© 2021 IEEE Robotics and Automation Letters

Personal use of this material is permitted. Permission from IEEE Robotics and Automation Letters must be obtained for all other uses, in any current or future media, including reprinting or republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.

DOI: 10.1109/LRA.2021.3098304

Large-Scale Exploration of Cave Environments by Unmanned Aerial Vehicles

Pavel Petráček $_{a}^{(b)}$, Vít Krátký $_{a}^{(b)}$, Matěj Petrlík $_{a}^{(b)}$, Tomáš Báča $_{a}^{(b)}$, Radim Kratochvíl $_{b}^{(b)}$, and Martin Saska $_{a}^{(b)}$

Abstract—This paper presents a self-contained system for the robust utilization of aerial robots in the autonomous exploration of cave environments to help human explorers, first responders, and speleologists. The proposed system is generally applicable to an arbitrary exploration task within an unknown and unstructured subterranean environment and interconnects crucial robotic subsystems to provide full autonomy of the robots. Such subsystems primarily include mapping, path and trajectory planning, localization, control, and decision making. Due to the diversity, complexity, and structural uncertainty of natural cave environments, the proposed system allows for the possible use of any arbitrary exploration strategy for a single robot, as well as for a cooperating team. A multi-robot cooperation strategy that maximizes the limited flight time of each aerial robot is proposed for exploration and search & rescue scenarios where the homing of all deployed robots back to an initial location is not required. The entire system is validated in a comprehensive experimental analysis comprising of hours of flight time in a real-world cave environment, as well as by hundreds of hours within a stateof-the-art virtual testbed that was developed for the DARPA Subterranean Challenge robotic competition. Among others, experimental results include multiple real-world exploration flights traveling over 470 m on a single battery in a demanding unknown cave environment.

Index Terms—Aerial Systems: Applications; Field Robots; Aerial Systems: Perception and Autonomy; Multi-Robot Systems; Mapping

MULTIMEDIA MATERIALS

The paper is supported by the multimedia materials available at mrs.felk.cvut.cz/papers/ral-2021-caves. The implementation is also publicly available at github.com/ctu-mrs.

I. INTRODUCTION

H UMAN exploration of complex cave systems has occurred for thousands of years. However, there are still entire cave systems and individual subterranean voids, shafts, and cavities that are yet uncovered. This is primarily due to the dangerous nature of subterranean exploration in environments like natural caves, although man-made cellars,

The work was supported by the Czech Science Foundation (GAČR) under research project no. 20-29531S, by CTU grant no. SGS20/174/OHK3/3T/13, by the Defense Advanced Research Projects Agency (DARPA), and by OP VVV funded project CZ.02.1.01/0.0/0.0/16 019/0000765 "Research Center for Informatics".

a Department of Cybernetics, Faculty of Electrical Engineering, Czech Technical University in Prague, 166 36 Prague 6, Czech Republic {pavel.petracek|vit.kratky|matej.petrlik|tomas.baca| martin.saska}@fel.cvut.cz,

b Brno University of Technology, Faculty of Civil Engineering, Institute of Geodesy, Czech Republic r kratochvil@fce.vutbr.cz.

Digital Object Identifier (DOI): see top of this page.



Fig. 1: Robotic exploration of the Bull Rock Cave (central Moravian Karst, Czech Republic) by a fully autonomous aerial vehicle.

drainages, and mines pose similar risks. These environments contain sediments such as debris, rocks, sand, clay, ice, decomposed organic matter, human waste, and even various forms of speleothems in limestone caves. Considering the absolute darkness, lack of GNSS signals, flowing and dripping water, humid air, and the possible presence of poisonous gases, wind gusts, hanging ropes, and wildlife, there is excessive risk to the lives of human explorers in the exploration of new environments, as well as in search & rescue missions. Given the current state-of-the-art technology in robotics, many dangerous areas of subterranean systems are safely reachable using mobile robots, with the greatest focus being on vertical exploration using aerial vehicles. In contrast to human exploration, the use of such technology presents several advantages in the form of accessibility, safety, speed, instantaneous environment visualization, and precise quantification. On the other hand, challenges to the operation of mobile robots in such an environment lies in the uncertainty, lack of light, high humidity, and diversity of space in the form of narrow and/or low passages, canyons, large domes, high chimneys, and deep abysses.

The challenges to deployment of aerial vehicles in subterranean environments with respect to robot control, communication, sensor fusion, and positioning are described thoroughly in [1]. These specific challenges continue to be relevant even after substantial progress in the field of mobile robotics. However, in contrast to [1], our motivation includes minimizing the need for communications required for operator control and instead focuses on the full autonomy of robots and autonomous cooperation among members of a robotic team. The restriction of communication in subterranean environments introduces challenges to the maximization of system robustness and the use of efficient decision making in the form of adaptable exploration strategies in harsh unknown environments.

Manuscript received March 1, 2021; Revised May 6, 2021; Accepted July 8, 2021. This paper was recommended for publication by Editor Pauline Pounds upon evaluation of the Associate Editor and Reviewers' comments.

A. Related work

In a non-robotic context, wild caves are explored by modernly termed *cavers*. However, the human surveying and mapping of caves is known to have existed for thousands of years for purposes ranging from dwelling to speleology. The significance of cave exploration and cave mapping to scientific research is a thoroughly studied inquiry in literature, e.g., in [2].

In the work presented here, we focus mainly on the robotic point of view within the scope of the application domain. One of the first cave-mapping approaches using robotic solutions was proposed in [3], where the authors employed hand-held laser scanners, which are limited in speed, accuracy, and safety. In the context of mobile robotics, topics like the automatic control of an unstable dynamic system such as an aerial multirotor vehicle [4], the fusion of inertial, visual, and laser information for localization and mapping [5], and path planning in dynamic environments [6] have been addressed in order to achieve faster and safer methodology than mapping done with hand-held devices, as proposed in [3].

Within the scope of subterranean environments, the DARPA Subterranean Challenge competition has pushed the state of the art of autonomous exploration in human-made mines [7]-[10]. Although these systems have provided interesting solutions with great potential, the authors of [7]–[10] rely on the predictable structure of underground mines, such as using the protraction of human-made tunnels to mark the furthest depth data as frontiers or predefining turns at junctions in [9]. Since the complexity and diversity of natural caves is extensive, more robust solutions with a minimum number of environmental assumptions are required. This was tackled in [11] where the authors introduced a possible way for applying autonomous drones as a technology to assist speleologists and archaeologists. Although an interesting read, the proposed methods only constitute a preliminary discussion that presents neither novel technology nor applied results. A similar discussion focusing on the state of robotic problems within the application of subterranean exploration with UAVs is presented in [12]. In contrast to [11], the authors of [12] present a set of preliminary experiments in laboratory conditions and two dimensional space. Unfortunately, the assumption of a planar world is highly restrictive within the scope of real-world deployment due to the complex character of natural subterranean environments.

The precise localization of mobile robots is crucial to autonomous navigation in such complex environments. Among existing state-of-the-art literature, the LOCUS algorithm [13] achieves the lowest localization error at the cost of high computational demands. Unlike with ground robots, this method might be unsuitable for aerial robots as the computational resources on lightweight UAVs are scarce due to their limited payload. In [14], the authors demonstrated that localization performance can be further improved by dropping range beacons. This is a viable strategy for heterogeneous robotic teams, but unfeasible for teams of only lightweight UAVs.

The use of robotic teams for cooperative exploration has been addressed mostly in planar worlds with recurrent connectivity constraints [15] or with the requirement of a centralized element [16]. A similarly defined task to our problem of team homing — respecting intermittent communication, need for decentralization, and limited operation time of aerial robots - is proposed in [17], where the robots gather and share data during the mission and return all the way back to the base before their operation times out. In contrast to [17], we propose homing coordination that lands each aerial robot at a position expanding a communication relay graph, thereby increasing the time for mere exploration in tasks where return to the starting position is not required. Related to the scope of search & rescue, the authors in [18] propose to re-position robots in a relay-chain formation to enable data transmission over longer distances once an object of interest is found. Our solution reports the position of the objects once the explorer robot connects to the relay graph during homing. The recently developed fast exploration technique in [19] maximizes explored volume over battery-limited flight time. The method is based on data only from an RGBD camera with a limited field of view (FoV). In comparison to LiDARbased methods, we have experimentally verified that RGBD cameras are sub-optimal sensors for the exploration of largescale caves due to their limited range and FoV.

B. Contributions

First, we propose a fully autonomous system enabling multimodal mapping, fast and efficient planning with sensoric fieldof-view constraints for safe movement in 3D, robust localization, and adaptable decision making. Second, a multi-robot cooperation for the efficient homing of a team of autonomous explorer robots is proposed. Third, the system has been validated through hundreds of hours of testing in a state-of-theart virtual testbed developed for the DARPA Subterranean Challenge robotic competition, as well as through hours of flight time in the real world. To the best of our knowledge, the presented large-scale experimental deployment of autonomous aerial robots in a natural cave environment goes beyond the current state of the art in autonomous robotics. Lastly, we present and share the experience obtained during this comprehensive experimental deployment that was carried out in close cooperation with speleologists.

II. EXPERIENCE GAINED

A. Speleology motivation

From the speleological point of view, aerial systems are crucial for pushing exploratory state-of-the-art methods to provide assistance in efficient scouting of difficult-to-access areas in vertical environments, as well as for the quick inspection of known areas using onboard sensors only. These systems minimize risks for humans by reducing the need to climb or to swim in cold water reservoirs, and also through the detection of poisonous gases or even radioactive waste. Furthermore, this enables the preservation and protection of natural environments against human influence, including ancient sediment forms, floor dripstone formations, paleontological and archaeological sites, and sources of potable water.

In contrast to well-established methods of subterranean documentation (i.e., theodolite and level/distance meter, compass, and clinometer), modern technology employs stationary and mobile laser scanners to produce a dense 3D model of the environment. Due to the complexity of natural environments, the use of stationary scanners is time-consuming because of the necessity of eliminating occluded spaces. Although handheld mobile scanners are more time-efficient in this context, their use is limited to areas accessible to humans. This limitation opens the door for mobile robotics which is able to tackle this challenge and to provide optimized 3D mapping. State-ofthe-art mapping in such environments reaches decimeter level precision, which is less precise than stationary scanners, yet sufficient for the majority of speleological needs. Moreover, the common issue of mapping drift accumulation in longcorridor spaces can be minimized using reference measurements by precise stationary scanners or man-measured control points to obtain accurate results.

B. System requirements

The primary prerequisite of a team of aerial explorers that can be deployed in caves involves the ability to adapt to diverse, unknown environments lacking sources of light and access to GNSS. This general description requires the abilities to

- be deployed in constrained cavities, as well as in open caverns of natural caves,
- map and visualize the environment in a fast, quantified manner in the form of dense point clouds and image streams,
- seamlessly infuse an arbitrary exploration strategy for more efficient mission operation within the scope of individual environments (policy selection is discussed in Sec. IV),
- return to the mission operator and promptly visualize the environment for human supervision, and
- maximize operation capabilities in terms of coverage when a team of robots is employed.

C. Depth estimation in high humidity

The performance of the PMD pico flexx time-of-flight (ToF) camera and the Intel Realsense D435 stereo camera have been analyzed as complementary sensors to the primary LiDAR for the purpose of improving the sensory FoV coverage. Although ToF cameras generally outperform stereo cameras in terms of distance measurement precision and density of measurement points [20], the high humidity typically present in natural caves causes dispersion of light emitted from ToF cameras by small water droplets. This effect significantly degrades the acquired measurements. As was verified empirically, ToF cameras can produce false-negative measurements of obstacles situated behind clouds of water droplets. The use of stereo cameras (e.g., Realsense) is recommended for its robustness to environmental conditions within natural caves. Nevertheless for large cave systems, such a sensor needs to be combined with 3D LiDARs in order to comply with the requirements of speleologists and first responders.

III. SYSTEM ARCHITECTURE

The system of the proposed autonomous explorer robot is divided into multiple groups of individual interconnected modules to be described in this section. All components and their relations are visualized in Fig. 2.



Fig. 2: Individual interconnected modules form the system architecture of the autonomous explorer robot. The *High-level planning* modules focus on achieving the mission objectives by generating references for the *Tracking & control* modules based on the map built by the *Perception* modules. This also provides a state estimate for closing the control feedback loop. All modules except the *Autopilot* group are handled by the main onboard computer.

A. Perception

The perception of the proposed system is based on a multichannel *LiDAR* sensor that is used for both building the spatial representation of the surrounding environment in the *Mapping* module, as well as for the motion estimation in the *LOAM* module. Obtaining the full-state estimate is realized within the *State estimation* module, where multiple sources of incomplete state measurements are fused together using a bank-offilters estimator.

The vertical navigation capabilities of the system can be greatly improved by equipping the robot with vertically-facing *RGBD* cameras that are able to fill in the blind spots in the limited vertical FoV of the *LiDAR*. Apart from navigation, these optional sensors may be used for detecting objects of interest in caves in search & rescue scenarios or for visual documentation of newly explored cave systems.

1) LiDAR: Even though our system is not tied to a specific LiDAR model, there are certain important parameters that can affect the performance and capabilities of the platform.

To reliably stabilize the UAV, the time delay of the estimated state must stay below the threshold of a certain critical value depending on the type of controller and gains. When this threshold is exceeded, the UAV begins oscillating and eventually automatically lands when the control error is too large to continue the mission safely. We have found experimentally that for most combinations of localization methods and controllers, the critical value ranges from 100 ms to 200 ms. Thus, 10 Hz is the lowest rotation frequency that can be used without employing methods of delay compensation.

The typical values of a vertical field of view (VFoV) of 3D LiDARs are in the 30° to 90° range. The higher VFoV values improve vertical mobility in constrained spaces, however with a low VFoV, it is impossible to safely navigate narrow vertical shafts as it is not known whether the space above the UAV is free and safe to fly through, or whether it contains an obstacle.

2) *RGBD*: The regions above and below the UAV that are not covered by the LiDAR can be captured using a depth camera or by spinning the LiDAR sensor around a vector that is orthogonal to the axis of scanning, as seen in [21]. However, such a solution adds additional weight to the sensor, which decreases the available flight time. A blind spot also still remains as part of the laser rays is blocked by the frame of the UAV. Alternatively, lightweight depth cameras can be mounted on opposite sides of the body frame in order to cover most blind spots of the LiDAR. Additional sensing modality is gained by combining an RGB and depth camera in a single sensor (RGBD) with a slight weight increase.

3) Localization: For localization of the UAV, we have adapted the LOAM algorithm [22]. This state-of-the-art method is very precise (0.55% translation error [23]) while attaining real-time performance. In our adapted version of the open-source implementation, the algorithm is optimized on CPU and employs parallel computing, which enables us to deploy and use the localization in the real-time position control feedback loop onboard fast-moving aerial vehicles.

4) Mapping: The LOAM-based algorithm builds a sparse internal representation of the environment consisting of edge and planar features. However, this sparse map is unsuitable for navigation purposes. Additionally, the LOAM map does not consider the probabilistic nature of the sensor, nor does it distinguish free and unknown space. Both of these factors are necessary in exploration techniques for reliable navigation and consistent frontier selection.

In the proposed system, the environment is represented by a dense probabilistic volumetric map, which consists of cubic cells with one of 3 states: free, occupied, or unknown. The map is kept in the octree structure to facilitate the Bayesian integration of new measurements and efficient access to individual cells of the probabilistic map. The high-level systems, such as grid-based path planning or inter-robot map registration, also benefit from quick access to the dense environment representation. This approach is capable of multi-modal fusion by integrating the data from all available onboard sensors and outputting point-cloud measurements. If high-level path planning is constrained by the field of view of onboard sensors (tackled in [24] and also in Sec. III-C), the multi-modality of mapping enables arbitrary movement in 3D.

5) Sensor processing: The targeted subterranean environments may have high humidity or may contain large clouds of whirling dust. The water and dust particles can then produce erroneous measurements for the LiDAR-based sensors. Assuming a partial reflection from water or dust particles and a large energy dissipation of distant reflections, these erroneous measurements can be filtered with respect to the measured intensity of returning light rays. As has been empirically verified, a simple threshold-based filtration over the intensity channel within the local neighborhood of the sensor is sufficient for filtering out false-positive measurements. The idea of the local filtration is to filter out particles gusting through the surrounding air due to the aerodynamic influence of the propellers. Although the cutoff threshold of the intensity magnitude is environment-specific, filtering out measurements below the 10th percentile of the intensity distribution per each laser scan proved to be a reliable solution, even in the dustiest real-world environments. Such processing is unavailable for camerabased systems that may require thorough, computationallyexpensive solutions to overcome these challenges.

6) State estimation: The reference controller (see Sec. III-B2) requires a position estimate of the UAV body frame in the world frame $\mathbf{r} = (x, y, z)$, the velocity of the

body frame $\dot{\mathbf{r}}$, rotation \mathbf{R} from the UAV body frame to the world frame, and angular velocity $\boldsymbol{\omega}$ in the body frame in order to close the feedback loop. The LOAM localization method provides 6-DoF pose estimate, i.e., \mathbf{r}_{LOAM} , \mathbf{R}_{LOAM} , which are fused in the *State estimation* block with interoceptive measurements from the *IMU* of the *Autopilot* to obtain the rest of the state variables.

The details about the estimation process are described in [25]. Nevertheless, it is worth highlighting the importance of the fusion of orientation \mathbf{R}_{IMU} and \mathbf{R}_{LOAM} in cave environments. While \mathbf{R}_{IMU} is very precise and without delay, the heading of the UAV (i.e., the measured direction of the bodyfixed, forward-facing axis) is unreliable due to the presence of ferromagnetic ores in the cave rocks that cause deviations in the magnetometer measurements. By correcting these errors with the heading from \mathbf{R}_{LOAM} in the estimation process, the resulting orientation \mathbf{R} is robust to changes in the erratic magnetic field in subterranean environments.

B. Tracking & control

The safe navigation of constrained environments with low obstacle clearance imposes the requirements of precise trajectory tracking with minimal control error, as any deviation from the desired state could potentially result in a collision. The *Reference controller* is responsible for minimizing the control error around the desired control reference that is provided by the *Reference tracker*. The controller outputs an attitude rate reference for the low-level *Attitude rate controller* in the *Autopilot*.

1) Reference tracker: The Reference tracker is essential in providing the Reference controller with smooth and feasible references to ensure a safe flight. The tracker based on the model predictive control (MPC) simulates an ideal virtual model of the UAV with constrained translational states up to jerk, together with heading and heading rate. The input can be either a single pair of desired 3D position \mathbf{p}_d and heading η_d , or a trajectory T_d in the form of a sequence of such pairs with a specified sampling rate. The full state of the virtual model is then sampled at 100 Hz, and \mathbf{r}_d , $\dot{\mathbf{r}}_d$, $\ddot{\mathbf{r}}_d$, η_d , η_d are passed to the Reference controller as reference \mathbf{x}_d .

2) Reference controller: The agile SE(3) geometric state feedback controller [26] minimizes the position and velocity errors. To compensate imperfect calibration and external forces acting upon the UAV, the controller is extended with the body and world disturbance terms described in [25]. The output attitude rate reference ω_d is tracked by the *Autopilot*.

C. Path planning

The planning approach used to safely navigate through apriori unknown environments must fulfill requirements of realtime responsiveness and efficient global planning in order to fully exploit the limited flight time of UAVs. For this purpose, fast iterative post-processing is applied to the output of an optimal grid-based planner in order to increase the UAVobstacle distance above a minimum threshold [27]. The gridbased planner and the iterative post-processing do not apply an optimistic assumption that the unknown space is collisionfree. Although this visibility-constrained precondition requires high sensory coverage around the robot to allow for arbitrary movement in 3D, it consequently prevents collisions of the trajectory being followed, even if replanning would fail. This methodology improves safety and robustness of the overall flight, allows for deployment in completely unknown environments without any apriori information, and permits seamless navigation in open spaces, as well as safe movement through narrow passages.

Common grid-based planning methods require preprocessing of an employed map representation, such as determining and applying the 3D distance transform for obstacle growing. This may introduce significant computational overhead by bottle-necking system performance, as the map must then be processed in every planning step. Such a computationally expensive task contradicts the requirements for responsiveness within evolving dynamic environments. To minimize the overall time required for a single planning iteration, a local KD-tree representation of the environment is used to decide the feasibility of particular cells within a voxel grid. This approach shifts the largest load from the pre-processing phase to the planning phase, which is beneficial especially to shorter plans that require searching only a small part of the environment. The low computational demands of the applied planning approach enable frequent replanning the global plan, which is also crucial for the efficient use of newly-discovered collision-free space.

To effectively exploit the limited flight time of aerial explorers, all mid-flight stops are eliminated by computing in parallel the next exploration goal during path following. The path to the next goal is efficiently appended to the rest of the current reference trajectory T_d using the prediction horizon of the MPC (see Sec. III-B). The need for precise locomotion control in complex natural caves makes uniform path-sampling unfeasible with respect to the dynamic constraints of a UAV and fast, collision-free trajectory tracking. Therefore, the reference trajectory T_d provided by the *Navigation & planning* module to the *Reference tracker* is computed based on the following process.

Given the dynamical constraints of the robot, the generated path is uniformly sampled with a sampling distance adapted to the maximum velocity magnitude v_{max} of the UAV. Based on this initial trajectory T_i , the required acceleration magnitudes a_n between consequent transition points are computed by velocity differentiation as

$$a_n(k) = \frac{||\mathbf{v}_i(k+1) - \mathbf{v}_i(k)||_2}{t_s},$$
(1)

where $\mathbf{v}_i(k)$ is the required velocity vector for transition from a transition point $t_i(k)$ to $t_i(k+1)$ on the initial trajectory T_i and t_s is a constant sampling period. The new velocity for a *k*-th segment is then given by

$$v_{k} = \begin{cases} \max\left(v_{max} \frac{a_{max}}{a_{n}}, v_{min}\right) & \text{if } a_{n}(k) > a_{max}, \\ v_{max} & \text{if } a_{n}(k) \le a_{max}, \end{cases}$$
(2)

where the minimum velocity v_{min} serves as a parameter balancing the precision and the time needed for trajectory tracking. By this step, the velocities for particular segments are set so that the maximum velocity is applied in straight segments, while lower velocities are applied in curved segments of any given path.

To further improve trajectory sampling and to achieve smoother changes in velocities, the sampling distance on particular segments is computed so that the motion along each segment has the constant acceleration

$$\overline{a}_k = \frac{|v_{k+1} - v_k|}{t_{acc,k}},\tag{3}$$

where $t_{acc,k}$ is the time available for acceleration on the k-th segment. The time $t_{acc,k}$ is obtained from the length l_k of the segment k and the required change of the velocity. The number of transition points N_k on the k-th segment of the initial trajectory T_i is given as

$$N_{k} = \begin{cases} \begin{bmatrix} \frac{l_{k}}{v_{k}t_{s}} \end{bmatrix} & \text{if} \quad \overline{a}_{k} = 0, \\ \begin{bmatrix} \frac{t_{acc}}{t_{s}} \end{bmatrix} & \text{if} \quad \overline{a}_{k} > 0, \end{cases}$$
(4)

where the desired constant acceleration is adapted to meet the velocity v_{k+1} at the end of each segment as

$$a_k = \frac{\overline{a}_k}{N_k t_s}.$$
(5)

The sequence of sampling distances for the k-th segment of T_i is then given by

$$U_{k,i} = v_k t_s + i a_k t_s^2, \ i \in \{1, \cdots, N_k\}.$$
 (6)

The trajectory sampled with sampling distances defined by (6) is passed to the *Reference tracker* [25] as a reference trajectory T_d in order to generate a feasible reference \mathbf{x}_d for the *Reference controller*. Despite its simplicity, the described sampling method achieves better results within the scope of the proposed application than the optimization-based trajectory generation methods proposed in [28], [29]. In contrast to the proposed method, the problem in [28], [29] is defined in such a way that the exact positions of all the path waypoints must be visited, generating significantly slower trajectories.

IV. EXPLORATION POLICY

Cave environments are naturally diverse and require various different mission strategies suitable for specific environments. Deriving the optimal policy is thereby dependent on various factors, such as the expected mission output, mission-specific constraints, the complexity and the specifics of the environment, and the number of available robots. For this reason, our system is designed so that any arbitrary policy can be utilized within the scope of an autonomous mission.

Nevertheless, two exploratory mission types are of the most use in practice: deep cave exploration and full-coverage exploration. These missions are used for scouting previously uncovered areas in order to obtain a general overview of the environment, monitor environmental changes such as gas leaks, detect natural water reservoirs, discover new possible passages, or assess the structural state of cavern walls and other objects of interest. The former approach maximizes the explored volume of space in the entire environment, while the latter minimizes the blind spots missed by onboard cameras with a constrained FoV.

The capabilities of a robotic mission are furthered with the use of multiple cooperating robots. To show an example of such improvement using a team of agents as opposed to a single agent, a homing strategy that maximizes the flight time of aerial robots during a multi-robotic mission is proposed in the following subsection. During the proposed coordination, continuous exploration is not assumed and distance-constrained ad-hoc communication is used. The robots are homogeneous and generate their behaviors in a decentralized manner based on their current state and the available information from other robots (only positions in a shared frame are required).

A. Multi-robot homing strategy

A cooperative operation maximizing the flight time of a multi-robot team is proposed for applications where homing all the deployed robots to an initial location is not required. This strategy is suitable for tasks where the possible gained information is superior to the cost of the robots, such as in search & rescue scenarios. This method assumes there is access to a low-bandwidth communication link among any two robots within an omnidirectional communication radius.

To maximize the flight time, the robots utilize local communication to plan the homing path such that a group of robots is able to build up a communication tree with the base station as the root communication node. This allows the robots to optimize their flight time by navigating back to a location in the proximity of another communication node (a landed robot, base station, or self-sustaining communication node deployed by other robots) when the battery capacity becomes drained. This entire homing strategy is showcased in an example scenario for two independent robots in Fig. 3.

In the proposed strategy, each robot constructs a navigation homing tree using nodes created from the set of past poses of the robot. This online-built tree has edges valued by the required flight time between two nodes and is used to estimate required homing time to the proximity of a communication node. The pose nodes are connected such that each path leafto-communication is the shortest (see Fig. 3a). A homing path is constructed recursively as a sequence of tree nodes from the current robot position (a leaf) to the nearest communication node, with the landing position being within communication range of the nearest communication node (see Fig. 3b). The tree is shared among the robots deployed in the same mission. The knowledge from the previous explorers is integrated to prolong their flight time (see Fig. 3c), thus causally maximizing the time capacity for the exploration task. When a communication node (e.g., a robot landing pose) is integrated into the homing tree, it is linked exclusively to another communication node to join the retranslation chain (see Fig. 3d). Consequently, the parents of neighboring pose nodes are updated so that each pose node has a parent with the minimal accumulated cost to any communication node (see Fig. 3b and Fig. 3d). The process of inserting pose nodes as well as communication nodes into the homing tree is described in Alg. 1.

V. EXPERIMENTAL ANALYSIS

The entire proposed system has been validated through hours of flight time in the real world, as well as in hundreds of hours in various virtual subterranean environments. The results of these experimental analyses are presented hereafter.





Fig. 3: An example scenario of the homing strategy for two robots (red and blue) that maximizes flight time by landing at feasible positions while building a communication chain to a base station.

Algorithm 1: Insertion of a node into the onboard-built homing tree. Function $cost(n_a, n_b)$ returns an estimate of flight time among nodes n_a and n_b , function accumulatedCost (n_a) returns the required flight time from node n_a to the nearest communication node, and function freeRay (n_a, n_b) returns true if a linear path between nodes n_a and n_b is collision-free in 3D.

1: procedure INSERTNODETOHOMINGTREE

	Input:	
:	Ν	▷ Node to be inserted
:	\mathcal{C},\mathcal{P}	Sets of communication and pose nodes
i:	d_e	Minimum edge length
:	if $N.type == COMMUI$	NICATION then
':	$N.parent \leftarrow \arg n$	$\min_{c \in \mathcal{C}} \operatorname{cost}(\mathbf{N}, c)$
3:	for $p \in \mathcal{P}$ do	Update parents of neighboring pose nodes
):	if $cost(N, p) < a$	ccumulated $Cost(p)$ then
0:	$p.parent \leftarrow$	- N
1:	$\mathcal{C} \leftarrow \mathcal{C} \cup \mathbb{N}$	
2:	else	
3:	$\mathcal{V} \leftarrow \mathcal{C} \cup \mathcal{P}$	
4:	if $\min_{v \in \mathcal{V}} (\mathbf{N} - \mathbf{v})$	$v _2) \ge d_e$ then
5:	$\overline{\mathcal{V}} = \{v \mid \text{freeRa}\}$	$y(N, v), \forall v \in \mathcal{V}\}$
6:	if $\overline{\mathcal{V}} \neq \emptyset$ then	
7:	$N.parent \leftarrow$	$- \arg \min_{v \in \overline{\mathcal{V}}} \left[\operatorname{cost}(N, v) + \operatorname{accumulatedCost}(v) \right]$
8:	$\mathcal{P} \leftarrow \mathcal{P} \cup \mathbb{N}$	-

A. Real-world environment

To analyze the properties of the system, a fully autonomous aerial robot (see Fig. 4) was deployed for several hours of flight time in the Bull Rock Cave located in the central Moravian Karst of the Czech Republic (see Fig. 1 and the attached multimedia materials).

During multiple autonomous exploratory missions, a single explorer (see the hardware components of the robot in Fig. 4) was deployed to validate the proposed system in various exploratory scenarios. The flight trajectories from all missions are visualized in Fig. 5a and the mission statistics and performance metrics of the mapping module are summarized in Table I. A greedy frontier-navigation policy was employed such that the frontier closest to the lateral direction of flight (A, B), the highest frontier (C), and frontier with the largest ratio of unknown to free cells in a bounded area (D) was selected as the next goal. With respect to these experiments in a harsh subterranean environment, we have

• validated the performance of the system by flying in large cave domes, as well as in narrow corridors just 70 cm wider than the dimensions of the robot,

	Eliabt	Turicatowy	Evalored	Mapping accuracy		
Trial	time (s)	length (m)	volume (m^3)	μ (m)	σ (m)	
А	731	476	7463	0.57	0.59	
В	935	473	11 403	0.53	0.56	
С	359	71	551	0.23	0.26	
D	749	602	3756	0.33	0.38	
Е	386	233	3055	0.39	0.39	
F	633	256	2579	0.21	0.22	
G	638	261	3650	0.27	0.33	
Н	297	142	1682	0.25	0.41	
Ι	129	121	3326	0.19	0.22	
J	425	233	4388	0.25	0.29	

TABLE I: Quantitative evaluation on multiple autonomous exploratory missions within the Bull Rock Cave system. The flight trajectories and qualitative analysis of the mapping accuracy are shown in Fig. 5.

- validated the real-time performance and robustness of the system in multiple autonomous horizontally-deep flights longer than 470 m using just a single battery and reaching a maximal velocity up to 2 m s⁻¹,
- validated the ability to autonomously explore natural domes in terms of vertical depth,
- verified the ability to perform a full mission and return to an initial location with the obtained information,
- quantified the accuracy of the onboard-built maps with respect to a ground truth map of the environment, and
- obtained feedback from speleologists in order to design the system following their requirements.

The dense onboard-built maps (20 cm resolution) from all the experiments were merged (manual global registration with local ICP refinement) during post-processing to obtain the map of the environment \mathcal{M} . The reference ground truth map \mathcal{M}_{gt} was built by registering over 100 largely overlapping scans taken by a Leica BLK360 terrestrial 3D scanner. The mapping accuracy over all the experiments reached mean $\mu = 0.37 \,\mathrm{m}$ and standard deviation $\sigma = 0.46 \,\mathrm{m}$ using the point-to-point Euclidean error metric between each point in \mathcal{M} and the corresponding closest point in \mathcal{M}_{gt} . The distribution of the mapping errors throughout all flights is visualized in Fig. 5b. As specified by the end-users, the decimeter-level mapping precision achieved over the course of these exceptionally fast and extensive flights is sufficient for the majority of speleological needs.



Fig. 4: General hardware components of an autonomous explorer robot. All data are processed and reasoned over with an onboard processing unit. The main source of data comes from the top-mounted LiDAR.

B. Virtual environment

To validate the proposed methodology for multi-robot coordination using a local low-bandwidth communication network, a team of aerial robots was deployed for hundreds of hours of flight in a virtual environment using a virtual testbed developed for the DARPA Subterranean Challenge competition.



(a) Overview of the cave environment with the trajectories of all exploration missions (see Table I) performed within Bull Rock Cave. The figure shows deep cave missions (A, B), vertical flight (C), and the thorough exploration of a bounded area (D).



(b) Visual analysis on the mapping accuracy – the distribution of mapping errors during all autonomous exploration tasks as summarized in Table I. The color bar legend represents the mapping error in meters using the point-to-point euclidean error metric.

Fig. 5: Full-coverage exploration of the Bull Rock Cave system (located in the central Moravian Karst, Czech Republic) with autonomous aerial explorers. Full resolution figure is available within the attached multimedia.

This state-of-the-art testbed consists of several large-scale cave environments containing dynamic obstacles and models of real-world interference, such as sensor discrepancies, communication schemes, and battery longevity.

In contrast to real-world experiments, the virtual environment is larger and allows for the seamless verification of multirobotic cooperation. To demonstrate the performance of the proposed homing strategy, a selected example scenario of such an operation is presented in Fig. 6. This experiment highlights the positive influence of the homing strategy in a search & rescue scenario where the three explorers were able to exploit the increased flight time. With a 50 m communication range and 1.2 m s⁻¹ average velocity for each robot, the homing cooperation increased the available flight time for exploration by 40 s and 80 s, respectively. Moreover, the experiment shows the influence of multi-sensor mapping, which allowed the black robot to single-handedly explore the upper floor of the virtual environment. The final exploratory trajectories of the cooperating robots during the presented mission reached lengths of 715 m, 1349 m, and 1405 m.

The influence of the homing strategy on the time available for mere exploration is also quantitatively analyzed in Table II. The results were averaged over six separate deployments, each with five cooperating robots. Identical mission parameters were set to all the robots for the baseline [17], as well as for the proposed method. The data show an increasing trend in the available mission time for belated explorers for which the effective exploration phase is consequently prolonged during their entire operation time.

VI. CONCLUSION

This letter presents a comprehensive study on the use of autonomous aerial explorers as an assisting technology for the exploration of natural cave environments. This study also



Fig. 6: Three autonomous explorers deployed in a virtual cave world within the DARPA simulation testbed. The robots finished their missions by building a communication tree with maximal edge length of $d_c = 50$ m.

Robot	1 st	2nd	3 rd	4 th	5 th
Exploration time before homing (secs)	316	330	360	380	389
Exploration time increase (%)	-1.5	2.8	12.2	18.4	21.0

TABLE II: Influence of the homing strategy on the flight time available for mere exploration. Comparison with a baseline time of 321s (averaged over 10 flights) where a robot returned to base before its operation timed out.

shares the experience acquired during the technology's development in close cooperation with a team of speleologists, cavers, and first responders.

The proposed self-sustaining system interconnects solutions for all crucial robotic tasks in order to enable full autonomy in complex unknown subterranean environments without access to GNSS. Among others, this includes laser-data processing which copes with high humidity and dustiness within subterranean environments and robust path-planning for unknown dynamic environments to allow for flights in constrained cavities, as well as in open caverns of natural caves. Moreover, a multi-robot cooperation is proposed for the efficient homing of a team of robots for applications where the possible information gain is superior to the costs of the robots, such as search & rescue scenarios in cave systems. The performance of the entire applicable system was validated in one of the most large-scale experimental analyses ever conducted, consisting of hours of flight time in Bull Rock Cave (Czech Republic, Moravian Karst) and in hundreds of hours in the state-ofthe-art virtual testbed developed for the DARPA Subterranean Challenge. This presented analysis of the entire system proves that it is a robust solution capable of reliable planning with sensoric field-of-view constraints and accurate mapping. The accuracy of localization and mapping was evaluated with respect to a ground-truth map of the cave environment and reached mean precision below 40 cm in real-world conditions. This performance has satisfied the requirements of speleologists and first responders.

REFERENCES

- R. La Scalea *et al.*, "Opportunities for autonomous UAV in harsh environments," in *IEEE ISWCS*, 2019, pp. 227–232.
- [2] P. Kambesis, "The importance of cave exploration to scientific research," *Journal of Cave and Karst Studies*, vol. 69, pp. 46–58, 04 2007.
- [3] R. Zlot and M. Bosse, "Three-Dimensional Mobile Mapping of Caves," Journal of Cave and Karst Studies, vol. 76, pp. 191–206, 2014.
- [4] T. P. Nascimento and M. Saska, "Position and attitude control of multirotor aerial vehicles: A survey," *Annual Reviews in Control*, vol. 48, pp. 129–146, 2019.
- [5] C. Debeunne and D. Vivet, "A Review of Visual-LiDAR Fusion based Simultaneous Localization and Mapping," *Sensors*, vol. 20, no. 7, p. 2068, 2020.
- [6] Y. Zhao et al., "Survey on computational-intelligence-based UAV path planning," *Knowledge-Based Systems*, vol. 158, pp. 54–64, 2018.
- [7] K. Ebadi *et al.*, "LAMP: Large-Scale Autonomous Mapping and Positioning for Exploration of Perceptually-Degraded Subterranean Environments," in *IEEE ICRA*, 2020, pp. 80–86.
- [8] T. Dang et al., "Explore Locally, Plan Globally: A Path Planning Framework for Autonomous Robotic Exploration in Subterranean Environments," in *IEEE ICAR*, 2019, pp. 9–16.
- [9] I. D. Miller *et al.*, "Mine Tunnel Exploration Using Multiple Quadrupedal Robots," *IEEE RA-L*, vol. 5, no. 2, pp. 2840–2847, 2020.
- [10] M. Petrlík *et al.*, "A Robust UAV System for Operations in a Constrained Environment," *IEEE RA-L*, vol. 5, no. 2, pp. 2169–2176, 2020.
- [11] G. Zhang et al., "SmartCaveDrone: 3D cave mapping using UAVs as robotic co-archaeologists," in *IEEE ICUAS*, 2017, pp. 1052–1057.
- [12] J. G. R. III et al., "Distributed subterranean exploration and mapping with teams of UAVs," in Ground/Air Multisensor Interoperability, Integration, and Networking for Persistent ISR VIII, vol. 10190, 2017, pp. 285–301.
- [13] M. Palieri *et al.*, "LOCUS: A Multi-Sensor Lidar-Centric Solution for High-Precision Odometry and 3D Mapping in Real-Time," *IEEE RA-L*, vol. 6, no. 2, pp. 421–428, 2020.
- [14] N. Funabiki et al., "Range-Aided Pose-Graph-Based SLAM: Applications of Deployable Ranging Beacons for Unknown Environment Exploration," *IEEE RA-L*, vol. 6, no. 1, pp. 48–55, 2020.
- [15] Y. Pei *et al.*, "Coordinated Multi-Robot Real-Time Exploration with Connectivity and Bandwidth Awareness," in *IEEE ICRA*, 2010, pp. 5460–5465.
- [16] W. Burgard et al., "Coordinated Multi-Robot Exploration," IEEE Transactions on Robotics, vol. 21, no. 3, pp. 376–386, 2005.
- [17] K. Cesare *et al.*, "Multi-UAV Exploration with Limited Communication and Battery," in *IEEE ICRA*, 2015, pp. 2230–2235.
- [18] J. Scherer *et al.*, "An Autonomous Multi-UAV System for Search and Rescue," in *DroNet*, 2015, p. 33–38.
- [19] B. Zhou *et al.*, "FUEL: Fast UAV Exploration Using Incremental Frontier Structure and Hierarchical Planning," *IEEE RA-L*, vol. 6, no. 2, pp. 779–786, 2021.
- [20] F. Pece et al., "Three depth-camera technologies compared," in First BEAMING Workshop, Barcelona, vol. 2011. Citeseer, 2011, p. 9.
- [21] J. Williams *et al.*, "Online 3D Frontier-Based UGV and UAV Exploration Using Direct Point Cloud Visibility," in *IEEE MFI*, 2020, pp. 263–270.
- [22] J. Zhang and S. Singh, "Low-drift and real-time lidar odometry and mapping," vol. 41, no. 2, 2017, pp. 401–416.
- [23] A. Geiger *et al.*, "Are we ready for Autonomous Driving? The KITTI Vision Benchmark Suite," in *IEEE CVPR*, 05 2012, pp. 3354–3361.
- [24] M. Beul *et al.*, "Fast Autonomous Flight in Warehouses for Inventory Applications," *IEEE RA-L*, vol. 3, no. 4, pp. 3121–3128, 2018.
- [25] T. Baca *et al.*, "The MRS UAV System: Pushing the Frontiers of Reproducible Research, Real-world Deployment, and Education with Autonomous Unmanned Aerial Vehicles," *J Intell Robot Syst*, vol. 102, no. 26, pp. 1–28, 2021.
- [26] T. Lee *et al.*, "Geometric tracking control of a quadrotor UAV on SE(3)," in *IEEE CDC*, 2010, pp. 5420–5425.
- [27] V. Krátký et al., "An Autonomous Unmanned Aerial Vehicle System for Fast Exploration of Large Complex Indoor Environments," JFR, vol. 5, no. 2, pp. 1–24, 2021.
- [28] C. Richter *et al.*, "Polynomial trajectory planning for aggressive quadrotor flight in dense indoor environments," in *Robotics Research*. Springer, 2016, pp. 649–666.
- [29] M. Burri et al., "Real-Time Visual-Inertial Mapping, Re-localization and Planning Onboard MAVs in Unknown Environments," in *IEEE IROS*, Sept 2015, pp. 1872–1878.