

Autonomous Fire Fighting with a UAV-UGV Team at MBZIRC 2020

Jan Quenzel*, Malte Splietker*, Dmytro Pavlichenko, Daniel Schleich
Christian Lenz, Max Schwarz, Michael Schreiber, Marius Beul, and Sven Behnke

Abstract—Every day, burning buildings threaten the lives of occupants and first responders trying to save them. Quick action is of essence, but some areas might not be accessible or too dangerous to enter. Robotic systems have become a promising addition to firefighting, but at this stage, they are mostly manually controlled, which is error-prone and requires specially trained personnel.

We present two systems for autonomous firefighting from air and ground we developed for the Mohamed Bin Zayed International Robotics Challenge (MBZIRC) 2020. The systems use LiDAR for reliable localization within narrow, potentially GNSS-restricted environments while maneuvering close to obstacles. Measurements from LiDAR and thermal cameras are used to track fires, while relative navigation ensures successful extinguishing.

We analyze and discuss our successful participation during the MBZIRC 2020, present further experiments, and provide insights into our lessons learned from the competition.

I. INTRODUCTION

Robotic systems exhibit an immense potential for future applications in search and rescue scenarios [1] to aid first responders, overtake life-threatening tasks, or improve accessibility. For example, when firefighters had to retreat from within the *Notre-Dame de Paris* cathedral fire due to structural damage, an unmanned ground vehicle (UGV) took over [2]. Fire brigades increasingly use remotely controlled unmanned aerial vehicles (UAV) equipped with thermal and color cameras for monitoring tasks [3] and for guidance during extinguishing, even without direct line of sight or through smoke. However, latency or communication loss to a remote-controlled robot, e.g., inside a building, can further aggravate the disaster.

To address these issues, Challenge 3 of the Mohamed Bin Zayed International Robotics Challenge (MBZIRC) 2020 [4] targeted an urban firefighting scenario for autonomous robots to foster and advance the state-of-the-art in perception, navigation, and mobile manipulation. Participants had to detect, approach, and extinguish multiple simulated fires around and inside a building with up to 3 UAVs and one UGV. Each fire provided a 15 cm circular opening with a heated plate recessed about 10 cm on the inside. Holes on the outside facade were surrounded by a propane fire ring,

*: equal contribution.

This work has been supported by a grant of the Mohamed Bin Zayed International Robotics Challenge (MBZIRC) and the German Federal Ministry of Education and Research (BMBF) in the project “Kompetenzzentrum: Aufbau des Deutschen Rettungsrobotik-Zentrums (A-DRZ)”

Institute for Computer Science VI, Autonomous Intelligent Systems, University of Bonn, Friedrich-Hirzebruch-Allee 8, 53115 Bonn, Germany, {quenzel, ..., behnke}@ais.uni-bonn.de



Fig. 1: Team of fire fighting UAV and UGV at MBZIRC 2020.

while indoor fires had a moving silk flame behind a heating element.

In this work, we present our approach with an autonomous firefighting UAV-UGV team (see Fig. 1) to solve this challenge. In addition to describing and analyzing our integrated systems for solving the tasks, we discuss lessons learned and our technical contributions, including

- a precise LiDAR registration method that is applicable both for UAVs and UGVs,
- robust 3D fire localization using thermal cameras and LiDAR, and
- thermal servoing for navigation relative to a fire.

II. RELATED WORK

Cooperative monitoring and detection of forest and wild-fires with autonomous teams of UAVs [5] or UGVs [6] gained significant attention in recent years [7]. While UGVs can carry larger amounts of extinguishing agents or drag a fire hose [8], payload limitations impede the utility of UAVs. Hence, Aydin et al. [9] investigated the deployment of fire-extinguishing balls by a UAV. Ando et al. [10] developed an aerial hose pipe robot using steerable pressurized water jets for robot motion and fire extinguishing.

In urban environments, thermal mapping [11] is commonly used for building inspection and relies on simultaneously




Fig. 2: Hardware design of our fire-fighting UAV.

captured color and thermal images from different poses and employs standard photogrammetry pipelines. Real-time assistance for firefighters is provided by Schönauer et al. [12] via thermal augmentation of live images within room-scale environments.

In contrast, Borrmann et al. [13] mounted a terrestrial LiDAR as well as a thermal imager and color camera on a UGV to obtain colorized point clouds, enabling automated detection of windows and walls [14]. Fritsche et al. [15] cluster high-temperature points from fused 2D LiDAR and mechanical pivoting radar to detect and localize heat sources. Similarly, Rosu et al. [16] acquire a thermal textured mesh using a UAV equipped with LiDAR and thermal camera and estimate the 3D position and extent of heat sources.

New challenges arise where the robots have to operate close to structures. UAVs and UGVs are often equipped with cameras and are remote-controlled by first responders. In contrast, autonomous execution was the goal for Challenge 3 of the MBZIRC 2020.

Team Skyeeye [17] used a 6-wheeled UGV and DJI Matrice 600 and 210 v2 UAVs equipped with color and thermal cameras, GPS and LiDAR. A map is prebuilt from LiDAR, IMU, and GPS data to allow online Monte Carlo localization and path planning with Lazy Theta*. Fires are detected via thresholding on thermal images. The fire location is estimated with an information filter from either projected LiDAR range measurements or the map. A water pump for extinguishing is mounted on a pan-tilt-unit on the UGV while being fixed in place on the UAV.

Although our general approach is similar to Team Skyeeye, we rely more heavily upon relative navigation for aiming the extinguisher at the target after initial detection and less on the quality of our map and localization. In comparison, the nozzle’s placement on the end effector of the robot arm on our UGV gives us a better reach, and the protruding nozzle on our UAV allows for a safer distance from fire and wall while aiming. Furthermore, where common distance measurements are unavailable, we use the known heat source size for target localization. We detect the holes of outdoor facade fires in LiDAR measurements and fuse these with thermal detections, which allows us to spray precisely on target, where thermal-only aiming would be offset by the

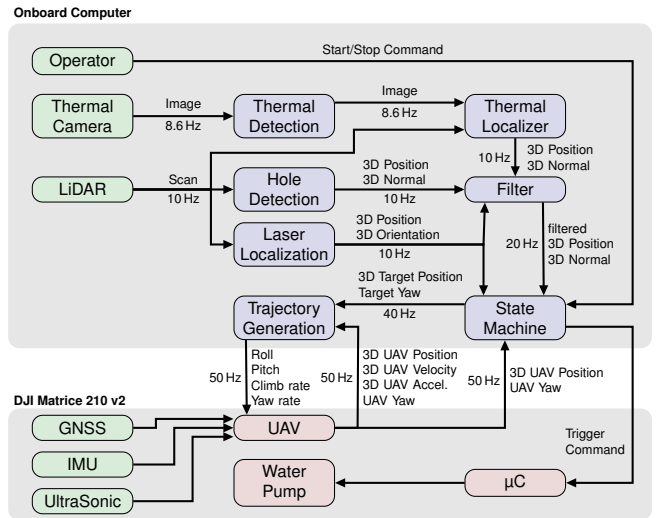


Fig. 3: UAV system overview.

surrounding flames.

III. UAV SOLUTION

We decided to build upon our experience from the first MBZIRC competition [18] and equipped commercially available DJI UAVs with sensors and computing power for easy adaptation to task-specific requirements. The limited payload of the previously used DJI M100 did not allow transporting sufficient water for extinguishing the fires. Hence, we upgraded to the more powerful DJI Matrice 210 v2.

A. Hardware Design

Fig. 2 shows the setup of our UAV. It is equipped with an Intel Bean Canyon NUC8i7BEH with Core i7-8559U processor and 32 GB of RAM. For perception and localization of the fires, we combine an Ouster OS1-64 LiDAR and a FLIR Lepton 3.5 thermal camera with 160×120 px resolution. The M210 provides GNSS-localized ego-motion estimates. We had to deactivate the obstacle avoidance in the forward direction to get close enough to fires for extinguishing.

Our water supply is stored in two downward-facing 1.5 L PET-bottles attached between the rear frame arms and the landing gear. The screw caps are connected via a flexible hose to a windscreen washer pump attached on top of the landing gear. We trigger the pump via an Arduino Nano microcontroller. We mounted a carbon tube on top of the UAV as an extended nozzle. It is supported by two additional carbon struts from the landing gear. The high location compensates for the height difference in the water jet parabola, allows to perceive the fire with the sensors below, and prevents the water from flowing out on its own during flight maneuvers. We chose a 10° downturn for the nozzle, LiDAR, and thermal camera to fly above the fire while maintaining observability. Fig. 3 gives an overview of our UAV software.

B. Perception

1) *Laser Localization*: GNSS-based localization is subject to local position drift and unreliable close to structures.




Fig. 4: Detected hole (red disk) inside a point cloud. The lower-right corner shows the points of the left plane projected into the virtual camera after morphological transformation as well as the detected circle.

The usage of RTK-/D-GPS was allowed but penalized and thus not used by us. Hence, we localize our UAV w.r.t. the building using LiDAR. In a test run, we collected data to generate a Multi-Resolution Surfel Map (MRSMap) of the arena [19]. We predefined the origin of our "field" arena coordinate system in GPS coordinates and transformed the map accordingly. During localization, we register the current LiDAR point cloud against the map. The UAV provides an estimate of its ego-motion from onboard sensors, which we use to update the initial pose. After registration, we update an offset transformation that compensates for the local position drift of the GPS. The translation between consecutive offset updates is bounded to be below 30 cm to prevent potential jumps from incorrect registration. Although the UAV ego-motion estimate is reliable in the short term, a continuous update is necessary. We experienced incorrect ego-motion estimates when approaching the building at specific heights. This issue is likely due to the ultra-sonic sensor measuring the distance to the building rather than the height above the ground. All subsequent navigation poses are expressed in the localized map frame.

2) *Hole Detection*: The water jet exiting the nozzle is very narrow; thus, precise aiming is required in order to extinguish a fire with a limited amount of water. As the global tracking accuracy is not sufficient for this task, we apply a hole detection algorithm on the LiDAR point clouds for relative navigation. First, we extract planes using RANSAC. For every such plane, we project the contained points into a virtual camera perpendicular to the plane. We apply morphological transformations to the resulting image in order to close gaps between projected points. Shape detection is used on the image to find potential holes. After re-projecting the circles into 3D, overly large or small holes are discarded, as the scenario specifies a constant 150 mm diameter for all target holes. An example is shown in Fig. 4.

3) *Thermal Detection and Heat Source Localization*: Our thermal detector applies lower and upper bounds to threshold the intensity. We find contours in the thresholded image. For

each contour, we compute the area, bounding box, center of intensity, as well as min, max, and mean intensity for further processing. Only contours of a specific size are retained.

The hole detection provides a reliable estimate at close range (<3 m) but fails to detect the opening from multiple meters away where the heat source is already visible in the thermal image and we have distance measurements on the wall. Hence, to obtain an approximate heat source position, we project the LiDAR point cloud into the thermal image. We use all points falling into the bounding box of the detected heat contour and calculate their mean position and normal from the covariance. Since hole and thermal detections on an equivalent heating element were very stable, we used them to calibrate the transformation between LiDAR and thermal camera. We recorded a heat source from multiple positions and minimized the reprojection error between the center of intensity of the heat source and the projected mean position of the detected hole. The resulting accuracy is sufficient to successfully approach the target.

4) *Filter*: Our thermal and hole detector both output lists of detections $d_i = (p_i, n_i)$, consisting of a position p_i and a normal n_i . A filtering module processes these detections to reject outliers and combines both detection types into an estimate of the position and orientation of the currently tracked fire. To do this, we keep track of a history $\mathcal{H} = ((p_1, n_1), \dots, (p_{10}, n_{10}))$ of ten recent valid detections and estimate the target position $p^* = \frac{1}{|\mathcal{H}|} \sum_{i=1}^{|\mathcal{H}|} p_i$ and normal $n^* := \frac{1}{|\mathcal{H}|} \sum_{i=1}^{|\mathcal{H}|} n_i$ as running averages over the detection history. Mind that the detection history may contain both thermal and hole detections. Thermal detections are necessary to determine which of the possible targets is currently on fire. However, we found that hole detections give a more precise estimation of the target position and especially of its orientation. Hence, we use thermal detections to initialize the target estimate and subsequently update it using hole detections if available. In the following, we detail the initialization process and the estimate update during target tracking:

a) *Target Initialization*: For the initialization process, we determine the feasibility of thermal detections by checking the angle between the detected normal and the line-of-sight from the UAV position to the detected fire position. If this angle exceeds 45° , the detection is rejected. The remaining detections are collected until ten out of the last twenty detected heat positions lie within a distance of 1 m. These ten detections are used to initialize the detection history \mathcal{H} more robustly against outliers.

b) *Target Tracking*: We filter thermal, and hole detections based on the same feasibility checks as above: detections for which the angle between the detected normal and the line-of-sight exceeds 45° are discarded. To remove outliers, we additionally only consider detections which positions lie within a ball of radius 1 m around the current position estimate and for which the angle between the detected normal and the estimated normal is lower than 45° .




Fig. 5: Flowchart of the UAV state machine.

Although we keep track of the most recent thermal detection, we only add it to the detection history \mathcal{H} if there has not been a valid hole detection within the last second. Thus, we ensure that we are still able to estimate target positions if the hole detector fails, but otherwise only use the more precise information from hole detections. In the case of multiple holes close to the target position, the estimate might drift away from the target if we only use hole detections. Furthermore, we have to recover from situations in which the initial heat detections were wrong or in which the fire has been extinguished in the meantime. To address these issues, we only add hole detections to the detection history \mathcal{H} if there has been a heat detection within the last second and the detected hole position lies within a radius of 1 m around the latest heat detection. If no detection was added to the history for more than 2 s, tracking times out, and we re-initialize the target estimate.

C. Control

We build upon our experience from the MBZIRC 2017 and reuse the method of Beul and Behnke [20] for trajectory generation and low-level control. The high-level control of the UAV is performed by a Finite State Machine (FSM)¹. The FSM uses inputs from components described above to produce navigation waypoints to locate and approach fires, as well as to control the water spraying during extinguishing. It also ensures that the UAV stays within arena limits and altitude corridor. The diagram of the FSM is shown in Fig. 5.

Search state: The UAV flies around the building in order to detect and localize a fire. The route around the building is defined as a linked list of waypoints. Each waypoint is defined in the building frame, obtained from laser localization. The route forms a convex hull around the building’s XY plane. Thus, transfer between two consequent waypoints can be safely performed without additional planning for collision avoidance, simplifying the overall system. While moving between waypoints, the detection and filter pipeline collects data. There are two types of waypoints: *Transfer* and *Observation*. Upon arrival at a *Transfer* waypoint, the UAV proceeds to the next waypoint immediately. When arriving at an *Observation* waypoint located in front of the known hole locations, the UAV hovers for a predefined duration of 2 s, looking for fires. The UAV switches into the *Extinguish* state if at any time a detection satisfies the following conditions:

- 1) The detection was observed at least five times in a row,
- 2) the detection Z coordinate lies within the allowed altitude interval, and

¹Open-source: https://github.com/AIS-Bonn/nimbro_fsm2




Fig. 6: Hardware design of our fire fighting UGV.

- 3) the angle between UAV heading and detection heading does not exceed the predefined 45° limit.

Extinguish state: The purpose of this state is to arrive at the extinguishing position without losing the detected fire on the way. Given the target from the previous state, the navigation goal is set at a predefined offset from the fire in the direction of its estimated normal (2.1 m away and 0.35 m above). This offset was determined empirically to suit the parabola of our water spraying system. The goal heading of the UAV at each time step is modified by adding an angle between the current heading of the UAV and the heading of the fire. By doing so, we try to continuously keep the fire in the center of the thermal camera’s field of view during the approach. This state performs relative navigation, with the UAV moving relative to the currently observed fire without relying on laser localization. This allows to achieve a precise placement of the UAV relative to the fire, which is necessary for consequent extinguishing. Once the UAV reaches the goal pose, the pump is started. It is stopped as soon as UAV deviates from the current goal pose more than a predefined threshold for either position or heading. The UAV keeps track of the fire and updates the goal pose accordingly. If the UAV does not receive new detections of the fire for longer than 5 s, the extinguishing is aborted, and it switches back to the *Search* state. If the UAV did not lose the target while extinguishing, the pump is stopped after all water has been sprayed. In that case, the UAV returns and hovers above the start position, waiting for manual landing and refueling of the water storage.

For additional safety, at all times, the building frame from laser localization is continuously tracked. If a sudden jump of more than 1 m is detected, the UAV immediately transitions into the *Stop* state, aborting any ongoing operation, and hovers at its current position until a safety pilot takes over control.

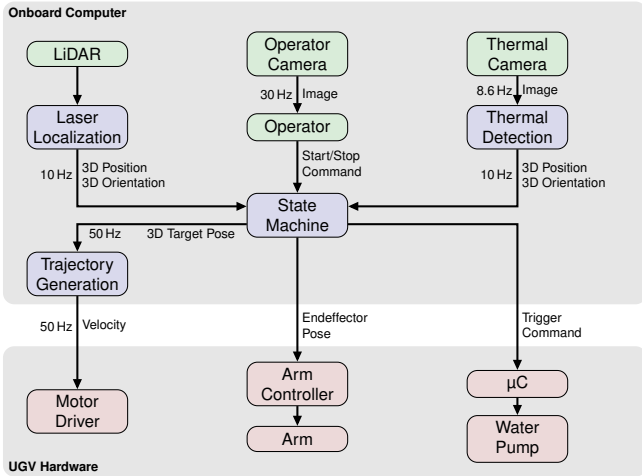


Fig. 7: UGV system overview.

IV. UGV SOLUTION

We mainly developed our ground robot Bob to solve Challenge 2 (pick and place of brick-shaped objects) [21] based on our very successful UGV Mario, which won the MBZIRC 2017 competition [22]. We added a thermal camera, water storage, and pumps to solve the firefighting challenge.

A. Hardware Design

The UGV hardware includes a four-wheeled omnidirectional base, a 6-DoF Universal Robots UR10e arm with a Velodyne VLP-16 3D LiDAR and a FLIR Lepton 3.5 thermal camera mounted on the wrist. A water storage for up to 10L, and two windscreen washer pumps, are installed in the robot base (see Fig. 6). A standard ATX mainboard with a quad-core Intel Core i7-6700 CPU and 64 GB RAM enables onboard computation. The robot is powered by an eight-cell LiPo battery with 20 Ah and 29.6 V nominal voltage for roughly one hour operation time. The robot's base has a footprint of 1.9×1.4 m and is driven by four direct-drive brushless DC hub motors. Each wheel is attached via a Dynamixel H54-200-S500-R yaw servo to enable omnidirectional driving capabilities. The two pumps are connected in series. This setup generates enough power to spray roughly 1 L of water in 10 s over a height difference of up to 1.5 m.

B. Perception

For heat source detection, we use the same algorithms as for our UAV. The simulated fires on the ground level are placed inside an acrylic glass enclosure, which brings new challenges: Firstly, only a fraction of the infrared rays penetrate the walls, so the thermal detection works only through the hole in the front, which results in a very narrow window for detection. Secondly, the LiDAR hardly ever perceives transparent acrylic glass, so hole detection cannot be reused for determining the relative position of the heat source. Assuming the heating element has reached target temperature and is observed at a nearly orthogonal angle through the hole, the bounding box of the thermal

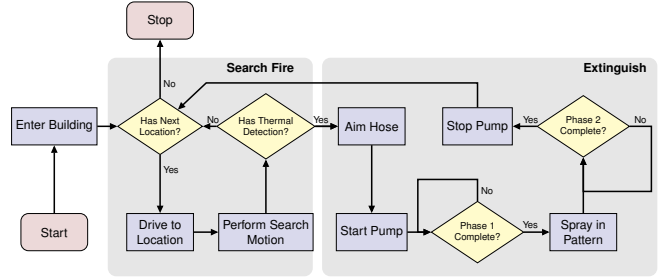


Fig. 8: The flowchart of our UGV state machine.

detections can be used to estimate the distance. We used an equivalent heating element for calibrating the measurements. Back projecting the detection's center point with this distance results in a 3D position estimate of the heat source.

We reused a modified version of the laser-based localization from Sec. III-B.1 for our UGV. Instead of estimating a correction transformation from the GPS-based "field" frame to the "map" frame, we calculate the offset transformation between the "field" origin and the "odometry" frame without GPS. The wheel encoder odometry is reset each time at startup but is accurate enough to be used as initialization during incremental scan matching.

The map is combined from two separately recorded segments a) from the start area to the kitchen entrance, and b) from a round trip inside the kitchen. Both segments are mapped as before and then aligned. While approaching the building, we use the outdoor map and seamlessly switch to the indoor map before passing the doorway.

C. High-level Control

Fig. 7 gives a system overview. Details for the central state machine are shown in Fig. 8. The UGV navigates from the start zone to the inside of the building using a set of static waypoints. The scenario specifies exactly two static fire locations within the kitchen, one of which is burning, which we utilized by driving the UGV to pre-determined positions roughly in front of the fire and performing a rectangular search motion with the arm. When a heat source is detected, the robot starts the extinguishing procedure. If not, it proceeds to the next location.

The nozzle is aimed so that the water jet will directly hit the heating element (the trajectory was determined experimentally). Since the relative position estimate of the heat source described in Sec. IV-B might differ from our calibration and vary with distance, the aiming process is performed in multiple 10 cm arm motions to incorporate new measurements. To further increase the chances of successfully extinguishing, we split the actual water spraying into two phases. For the first half of the assigned water, the nozzle is kept still. In the second phase, the remaining water is sprayed in an hourglass-shaped pattern by tilting the nozzle. During the Grand Challenge run, the first phase already delivered a sufficient amount of water to extinguish the fire. Even though we never observed the thermal detector producing any false-positives during our runs, the robot will continue proceeding to the second fire after extinguishing the

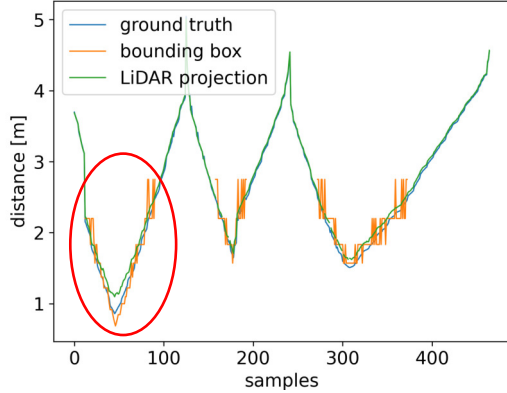


Fig. 9: Comparison of distance estimates from bounding box (Sec. IV-B), LiDAR projection (Sec. III-B.3), and real distance.

first one. The water capacity is more than sufficient, with approximately 4 L per fire.

V. EVALUATION

The two variants for estimating the position of the heat source described above differ in their application range. The LiDAR projection approach explained in Sec. III-B.3 is required for distances above 3 m, while the bounding box method described in Sec. IV-B works best for distances smaller than 2 m. Fig. 9 depicts a comparison of both methods to the ground truth distance obtained by fitting a plane to the test wall. The red encircled part indicates the benefits of the bounding box method for close distances. Estimates below 1 m would be even more accurate, but the minimum distance of the LiDAR for acquiring ground truth data does not permit this measurement. The large jumps in the bounding box estimates above 2 m originate from the discrete nature of the bounding boxes in the low-resolution thermal images. This figure does not compare the quality but shows that both variants perform well in their designated operating range.

A. UAV

On the first scored challenge run, software issues and wrong predefined search poses prevented the UAV from detecting any fire. Similarly, we experienced incorrect height estimates on the second challenge day. We attribute this to the ultrasonic sensor measuring the building wall rather than the ground, thus estimating the height too low. After the Grand Challenge, we also noticed that during all trials, the LiDAR localization was disabled. As a result, the UAV flew too high to detect the fire. We then switched to manual mode and were able to fill the container of the windy ground-level facade fire with 322 mL of water. The manually flown trajectory and localized detections of the fire are shown in Fig. 10.

Shortly after the Grand Challenge started, a series of unfortunate circumstances led to the UAV crashing into the building, as shown in Fig. 11. The laser localization was disabled due to human error, and position estimates




Fig. 10: Fire detection and UAV localization during manual fire extinguishing of the windy ground level facade fire during the second challenge day. The trajectory is color-coded by time (yellow to blue).

relied solely on the GNSS-based ego-motion estimates from the UAV. To compensate for position drift, we added a static offset to the GPS poses, which was calculated as the difference between the predefined start position and the GPS pose at the time of the challenge start. However, the GPS signal strongly drifted while computing the offset. Thus, a large incorrect offset was added to the GPS pose, resulting in an initial localization error of around 20 m. Moreover, it continued to drift afterwards. The canyon-like starting position between the building and the operator booth, as well as interference with the UGV, might be responsible for the drift.

After repair, we performed further tests to showcase our UAV's ability with the mockup shown in Fig. 12. After lift-off, the UAV flies towards multiple predefined search waypoints (Fig. 12 left, Poses 1-3). The heat source is first detected at Waypoint 3. Now the UAV begins to navigate relative to the detection and turns towards it before flying closer (Pose 4) and extinguishing the fire (Fig. 12 middle and right). After extinguishing, the UAV proceeds to Pose 5.

B. UGV

Fig. 13 shows an overview of the UGV's behavior during the scored runs. On the first challenge day, the robot could not compete due to software issues. During the second challenge run, the autonomy worked well to the point where the UGV tried to position the arm for scanning the first




Fig. 11: Incorrect GNSS-based UAV localization during the Grand Challenge. The laser localization shows the trajectory before crashing into the building, color-coded by time (yellow to blue).




Fig. 12: Autonomous fire extinguishing on mockup target with heating element. Red line segments indicate the UAV’s heading at various numbered waypoints. When the heat source is detected [c)], the UAV approaches [orange rectangle, a)] and starts to extinguish [b-d)].

location for heat signatures. Because the manually configured approximate fire location was wrong (c.f. Fig. 13 left), the arm collided with the target, which triggered an emergency stop. While manually maneuvering the robot out of the building during a reset, the not sufficiently tested motor controller malfunctioned, and the robot drove into a cupboard (light-pink trajectory). After resetting to the start zone, the faulty controller caused the robot to drive towards the arena boundary (magenta trajectory), even though localization and waypoints were correct. Near the end of the run, we switched to manual mode. With arm motions and a pump macro prepared in advance, the robot managed to spray 470 mL water into the container just before the timeout.

After analyzing the collected data, we found localization inaccuracies while entering the building (see Fig. 13 left) due to the proximity of inside and outside points in the map. After creating separate maps and switching on driving inside, the localization performed accurately.

During the Grand Challenge run, the UGV successfully solved the task autonomously. There was an issue with a disconnected tf-tree, which led to the state machine stopping in front of the first fire. After a reset, we solved the issue, and the UGV successfully performed as intended: it stopped in front of the first location, scanned for the heat source, aimed and sprayed sufficient water to achieve the full score. Fig. 13 right shows driven trajectories during the first (purple) and second (magenta) attempt. The deviation after the first waypoint is due to an edge on the uneven ground; otherwise, the overlap of planned- and driven trajectories demonstrates our method’s high accuracy.

A video showcasing the evaluation can be found on our website²

VI. LESSONS LEARNED

Robotic competitions not only allow evaluation of the performing systems but also represent the opportunity to discover systematic problems and to learn from these.

As demonstrated by our problematic runs during the competition, it is highly challenging to develop such a complex system to a robustness level such that it immediately performs well in competition situations. Of course, having more testing time *with a similar environment setup* as well

as having proper scenario specifications early on with a limited number of sub-challenges would have helped this point immensely.

While we tested our UAV system with a mockup in our lab and on an open field, the competition environment was unexpectedly dissimilar and led to system failures. For example, GPS visibility was severely limited near the mandatory start position, which led to initialization errors. An initialization-free or delayed initialization scheme, which we implemented later, or even a GPS drift detection would have improved robustness. Furthermore, relative navigation enables task fulfillment with unreliable pose information, and low-level obstacle avoidance is mandatory in such a scenario. Furthermore, improved visualization of the UAV state and perception could have helped to detect problems early on.

The competition also placed an enormous strain on the involved personnel. For safe operation, one safety pilot was required per UAV, plus at least one team member supervising the internal state of the autonomous systems and being able to reconfigure the system during resets. While reducing the number of required human operators per robot is an admirable research goal, this is not possible in the near future due to safety regulations. We thus feel that future competitions should keep the required human personnel in mind so that small teams can continue to participate.

UGV motor controller failures led to crashes in the competition and highlighted a lack of proper subsystem-level testing. The controllers were tested only with manual driving commands, where integrator wind-up is less likely to occur. The lack of testing with autonomous control led to the overloading of the responsible developers. In conclusion, our team size was too small for the numerous different tasks.

What proved to be useful in the context of competitions is to prepare low-effort backup strategies in advance, like arm motions for manual mode. Also, reducing high-level control to a bare minimum, as we did for the UGV with two fixed fire locations instead of a universal search-and-extinguish, makes the systems easier to test and reduces potential errors.

VII. CONCLUSION

This paper presented an autonomous UAV-UGV team for fire extinguishing in Challenge 3 of the MBZIRC 2020. Our UGV Bob successfully delivered the maximum amount of water into the fire in the Grand Challenge. Even though our

²https://www.ais.uni-bonn.de/videos/icuas_2021_mbzirc/




Fig. 13: Navigation map and UGV trajectories on the second trial day (left) with active Fire 2 and the Grand Challenge run (right) with active Fire 1. The light green straight lines show the planned trajectory, where each branching dark green segment is a waypoint with the direction indicating the robot's heading. The other lines are the actual driven trajectory according to localization, colored differently after each reset. Two red rectangles predefine the suspected fire locations to scan for heat signatures.

UAV did not work during the competition, after analyzing and fixing the issues, our experiments show that the system functions reliably.

The shared LiDAR localization works accurately on both platforms with different LiDAR sensors. This was especially important to maneuver the oversized UGV chassis, initially developed for Challenge 2, through the door and within the restricted space inside the building. The UGV high-level control was tailored to the scenario with two fixed fire locations but can be easily extended by driving an exploratory path while scanning the environment for heat sources with continuous arm motions. The methodology proved to be sound, and failures during the competition originated from insufficient testing and last-minute changes.

We detected the fire targets from complementary modalities and fused geometric hole features from LiDAR with heat detections from thermal cameras. Our filter allowed us to approach the fire, while both robots utilize relative navigation to reliably and repeatedly spray water into the fire.

Aside from analyzed and corrected issues during the competition, we showed that both robots can perform their designated tasks autonomously and that the perception, motion planning, and fire extinguishing works reliably. The knowledge gained and components developed will also prove useful in future competitions.

REFERENCES

- [1] F. Nex and F. Remondino, "UAV for 3D mapping applications: A review," *Applied Geomatics*, vol. 6, no. 1, pp. 1–15, 2014.
- [2] L. Peskoe-Yang, "Paris firefighters used this remote-controlled robot to extinguish the Notre Dame blaze," <https://spectrum.ieee.org/automaton/robotics/industrial-robots/colossus-the-firefighting-robot-that-helped-save-notre-dame>, 2019, accessed: 2020-07-15.
- [3] G. M. Kruijff, F. Pirri, M. Gianni, P. Papadakis, M. Pizzoli, A. Sinha, V. Tretyakov, T. Linder, E. Pianese, S. Corrao, F. Priori, S. Febrini, and S. Angeletti, "Rescue robots at earthquake-hit Mirandola, Italy: [4] MBZIRC, "MBZIRC challenge description," <https://www.mbzirc.com/challenge/2020>, 2019, accessed: 2020-03-16.

- A field report," in *IEEE International Symposium on Safety, Security and Rescue Robotics (SSRR)*, 2012.
- [5] R. Bailon-Ruiz and S. Lacroix, "Wildfire remote sensing with UAVs: A review from the autonomy point of view," in *International Conference on Unmanned Aircraft Systems (ICUAS)*. IEEE, 2020.
- [6] K. A. Ghamry, M. A. Kamel, and Y. Zhang, "Cooperative forest monitoring and fire detection using a team of UAVs-UGVs," in *International Conference on Unmanned Aircraft Systems (ICUAS)*. IEEE, 2016, pp. 1206–1211.
- [7] J. Delmerico, S. Mintchev, A. Giusti, B. Gromov, K. Melo, T. Horvat, C. Cadena, M. Hutter, A. Ijspeert, D. Floreano, *et al.*, "The current state and future outlook of rescue robotics," *Journal of Field Robotics (JFR)*, vol. 36, no. 7, pp. 1171–1191, 2019.
- [8] P. Liu, H. Yu, S. Cang, and L. Vladareanu, "Robot-assisted smart firefighting and interdisciplinary perspectives," in *International Conference on Automation and Computing (ICAC)*. IEEE, 2016, pp. 395–401.
- [9] B. Aydin, E. Selvi, J. Tao, and M. J. Starek, "Use of fire-extinguishing balls for a conceptual system of drone-assisted wildfire fighting," *Drones*, vol. 3, no. 1, p. 17, 2019.
- [10] H. Ando, Y. Ambe, A. Ishii, M. Konyo, K. Tadokuma, S. Maruyama, and S. Tadokoro, "Aerial hose type robot by water jet for fire fighting," *IEEE Robotics and Automation Letters (RA-L)*, vol. 3, no. 2, pp. 1128–1135, 2018.
- [11] Y. K. Cho, Y. Ham, and M. Golpavar-Fard, "3D as-is building energy modeling and diagnostics: A review of the state-of-the-art," *Advanced Engineering Informatics*, vol. 29, no. 2, pp. 184 – 195, 2015.
- [12] C. Schönauer, E. Vonach, G. Gerstweiler, and H. Kaufmann, "3D building reconstruction and thermal mapping in fire brigade operations," in *4th Augmented Human International Conference*. ACM, 2013, pp. 202–205.
- [13] D. Borrmann, J. Elseberg, and A. Nüchter, "Thermal 3D mapping of building façades," in *Int. Conf. on Intelligent Autonomous Systems (IAS)*. Springer, 2013, pp. 173–182.
- [14] G. Demisse, D. Borrmann, and A. Nüchter, "Interpreting thermal 3D models of indoor environments for energy efficiency," *Journal of Intelligent & Robotic Systems (JINT)*, vol. 77, no. 1, pp. 55–72, 2015.
- [15] P. Fritsche, B. Zeise, P. Hemme, and B. Wagner, "Fusion of radar, LiDAR and thermal information for hazard detection in low visibility environments," in *IEEE International Symposium on Safety, Security and Rescue Robotics (SSRR)*, 2017, pp. 96–101.
- [16] R. A. Rosu, J. Quenzel, and S. Behnke, "Reconstruction of textured meshes for fire and heat source detection," in *IEEE International Symposium on Safety, Security and Rescue Robotics (SSRR)*. IEEE, 2019, pp. 235–242.
- [17] R. Suarez Fernandez, A. Rodríguez Ramos, A. Alvarez, J. Rodríguez-Vázquez, H. Bavle, L. Lu, M. Fernandez, A. Rodelgo, A. Cobano, D. Alejo, D. Acedo, R. Rey, S. Martinez-Rozas, M. Molina, L. Merino, F. Caballero, and P. Campoy, "The SkyEye team participation in the 2020 Mohamed Bin Zayed International Robotics Challenge," in *Mohamed Bin Zayed International Robotics Competition (MBZIRC) Symposium*, 2020.
- [18] M. Beul, M. Nieuwenhuisen, J. Quenzel, R. A. Rosu, J. Horn, D. Pavlichenko, S. Houben, and S. Behnke, "Team NimbRo at MBZIRC 2017: Fast landing on a moving target and treasure hunting with a team of micro aerial vehicles," *Journal of Field Robotics (JFR)*, vol. 36, no. 1, pp. 204–229, 2019.
- [19] D. Droeschel and S. Behnke, "Efficient continuous-time SLAM for 3D lidar-based online mapping," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2018.
- [20] M. Beul and S. Behnke, "Fast full state trajectory generation for multirotors," in *International Conference on Unmanned Aircraft Systems (ICUAS)*. IEEE, 2017.
- [21] C. Lenz, M. Schwarz, A. Rochow, J. Razlaw, S. Arul Periyasamy, M. Schreiber, and S. Behnke, "Autonomous wall building with a UGV-UAV team at MBZIRC 2020," *International Symposium on Safety, Security, and Rescue Robotics (SSRR)*, 2020.
- [22] M. Schwarz, D. Droeschel, C. Lenz, A. S. Periyasamy, E. Y. Puang, J. Razlaw, D. Rodriguez, S. Schüller, M. Schreiber, and S. Behnke, "Team NimbRo at MBZIRC 2017: Autonomous valve stem turning using a wrench," *Journal of Field Robotics (JFR)*, vol. 36, no. 1, 2019.