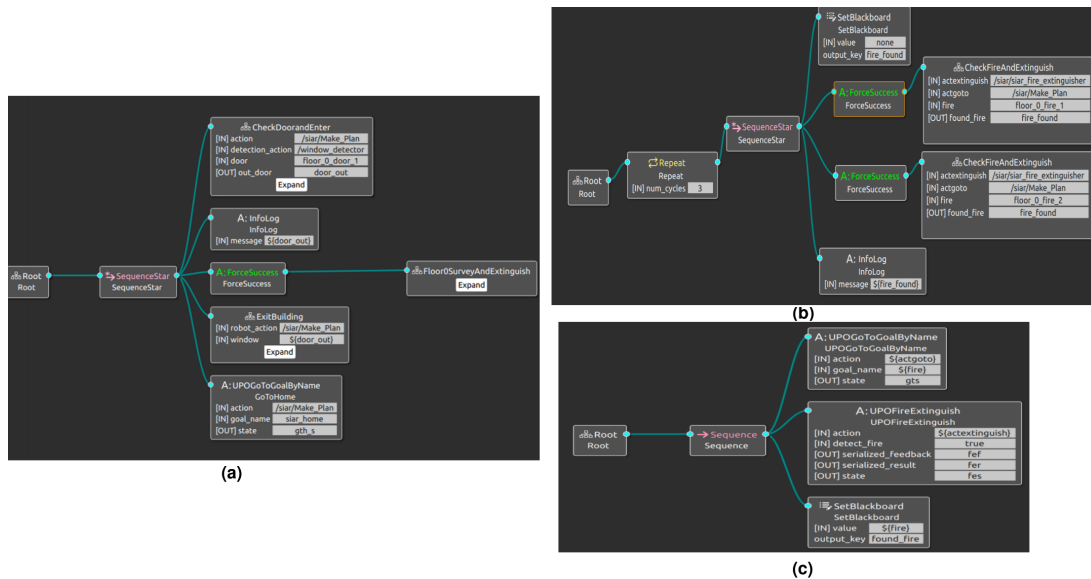
## 7.2. Trial 2

In Trial 2, we included our second UAV, M210, and tasked it to attack a fire in a different facade of the building. As before, the ground robot was tasked to look for and extinguish the interior, ground floor fire, and the M600 was tasked to extinguish one outdoor fire with the blanket and one facade fire with water. Figure 18 shows the general evolution of Trial 2.

After Trial 2 takeoff, the M600 unfortunately dropped its blanket due to a magnet malfunction. A reset was called and then the UGV SIAR was able to carry out its mission successfully. Figure 19 presents the mission of SIAR. First, the UGV is commanded to enter the building, which occurs at timestamp  $t_2$  in Figure 18. Then, the Floor0SurveyAndExtinguish submission is activated. It consists of inspecting two different points using the CheckFireAndExtinguish sub-tree (see Figure 19, b and c). When a inspection point is reached, the block UPOFireExtinguish is called. The onboard, infrared, thermal camera checks whether a fire exists using the procedure of Section 5.3.1. If no fire is present, the module returns FAILURE (this result occurs, correctly, as there is no fire at the first inspection point between times  $t_5$  and  $t_6$ ). Otherwise, first a centering action is performed and finally the fire extinguishing system is activated to deploy 1 litre (see Figure 20, a). This latter action occurs at the second inspection point (between times  $t_9$  and  $t_{10}$ ). Figure 20 shows the UGV extinguishing this indoor fire during Trial 2. We successfully dispensed only 350 ml



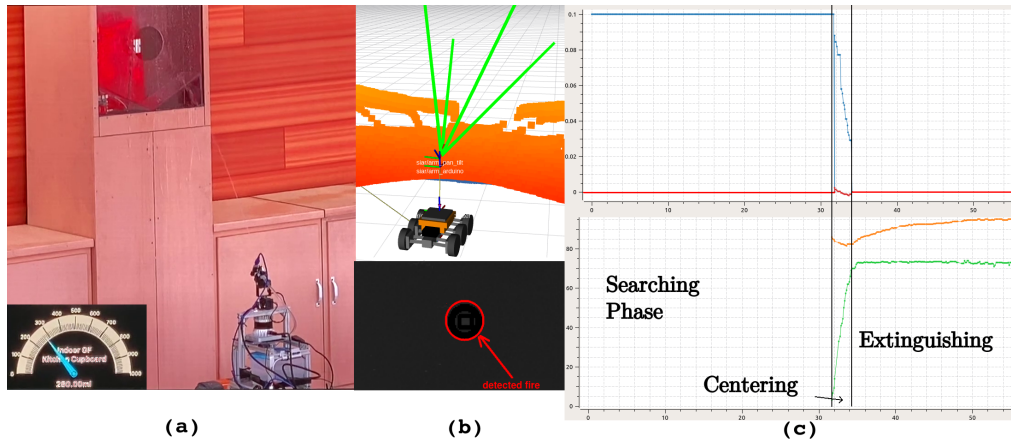
**Figure 18.** (a) Trajectories of the 3 robots employed in Trial 2. The M600 UAV was tasked with extinguishing an outdoor fire (green trajectory). The M210 UAV was commanded to extinguish a fire in the opposite facade of the building (black trajectory). The UGV SIAR was commanded to look and extinguish the indoor fire in the ground floor (red trajectory). (b) Task performed by each robot with the corresponding timestamp.




**Figure 19.** a) General tree of a (simplified) BT for the SIAR mission to put out one fire on the ground floor of the building. b) Subtree Floor0SurveyandExtinguish. c) Subtree CheckFireandExtinguish.

on target from the whole liter, because of a small misalignment of the water ejector and the thermal camera. The mission is not designed to deploy all the water on the UGV, allowing it to repeat exploring the environment for more fires and compensate for potential false alarms. This limitation out to be too conservative, as the system was actually capable of detecting the real fire present and no false alarms were created.

After the UGV put out the inside fire, a new reset is called (in time  $t_{11}$ ) to initiate the UAVs. The M210 takes off first at  $t_{12}$  and is sent to explore the farthest facade. 20 seconds later, at  $t_{13}$ , the M600 takes off. The M600 is tasked the same mission as in Trial 1, with a change in the order of the exploration path waypoints to add a safety buffer with the M210 UAV. In this case, the outdoor fire is again correctly detected and located, and the blanket is deployed (at  $t_{18}$ ), but misses the



**Figure 20.** Fire extinguishing action performed by the UGV in Trial 2. (a) Snapshot of the video recorded by the organization in the extinguishing action. The amount of water poured in can be seen below. (b-up) Internal state of the SIAR and FOV of the thermal camera. (b-bottom) Detected fire on the IR image. (c) The upper plot represents the pan and tilt velocity commands on blue and red lines, respectively. In the bottom plot, the image coordinates are plotted in light green and orange, respectively.



**Figure 21.** (a) Trajectories of the M600 (green and yellow) and SIAR (purple and black) in their first and second missions, respectively. (b) Task performed for each robot with the corresponding timestamp.

target by a few centimeters. We actually increased the height at which the deployment manoeuvre was performed from Trial 1 to 2 to add some safety buffer, and this ends up being detrimental. There is no further time to undertake the facade part of the M600 mission.

The M210 is able to explore part of the facade looking for the fires ( $t_{16}$  until  $t_{21}$ ). However, we ran out of time during Trial 2 before it could detect any fire. Please notice that this mission is the same as the M600 facade sub-mission described above, changing the exploration path waypoints. This comparison illustrates one of the main advantages of using BTs: Modularity allows reusing BTs from different platforms with minimal changes.

### 7.3. Grand Challenge

In the Grand Challenge, a reduced version of the three MBZIRC Challenges are carried out at the same time. Regarding Challenge 3, the fires to put out with a blanket are eliminated. There is also

a limit in the number of robots for tackling all challenges, and thus we could only use our M600 platform and the UGV. The total amount of time was raised from 15 to 25 minutes, but resets are global to all challenges.

Figure 21 shows the main events that occurred during the Challenge 3 entry to the Grand Challenge. The mission for the M600 drone is a stripped down version of the mission in Trials 1 and 2, eliminating the blanket deployment (again, this was very easily carried out given the composition of the associated BT). In contrast to the experiments in Trial 2, in this case our M600 drone manages to partially extinguish a facade fire. After a first reset at  $\tau_4$  (due to other of the three Challenges), the M600 follows an exploratory path until it detects an active fire with its thermal camera on front at  $\tau_8$ . As soon as it detects the fire, it switches into fire-centering mode and starts to eject water to extinguish it (see Figure 13). This particular fire is given larger weights for the scoring system as it was subjected to lateral wind gusts of up to 8 m/s.

Regarding the UGV platform, the mission is the same as before, with the main change of deploying the whole tank of water and returning to home position instead of searching for the remaining fires. The UGV does not perform as well as in the previous day, as the navigation system gets stuck when trying to enter in the building at timestamp  $\tau_7$ . Since requesting a reset affects and stops every challenge, it is necessary to wait for the opportune moment and coordinate the whole team to intervene the platform and restart its system.

Later, at time  $\tau_{12}$  we reinitialize the UGV and thus finally complete the second mission successfully. The UGV heads to the first inspection point, and then it executes the searching phase (Section 5.1) looking for fire at time  $\tau_{14}$ . As soon as it detects the fire (located at a different point than the previous day), the centering action is executed ( $\tau_{15}$ ) and the UGV is able to partially extinguish the fire on the ground floor. Unfortunately, the time expires only a few seconds after the start of the extinguishing procedure ( $\tau_{16}$ ), and thus only a small quantity of water is deployed on target.

## 8. Conclusions and Lessons Learned

We present in this article our entry to tackle MBZIRC Challenge 3: "Team of Robots to Fight Fire in High Rise Building", using a cooperative, multi-robot framework. Our system achieved the 7th place (same score as 6th, longer mission time) in Challenge 3 and third place in the Grand Challenge of the team. The difficulty of this Challenge can be seen in that only nine teams were able to score in autonomous mode during the two trials of Challenge 3, and only five teams during the Grand Finale trial of the same Challenge 3. We managed to score in autonomous mode in all three occasions. While the system and techniques presented in the paper were developed with the MBZIRC challenge in mind, we believe the localization, navigation and mission executive can be adapted and applied to other urban scenarios.

Nevertheless, our system did not prove as robust as we expected, and some other teams fared better. Some of the lessons learned are presented here:

### 8.1. Hardware matters

In a robotics competition, software obviously matters, but hardware selection and design decisions are key. Our use of off-the-shelf platforms was a good decision, in particular for the drones. Platforms we selected offered a ready to fly system with direct ROS integration. This choice opened the door to focus on the software-development tools required for the robot autonomy, significantly reducing the issues and problems with integrating hardware into drones. It is also worth noting that we experimented with DJI software issues before and during the competition: onboard firmware mismatch problems, random GPS no-fly area restrictions, and motor overheating. Nevertheless, all in all, we consider that the advantages of using a commercial off-the-shelf drone platform was a good decision and that it is, in general, preferable to avoid the design/implementation of custom drones in order to focus on autonomy problem. The platforms were adapted for the water-spraying and blanket-deploying systems. The results could have been improved by more robust hardware

adaptation. While all in all the water ejection mechanisms were adequate, pumping at higher pressure would have improved accuracy. Regarding the blanket deployment mechanism, while we were able to deploy the blanket on the fire once, it was affected by the air turbulence generated by the own UAV, which situation led to some blanket drops, as described above. Finally, our ground robot used a pan-and-tilt mount, for budget and weight reasons, whereas most teams used an arm. The use of an arm clearly helps in the final extinguishing maneuver.

## 8.2. Versatile localization

Using a map-based localization system allowed considerable versatility, even permitting us, finally, to discard GPS. This approach allowed fast system deployment, so that we did not need a precise GPS localization to localize the robot into the map, just the minimum required by the DJI drones to take off safely. The localization had a stable performance and a good precision during the competition, relying only on the use of LIDAR measures that were matched against a 3D map obtained *a priori*. It also was very important in order to be able to safely navigate in the scenario, as the localization of our platforms was performed taking into account the different elements on the stage, e.g. columns or safety nets, allowing us to locate them with good precision against the obstacles in the environment. The localization supported the robot autonomous navigation in the scenario, and the map-based system made the indoor-outdoor transition on localization seamless. However, only the ground robot entered the building. Even though we performed missions with the drones inside the building in simulation, finally we did not attempt to enter the building during the competition. This latter task was one of the most challenging, and no team was actually able to score inside the building during the competition, and only some manual trials were performed. Such missions clearly require techniques with a higher level of maturity.

## 8.3. Local operational autonomy

Robot local autonomy was also key. Each robot was able to carry out its mission with all processing onboard, the central station only in charge of launching the local missions of the robots in adequate order and relaying some global information. Even under communication dropouts, the robots were able to carry out their missions. The use of BTs as a mission executive on board the robots offered us also a great flexibility in designing and composing missions on the basis of the atomic tasks and sub-behaviours already defined. The BT engine also allowed the robots to react to the events detected during the runs.

## 8.4. Rapid and robust deployment

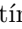
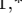
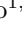

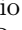
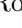
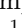
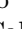
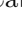
System robustness is critical to operating several robots at the same time. During trials, preparation time was very short. While we were able to perform missions with three robots at the same time, the system start-up was not robust enough to fully utilize the 15-minute mission limit. While the system can operate under communication dropouts, communication with the robots was needed to start them and, some times, delays in the starting sequence caused by communication problems prevented achieving more tasks in the mission. However, we feel that we are moving in the right direction and we look forward to subsequent competitions to improve the robustness and accuracy of our systems and to, hopefully, win the challenge!

## Acknowledgments

This work has been supported by the Spanish Ministry of Science, Innovation and Universities (COMCISE RTI2018-100847-B-C22, MCIU/AEI/FEDER, UE) and the Khalifa University under contract no. 2020-MBZIRC-10. The work of S. Martinez-Rozas is partially supported by the Antofagasta University.

UPO would like to acknowledge Paulo Alvito and the IDMind<sup>10</sup> company for the support with the ground robot hardware, Marcel Richter and the AeroGuillena airfield for the support during the tests, and Danilo Tardioli, from the University of Zaragoza, Spain, for his multi-master ROS software for multi-robot teams.

## ORCID

Simón Martínez-Rozas<sup>1,\*</sup>  <https://orcid.org/0000-0002-3289-3371>  
 Rafael Rey<sup>1,\*</sup>  <https://orcid.org/0000-0002-8708-9735>  
 David Alejo<sup>1,\*</sup>  <https://orcid.org/0000-0001-6664-6216>  
 Domingo Acedo<sup>1,\*</sup>  <https://orcid.org/0000-0003-3168-9173>  
 José Antonio Cobano<sup>1</sup>  <https://orcid.org/0000-0001-5094-6943>  
 Alejandro Rodríguez-Ramos<sup>2</sup>  <https://orcid.org/0000-0002-3257-4602>  
 Pascual Campoy<sup>2</sup>  <https://orcid.org/0000-0002-9894-2009>  
 Luis Merino<sup>1</sup>  <https://orcid.org/0000-0003-4927-8647>  
 Fernando Caballero<sup>1</sup>  <https://orcid.org/0000-0001-8869-2846>

## References

- Alejo, D., Cobano, J. A., Heredia, G., and Ollero, A. (2014). Collision-free 4d trajectory planning in unmanned aerial vehicles for assembly and structure construction. *Journal of Intelligent Robot Systems*, 73(1-4):783–795.
- Alejo, D., Mier, G., Marques, C., Caballero, F., Merino, L., and Alvito, P. (2020). Siar: A ground robot solution for semi-autonomous inspection of visitable sewers. In *Advances in Robotics Research: From Lab to Market*, pages 275–296. Springer.
- Amano, H., Osuka, K., and Tarn, T. . (2001). Development of vertically moving robot with gripping handrails for fire fighting. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, volume 2, pages 661–667.
- Ambrosia, V. G., Wegener, S. S., Sullivan, D. V., Buechel, S. W., Dunagan, S. E., Brass, J. A., Stoneburner, J., and Schoenung, S. M. (2003). Demonstrating uav-acquired real-time thermal data over fires. *Photogrammetric engineering & remote sensing*, 69(4):391–402.
- Ando, H., Ambe, Y., Ishii, A., Konyo, M., Tadakuma, K., Maruyama, S., and Tadokoro, S. (2018). Aerial hose type robot by water-jet for fire fighting. *IEEE Robotics and Automation Letters*, 3:1128–1135.
- Aydin, B., Selvi, E., Tao, J., and Starek, M. J. (2019). Use of fire-extinguishing balls for a conceptual system of drone-assisted wildfire fighting. *Drones*, 3(1).
- Bailon-Ruiz, R., Lacroix, S., and Bit-Monnot, A. (2018). Planning to monitor wildfires with a fleet of uavs. In *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 4729–4734. IEEE.
- Bohren, J. and Cousins, S. (2010). The smach high-level executive [ros news]. *IEEE Robotics & Automation Magazine*, 17(4):18–20.
- Bradshaw, A. (1991). The UK security and fire fighting advanced robot project. In *IEE Colloquium on Advanced Robotic Initiatives in the UK*.
- Casbeer, D. W., Kingston, D. B., Beard, R. W., and McLain, T. W. (2006). Cooperative forest fire surveillance using a team of small unmanned air vehicles. *International Journal of Systems Science*, 37(6):351–360.
- Chen, R., Cao, H., Cheng, H., and Xie, J. (2019). Study on urban emergency firefighting flying robots based on uav. In *2019 IEEE 4th Advanced Information Technology, Electronic and Automation Control Conference (IAEAC 2019)*, pages 1890–1893.
- Colledanchise, M. and Ögren, P. (2017). How behavior trees modularize hybrid control systems and generalize sequential behavior compositions, the subsumption architecture, and decision trees. *IEEE Transactions on Robotics*, 33(2):372–389.
- Colledanchise, M. and Ögren, P. (2018). *Behavior trees in robotics and AI: An introduction*. CRC Press.

<sup>10</sup> <https://www.idmind.pt>

- Conde, R., Alejo, D., Cobano, J. A., Viguria, A., and Ollero, A. (2012). Conflict detection and resolution method for cooperating unmanned aerial vehicles. *Journal of Intelligent Robot Systems*, 65:495–505.
- Cucci, D. A., Rehak, M., and Skaloud, J. (2017). Bundle adjustment with raw inertial observations in uav applications. *ISPRS Journal of Photogrammetry and Remote Sensing*, 130:1–12.
- Delmerico, J., Mintchev, S., Giusti, A., Gromov, B., Melo, K., Horvat, T., Cadena, C., Hutter, M., Ijspeert, A., Floreano, D., Gambardella, L. M., Siegwart, R., and Scaramuzza, D. (2019). The current state and future outlook of rescue robotics. *Journal of Field Robotics*, 36:1171–1191.
- Dias, A., Capitan, J., Merino, L., Almeida, J., Lima, P., and Silva, E. (2015). Decentralized target tracking based on multi-robot cooperative triangulation. In *IEEE International Conference on Robotics and Automation, ICRA*, page 3449–3455.
- Endres, F., Hess, J., Engelhard, N., Sturm, J., Cremers, D., and Burgard, W. (2012). An Evaluation of the RGB-D SLAM System. In *Robotics and Automation (ICRA), 2012 IEEE International Conference on*, pages 1691–1696.
- Fallon, M. F., Johannsson, H., and Leonard, J. J. (2012). Efficient scene simulation for robust monte carlo localization using an rgb-d camera. In *Robotics and Automation (ICRA), 2012 IEEE International Conference on*, pages 1663–1670. IEEE.
- Filippis, L. D., Guglieri, G., and Quagliotti, F. (2012). Path planning strategies for uavs in 3d environments. *Journal of Intelligent Robot Systems*, 65:247–264.
- Gancet, J. and Lacroix, S. (2007). Embedding heterogeneous levels of decisional autonomy in multi-robot systems. In *Distributed Autonomous Robotic Systems 6*, pages 263–272. Springer.
- Geiger, A., Ziegler, J., and Stiller, C. (2011). StereoScan: Dense 3D Reconstruction in Real-Time. In *Intelligent Vehicles Symposium (IV), 2011 IEEE*, pages 963–968.
- Gridseth, M. and Barfoot, T. (2019). Towards direct localization for visual teach and repeat. In *2019 16th Conference on Computer and Robot Vision (CRV)*, pages 97–104.
- Hornung, A., Wurm, K. M., and Bennewitz, M. (2010). Humanoid robot localization in complex indoor environments. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*.
- Imdoukh, A., Shaker, A., Al-Toukhy, A., Kablaoui, D., and El-Abd, M. (2017). Semi-autonomous indoor firefighting uav. In *2017 18th International Conference on Advanced Robotics (ICAR), Hong Kong, 2017*, pages 310–315.
- Innocente, M. S. and Grasso, P. (2019). Self-organising swarms of firefighting drones: Harnessing the power of collective intelligence in decentralised multi-robot systems. *Journal of Computational Science*, 34:80–101.
- Kanai, S., Hatakeyama, R., and Date, H. (2015). Improvement of 3d monte carlo localization using a depth camera and terrestrial laser scanner. *ISPRS - International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, XL-4/W5:61–66.
- Kelly, A. (1994). A feedforward control approach to the local navigation problem for autonomous vehicles. Technical Report CMU-RI-TR-94-17, Carnegie Mellon University, Pittsburgh, PA.
- Kerl, C., Sturm, J., and Cremers, D. (2013). Robust Odometry Estimation for RGB-D Cameras. In *In ICRA*.
- Kitt, B., Geiger, A., and Lategahn, H. (2010). Visual Odometry based on Stereo Image Sequences with RANSAC-based Outlier Rejection Scheme. In *Intelligent Vehicles Symposium (IV)*.
- Lin, Y. and Saripalli, S. (2017). Sampling-based path planning for uav collision avoidance. *IEEE Transactions on Intelligent Transportation Systems*, 18(11):3179–3192.
- Maza, I., Caballero, F., Capitan, J., Martínez-de Dios, J. R., and Ollero, A. (2011). Experimental results in multi-UAV coordination for disaster management and civil security applications. *Journal of Intelligent & Robotic systems*, 61(1-4):563–585.
- MBZIRC (2020). Mohamed Bin Zayed International Robotics Challenge - The Challenge 2020. <https://www.mbzirc.com/challenge/2020>.
- Merino, L., Caballero, F., Martínez-de Dios, J. R., Ferruz, J., and Ollero, A. (2006). A cooperative perception system for multiple uavs: Application to automatic detection of forest fires. *Journal of Field Robotics*, 23(3-4):165–184.
- Merino, L., Caballero, F., Martínez-De-Dios, J. R., Maza, I., and Ollero, A. (2012). An unmanned aircraft system for automatic forest fire monitoring and measurement. *Journal of Intelligent & Robotic Systems*, 65(1-4):533–548.
- Michael, N., Shen, S., Mohta, K., Kumar, V., Nagatani, K., Okada, Y., Kiribayashi, S., Otake, K., Yoshida, K., Ohno, K., et al. (2014). Collaborative mapping of an earthquake damaged building via ground and aerial robots. In *Field and service robotics*, pages 33–47. Springer.

- Mohta, K., Watterson, M., Mulgaonkar, Y., Liu, S., Qu, C., Makineni, A., and Kumar, V. (2018). Fast, autonomous flight in gps-denied and cluttered environments. *Journal of Field Robotics*, 35:101–120.
- Molina, M., Carrera, A., Camporredondo, A., Bavle, H., Rodriguez-Ramos, A., and Campoy, P. (2020). Building the executive system of autonomous aerial robots using the aerostack open-source framework. *International Journal of Advanced Robotic Systems*, 17(3):1729881420925000.
- Nash, S. K. and Tovey, C. (2010). Lazy theta\*: Any-angle path planning and path length analysis in 3d. In *Proceedings of the Twenty-Fourth AAAI Conference on Artificial Intelligence (AAAI)*, page 147–154.
- Ogren, P. (2012). Increasing modularity of uav control systems using computer game behavior trees. In *Aiaa guidance, navigation, and control conference*, page 4458.
- Pagliari, D., Cazzaniga, N. E., and Pinto, L. (2015). Use of Assisted Photogrammetry for Indoor and Outdoor Navigation Purposes. *ISPRS - International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, pages 113–118.
- Park, M., Lee, S., and Han, W. (2014). Development of lateral control system for autonomous vehicle based on adaptive pure pursuit algorithm. In *2014 14th International Conference on Control, Automation and Systems (ICCAS 2014)*, pages 1443–1447.
- Paxton, C., Hundt, A., Jonathan, F., Guerin, K., and Hager, G. D. (2017). Costar: Instructing collaborative robots with behavior trees and vision. In *IEEE International Conference on Robotics and Automation (ICRA)*, page 564–571.
- Pecho, P., Magdolenová, P., and Bugaj, M. (2019). Unmanned aerial vehicle technology in the process of early fire localization of buildings. *Transportation Research Procedia*, 40:461–468. TRANSCOM 2019 13th International Scientific Conference on Sustainable, Modern and Safe Transport.
- Perez-Grau, F., Caballero, F., Merino, L., and Viguria, A. (2017a). Multi-modal mapping and localization of unmanned aerial robots based on ultra-wideband and rgb-d sensing. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*.
- Perez-Grau, F., Caballero, F., Viguria, A., and Ollero, A. (2017b). Multi-sensor three-dimensional monte carlo localization for long-term aerial robot navigation.
- Perez-Grau, F., Ragel, R., Caballero, F., Viguria, A., and Ollero, A. (2018). An architecture for robust uav navigation in gps-denied areas. *Journal of Field Robotics*, 35(1):121–145.
- Pratt, J. and Manzo, J. (2013). The DARPA robotics challenge. *IEEE Robotics & Automation Magazine*, 20:10–12.
- Quigley, M., Gerkey, B., Conley, K., Faust, J., Foote, T., Leibs, J., Berger, E., Wheeler, R., and Ng, A. (2009). Ros: an open-source robot operating system. In *Proc. of the IEEE Intl. Conf. on Robotics and Automation (ICRA) Workshop on Open Source Robotics*, Kobe, Japan. <http://ros.org>.
- Royer, E., Lhuillier, M., Dhome, M., and Lavest, J.-M. (2007). Monocular vision for mobile robot localization and autonomous navigation. *International Journal of Computer Vision*, 74(3):237–260.
- Sanchez-Lopez, J., Wang, M., Olivares-Mendez, M., Molina, M., and Voos, H. (2019). A real-time 3d path planning solution for collision-free navigation of multirotor aerial robots in dynamic environments. *Journal of Intelligent Robot Systems*, 93:33–53.
- Schmid, K., Tomic, T., Ruess, F., Hirschmüller, H., and Suppa, M. (2013). Stereo vision based indoor/outdoor navigation for flying robots. In *2013 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 3955–3962.
- Schneider, F. E. and Wildermuth, D. (2017). Using robots for firefighters and first responders: Scenario specification and exemplary system description. In *2017 18th International Carpathian Control Conference (ICCC)*, pages 216–221. IEEE.
- Shen, R., Schwertfeger, S., and Visser, A. (2016). 16 years of robocup rescue. *KI-Künstliche Intelligenz*, 30:267–277.
- Shen, S., Mulgaonkar, Y., Michael, N., and Kumar, V. (2014). Multi-sensor fusion for robust autonomous flight in indoor and outdoor environments with a rotorcraft mav. In *2014 IEEE International Conference on Robotics and Automation (ICRA)*, pages 4974–4981. IEEE.
- Skeele, R. C. and Hollinger, G. A. (2016). Aerial vehicle path planning for monitoring wildfire frontiers. In *Field and Service Robotics*, pages 455–467. Springer.
- Tardioli, D., Parasuraman, R., and Ögren, P. (2019). Pound: A multi-master ros node for reducing delay and jitter in wireless multi-robot networks. *Robotics and Autonomous Systems*, 111:73–87.
- Thrun, S., Fox, D., Burgard, W., and Dellaert, F. (2001). Robust monte carlo localization for mobile robots. *Artificial Intelligence*, 128(1):99 – 141.



- TrinityCollege (2020). Trinity College International Fire Fighting Robot Contest. <https://trinityrobotcontest.org/>.
- Usenko, V., von Stumberg, L., Stückler, J., and Cremers, D. (2020). Tum flyers: Vision—based mav navigation for systematic inspection of structures. In *Bringing Innovative Robotic Technologies from Research Labs to Industrial End-users*, pages 189–209. Springer.
- Winfield, A. F., Franco, M. P., Brueggemann, B., Castro, A., Ferri, G., Ferreira, F., Liu, X., Petillot, Y., Roning, J., Schneider, F., et al. (2017). euRathlon and ERL emergency: A multi-domain multi-robot grand challenge for search and rescue robots. In *Iberian Robotics conference*, pages 263–271. Springer.
- Witt, R. R., Roff, A., Beranek, C. T., and Howell, L. G. (2020). Heat-detecting drones are a cheaper, more efficient way to find koalas, the conversation. <https://theconversation.com/heat-detecting-drones-are-a-cheaper-more-efficient-way-to-find-koalas-140332>.
- Yamada, Y. and Nakamura, T. (2016). Ummanned fire-fighting aerial vehicle to enable continuous water discharge. In *JSME Conference on Robotics and Mechatronics*.
- Zammit, C. and van Kampen, E. (2020). Comparison of a\* and rrt in real-time 3d path planning of uavs. In *Proceedings of the 2020 AIAA SciTech Forum*, pages 1–25.
- Zhan, W., Wang, W., Chen, N., and Wang, C. (2014). Efficient uav path planning with multiconstraints in a 3d large battlefield environment. *Mathematical Problems in Engineering*, page 1–12.
- Zhang and Singh, S. (2014). LOAM: Lidar Odometry and Mapping in Real-time. In *Robotics: Science and Systems Conference (RSS)*.
- Zhang, J. and Singh, S. (2017). Low-drift and real-time lidar odometry and mapping. *Auton. Robots*, pages 401–416.

**How to cite this article:** Martínez-Rozas, S., Rey, R., Alejo, D., Acedo, D., Cobano, J. A., Rodríguez-Ramos, A., Campoy, P., Merino, L., & Caballero, F. (2022). An aerial/ground robot team for autonomous firefighting in urban GNSS-denied scenarios. *Field Robotics*, 2, 241–273.

**Publisher's Note:** Field Robotics does not accept any legal responsibility for errors, omissions or claims and does not provide any warranty, express or implied, with respect to information published in this article.