

Special Issue: DARPA Subterranean Challenge, Advancement and Lessons Learned from the Finals (DARPA SubT Final)

Regular Article

# Flexible Supervised Autonomy for Exploration in Subterranean Environments

Harel Biggie<sup>\*1</sup>, Eugene R. Rush<sup>\*2</sup>, Danny G. Riley<sup>1</sup>, Shakeeb Ahmad<sup>3</sup>,  
Michael T. Ohradzansky<sup>3</sup>, Kyle Harlow<sup>1</sup>, Michael J. Miles<sup>2</sup>, Daniel Torres<sup>1</sup>,  
Steve McGuire<sup>4</sup>, Eric W. Frew<sup>3</sup>, Christoffer Heckman<sup>1</sup> and J. Sean Humbert<sup>2</sup>

<sup>1</sup>Computer Science, University of Colorado Boulder

<sup>2</sup>Mechanical Engineering, University of Colorado Boulder

<sup>3</sup>Aerospace Engineering Sciences, University of Colorado Boulder

<sup>4</sup>Electrical & Computer Engineering, University of California Santa Cruz

**Abstract:** While the capabilities of autonomous systems have been steadily improving in recent years, these systems still struggle to rapidly explore previously unknown environments without the aid of GPS-assisted navigation. The DARPA Subterranean (SubT) Challenge aimed to fast track the development of autonomous exploration systems by evaluating their performance in real-world underground search-and-rescue scenarios. Subterranean environments present a plethora of challenges for robotic systems, such as limited communications, complex topology, visually-degraded sensing, and harsh terrain. The presented solution enables long-term autonomy with minimal human supervision by combining a powerful and independent single-agent autonomy stack, with higher level mission management operating over a flexible mesh network. The autonomy suite deployed on quadruped and wheeled robots was fully independent, freeing the human supervision to loosely supervise the mission, and make high-impact strategic decisions. We also discuss lessons learned from fielding our system at the SubT Final Event, relating to vehicle versatility, system adaptability, and reconfigurable communications.

**Keywords:** subterranean robotics, navigation, emergency response, exploration, GPS-denied operation

## 1. Introduction

Despite a myriad of developments in sensing, planning, control and state estimation over the last few decades, deploying robots in harsh subterranean environments for the purpose of rapid situational awareness presents a number of new challenges to robot autonomy. Traditionally, robots rely on a number of complex, interconnected subprocesses, such as localization, mapping, and planning, to navigate unknown environments. Maintaining accurate state estimates, a process critical to mapping

---

\*These authors contributed equally to this work.

Received: 29 June 2022; revised: 30 September 2022; accepted: 17 November 2022; published: 10 January 2023.

**Correspondence:** Harel Biggie, Computer Science, University of Colorado Boulder, Email: [Harel.Biggie@colorado.edu](mailto:Harel.Biggie@colorado.edu)

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 © 2023 Biggie, Rush, Riley, Ahmad, Ohradzansky, Harlow, Miles, Torres, McGuire, Frew, Heckman and Humbert

DOI: <https://doi.org/10.55417/fr.2023004>

and exploration, is exceptionally challenging in subterranean environments. GPS is unavailable for obtaining position estimates and visual-based localization methods can be affected by varied lighting conditions and environmental factors such as heavy dust, fog, or smoke. Subterranean environments, such as mines and caves, are often unstructured and contain hazardous obstacles, making navigation with ground vehicles challenging. Additionally, aerial vehicles can be exceptionally difficult to deploy in tight constrained underground spaces due to self-induced propeller wash. The DARPA Subterranean Challenge (SubT) (DARPA, 2022) aimed to spark new developments in the areas of autonomy, perception, mobility, and networking in subterranean environments. In the following work, a scalable multiagent autonomy solution for subterranean exploration developed by the University of Colorado’s Team MARBLE for the SubT Challenge is presented, along with critical lessons learned and developments made along the way.

DARPA designed the SubT Challenge to simulate search-and-rescue scenarios in unknown subterranean environments, and consisted of three domain-specific circuit events, Tunnel Circuit, Urban Circuit, and Cave Circuit, followed by the Final Event, which was a combination of three subterranean domains. Teams were challenged with developing robot platforms to deploy in each of the events in search of sets of predefined artifacts, such as backpacks or a Bluetooth signal produced by a cell phone. Correct identifications, consisting of an artifact classification and location to within a 5m sphere of the ground truth location, resulted in a point scored. Placement in the competition was determined by the team which could score the most points over a series of one hour deployments. For the Final Event, teams were limited to a single “human supervisor” (HS) who was able to interact with the systems and visualize any incoming data. Adding to the challenge, teams had a limited window of time in which five team members could set up and initialize robots at the entrance to the course.

Team MARBLE’s initial approach to subterranean exploration for the Tunnel and Urban Circuit Events is presented in (Ohradzansky et al., 2021). Initially, a graph-based planning and exploration strategy was implemented, the details of which are presented in (Ohradzansky et al., 2020). This solution is suitable for tunnel-like mines that have mostly planar corridor-junction structures, because the environment can be easily represented by a graph of nodes and edges. A scanning lidar was used to center robots in corridors while navigating edge sections as well as avoid obstacles in the environment. However, this approach lacked multiagent coordination, resulting in significant overlap of explored regions by different agents. For the Urban and Cave Circuit Events, a three-dimensional volumetric map representation of the environment was generated and used in a frontier-based exploration strategy (Ahmad et al., 2021a; Ahmad et al., 2021b). In this approach, the exploration rate of the robot is maximized using a frontier-based (Yamauchi, 1997) sampling technique and a fast marching cost-to-go calculation (Sethian, 1999) to select goal poses and plan paths to them in three dimensional space. An artificial potential function based obstacle avoidance algorithm enables the robot to path follow while avoiding small obstacles in the environment. Our initial approach also implemented limited forms of multiagent coordination in the form of agents sharing goal points and paths.

Other teams developed impressive solutions to the initial Tunnel and Urban Circuit challenges. Team CSIRO, a collaboration between the Commonwealth Scientific and Industrial Research Organization (CSIRO), Emesent, and Georgia Tech, presents a unique homogeneous sensing solution (Hudson et al., 2021). In this approach, heterogeneous teams of robots, including both ground and aerial platforms, share sensor information as a part of a decentralized multiagent SLAM system. Initially, exploration was handled through manual waypoints commanded by the human supervisor, but eventually an autonomous exploration algorithm was implemented (Williams et al., 2020). A common perception module, the CatPack, is used across all ground vehicles for easy reuse of the autonomy stack across different platforms. Similar to Team CSIRO, Team CERBERUS also used a heterogeneous team of ground and aerial platforms (Tranzatto et al., 2022a; Papachristos et al., 2019; Tranzatto et al., 2022b). In their approach, map information from different agents is fused into an optimized global map that is shared back to the agents (Khattak et al., 2020). Similar to the work presented in (Ohradzansky et al., 2020), other teams used graph-based planning approaches

for global navigation (Dang et al., 2019; Dang et al., 2020). Other noteworthy teams and their approaches to autonomous subterranean navigation include Team CoSTAR (Santamaria-Navarro et al., 2020; Ebadi et al., 2020; Agha et al., 2021; Otsu et al., 2020), CTU-CRAS (Rouček et al., 2020), Team Explorer (Scherer et al., 2021), and NCTU (National Chiao Tung University) (Huang et al., 2019). Additional discussions on the challenges, novel developments, and lessons learned from the Tunnel and Urban Circuit Events are included in the following works (Miller et al., 2020; Lajoie et al., 2020). One common theme common to many of these approaches is the use of heterogeneous teams of agents with multimodal sensing solutions. By diversifying the sizes, types of locomotion, and sensor modalities of individual robots, teams can be more versatile when faced with varied environments, each with a unique set of challenges. This ability to be flexible and adapt to the needs of the mission is one of Team MARBLE’s driving philosophies.

The format of the challenge necessitated advancements in platform design, robust communication networks, intelligent planners, and a balance between autonomy and human decision making for robot fleets. Team MARBLE’s solution, which was developed over the course of three different circuit events and showcased at the Final Event, led to a third place finish. Specifically, we develop a heterogeneous fleet of autonomous robots, each capable of operating independently of human intervention. Our autonomy-first approach employs a lightweight graph-based planner that scales to large environments, and has been adapted to reason over dynamic environments, such as closing doors and passing agents, as well as take advantage of multiagent coordination. Information sharing between agents is accomplished via a custom mesh networking solution with fast reconnect times and configurable message prioritization. The same network provides the human supervisor real-time visibility and control of the mission. Our autonomy-first philosophy inspired an autonomous mission management system that frees the human supervisor from agent-level micromanagement and deepens the opportunity to strategically accomplish mission goals. In this work we will present each of the components of our proposed autonomy system, as well as a detailed performance analysis of our solution and lessons learned along the way.

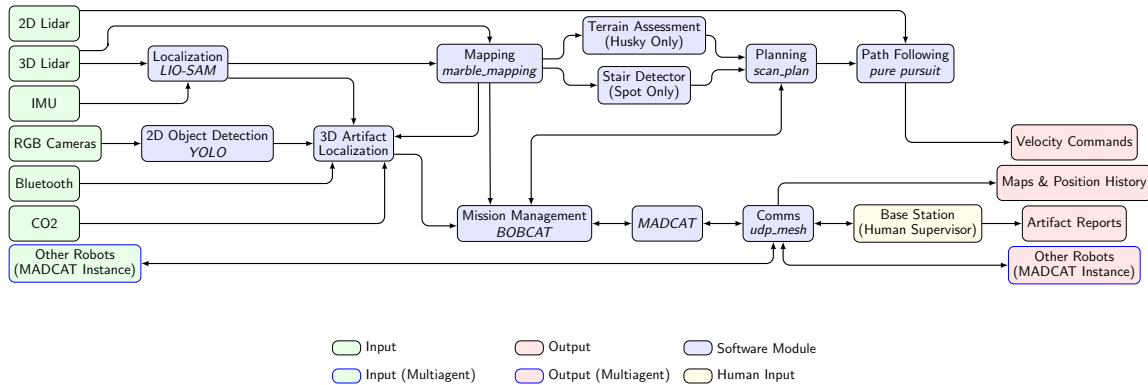
This paper is organized with the following structure. First, an overview of our system is provided in Section 2. The robot platforms developed for the Final Event are described in Section 3, followed by a description of the localization system in Section 4, and the artifact detection system in Section 5. The multiagent components of our system include the volumetric mapping pipeline in Section 6 and the graph-based path planning over those maps in Section 7. Information transmitted among agents, as well as to and from the human supervisor, is mediated by the wireless mesh network communications system described in Section 8 and handled by the autonomous mission management system described in Section 9. Finally, we analyze how these systems performed in the DARPA SubT Challenge Final Event in Section 10 and discuss lessons learned in Section 11.

## 2. System Overview

In the following subsections, we present a high-level description of Team MARBLE’s approach to the DARPA Subterranean Challenge. First, the general concept of operation is described, followed by an overview of each of the major components of the developed autonomy solution. A full high level summary of the autonomy system can be seen in Figure 1.

### 2.1. Concept of Operations

Team MARBLE has emphasized development of multiagent autonomy solutions that are able to operate without requiring intervention from a human operator. This aligns with the goals of the DARPA Subterranean Challenge where intermittent or unavailable communications with agents from a base station or between agents is expected. Therefore, our solution is centered around robust single-agent autonomy, where independent robots are able to explore unknown environments and report back to a base station with collected information about the environment including map data and detected artifact locations. In communication-limited environments where information sharing



**Figure 1.** Overall block diagram showing the high level functionality of the autonomy stack. Inputs are shown in green and outputs are shown in red. Software package names are *italicized* and inputs and outputs which are shared between agents are outlined in blue. Terrain assessment and stair detection both add semantic information to the map but are only run on the Husky and Spot, respectively.

with other agents and the human supervisor may not be available, our agents persist and continue to execute the mission.

While it is important for single, independent agents to be able to explore autonomously, our solution incorporates several multiagent components to improve exploration efficiency. Our fleet is also designed to be opportunistic, capitalizing on communication links when they are available to amplify fleet performance. Multiagent coordination is an auxiliary capability that reduces redundant efforts when agents enter communication range with one another, and inform each other where they have been and where they plan to go next.

Our system’s performance can further be improved when communications are available which enables the human supervisor to have a holistic perspective of the specific search-and-rescue scenario. This holistic perspective empowers the human to make high-level contextual decisions through two types of intervention: directing the agent to a specific location by commanding a high-level waypoint, or teleoperating the agent by commanding low-level velocity signals. During the 60-minute Final Event Prize Run, Team MARBLE’s robot fleet was completely autonomous, with the exception of five strategic low-level human supervisor interventions. The balance between human input and autonomy is further discussed in Section 9.

## 2.2. Perception

A modular perception suite, shown in Figure 5b, was designed as the basis for the autonomy stack. The primary sensor is the Ouster OS1-64 lidar (Light Detection and Ranging), which provides 3D point clouds for mapping, localization, semantic mapping, and obstacle avoidance. A LORD Microstrain 3DM-GX5-15 IMU (Inertial Measurement Unit) is used to measure linear and angular acceleration of the sensor head for use in lidar-inertial state estimation. To identify visual artifacts, the systems are equipped with several FLIR Blackfly PGE-05S2C-CS cameras and an array of 5W dimmable LEDs for self illumination. The Husky platforms are equipped with four cameras facing forward, backward, to the left, and to the right. The Spot robots had a similar configuration, save omitting the rear camera due to occlusions caused by the custom-built compute and power management system.

## 2.3. Localization

Localization provides consistent pose information for many downstream autonomy processes including volumetric mapping, path planning, artifact detection, and multiagent coordination. However, ensuring reliable localization is difficult in austere underground environments. Because conventional vision-based solutions can be unreliable due to dark, feature-poor settings, Team MARBLE utilized

lidar-based methods and specifically tested and integrated LIO-SAM (Shan et al., 2020), which has fast online loop closures during long-duration missions. Several modifications were made to the system to improve localization accuracy and reliability, which are further discussed in Section 4. Methods used to align the robots with into a common reference frame are also presented in this section.

## 2.4. Exploration

The exploration algorithm generates safe and traversable paths that lead agents toward unexplored areas of previously unseen environments. The developed sampling based path planning algorithm is designed to be lightweight, so that it can operate on rapid exploration timescales, regardless of the extent of the environment. Computational efficiency is achieved by employing a bifurcated local-global graph for sampling unseen frontiers as well as a *good enough* strategy for final selection. In such time-constrained search and rescue settings, even humans will often make rapid decisions rather than dwell for long periods of time to make globally optimal ones. The details of the planning algorithm is described in Section 7.

The planner is also constructed to be flexible, so that additional capabilities could be scaffolded on top of the core algorithm. Integration of the planning algorithm with semantic mapping was critical for rough terrain and stair traversal, as explained in Section 6. Section 7.4 explains how this planning algorithm re-plans in dynamic environments, whether due to doorways that are being opened or closed, or fellow agents that are passing by. This capability is crucial for robust operation in real-world environments, which cannot assumed to be static. Finally, the planner is able to run more efficiently with multiple robots using multiagent coordination is covered in Section 7.3.

Agents follow paths via a modified pure pursuit controller (Coulter, 1992). The yaw rate command is computed by comparing the current agent’s heading against a lookahead point that is a fixed distance along the path. Forward speed is regulated based on local proximity to obstacles in the environment and the relative heading error to the lookahead point. This results in slower speeds when agents are in cluttered environments or experiencing large heading errors. A 2D (RP Lidar) was used for local obstacle avoidance on the Husky platform and the Spot platform had built in obstacle avoidance.

## 2.5. Mapping

Team MARBLE’s mapping framework is based on the open source *Octomap* package (Hornung et al., 2013), and has been customized with additional capabilities including map merging, transmission of difference maps, and encoding of semantic information. The core of the mapping framework is a log-odds based probability metric for occupied and unoccupied voxels or cells. These cells provide a 3D representation of the environment that is later used for navigation. This flexible framework enabled transmission of key environmental features such as rough terrain and the location of stairs efficiently through a bandwidth limited communication system. The details of our mapping system are provided in Section 6.

## 2.6. Artifact Detection

The artifact detection system precisely localizes visual artifacts using RGB sensors for visual classification and detection and lidar for the depth estimation. Nonvisual artifacts such as cellphones and gas are localized based on the position of the robot. A weighted median filter fuses together all detections which are sent to the human supervisor for final validation as described in Section 5.

## 2.7. Communication Systems

A mesh networking solution transmits data between robots and back the base station for the human supervisor to review. Standard 2.4 GHz 802.11 wireless radios based on the *ath9k* chipset are used

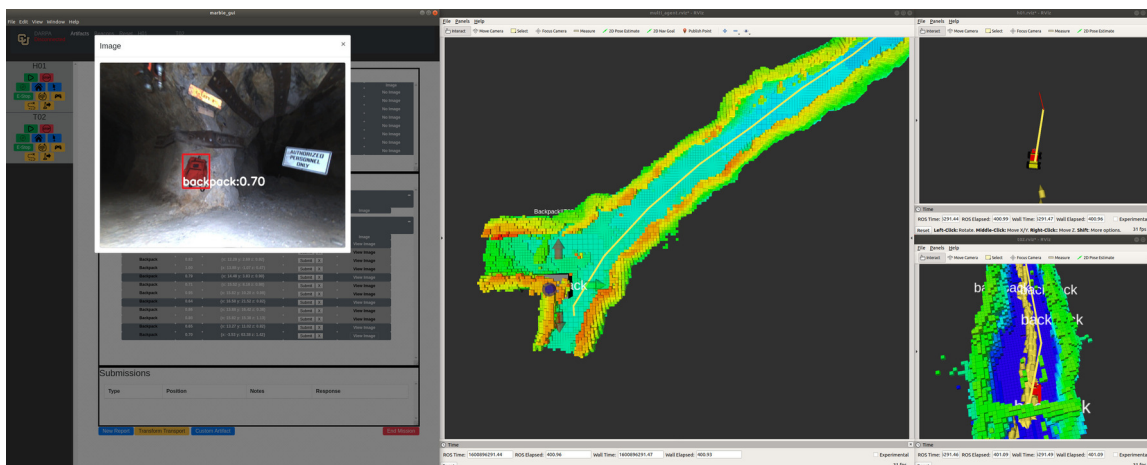
for the physical layer. The wireless radios are embedded in beacons that can be deployed from the back of the Husky platforms, allowing for ad-hoc mesh networks to be established. Meshing technology was provided by Meshmerize (Pandi et al., 2019) and a custom UDP based transport layer (*udp\_mesh*) was developed. Details of this innovative layer can be seen in Section 8.

## 2.8. Multiagent Coordination

Exploring unknown environments with multiple agents can be made more efficient through coordination, especially when agents are not within communication range. Sharing information across agents, such as explored regions, discovered artifacts, and current behavioral states, allows for more intelligent management of multiagent exploration. The framework called Multiagent Data Collaboration for Autonomous Teams (MADCAT), provides the multiagent data sharing capabilities required for the Subterranean Challenge mission (Riley and Frew, 2021), including transmission of relevant coordination data and maps, as well as map merging functionality and decision making for each agent. Additionally, MADCAT implements Behaviors, Objectives and Binary states for Cooperative Autonomous Tasks (BOBCAT), originally presented in (Riley and Frew, 2022), for high-level autonomy, decision making, and interfacing with the human supervisor. The MADCAT algorithm is discussed in more detail in Section 9.

## 2.9. Mission Management

The human supervisor is able to monitor the fleet’s progression through the unknown subterranean environment using a custom GUI operating on a computer at the entrance to the environment (base station). Current mission status of all agents in the field (“Reporting,” “Exploring,” “Home,” etc.) as well as their location in the global map are displayed whenever robots are within communication range of the base station. Additionally, the goal point and goal path for each robot is visible, allowing the supervisor to see the intent of each robot. The human supervisor is able to take over control of a given agent by either sending manual goal points or teleoperating the vehicle using a joystick interface with the base station. Reported artifacts are displayed per robot, with the type (survivor, cell phone, backpack, helmet, rope), position, confidence, submission result, and an image. The GUI also enables the modification of artifact classes and locations prior to submission to the DARPA scoring server. An example of the GUI interface is shown in Figure 2.




**Figure 2.** Example of the Human Supervisor’s interface. The custom GUI is on the left, showing a received artifact image. The middle is the multiagent RViz view, with all robots and the complete merged map. The right is a third-person follower for each robot (two in this case) with that robot’s original unmerged map.

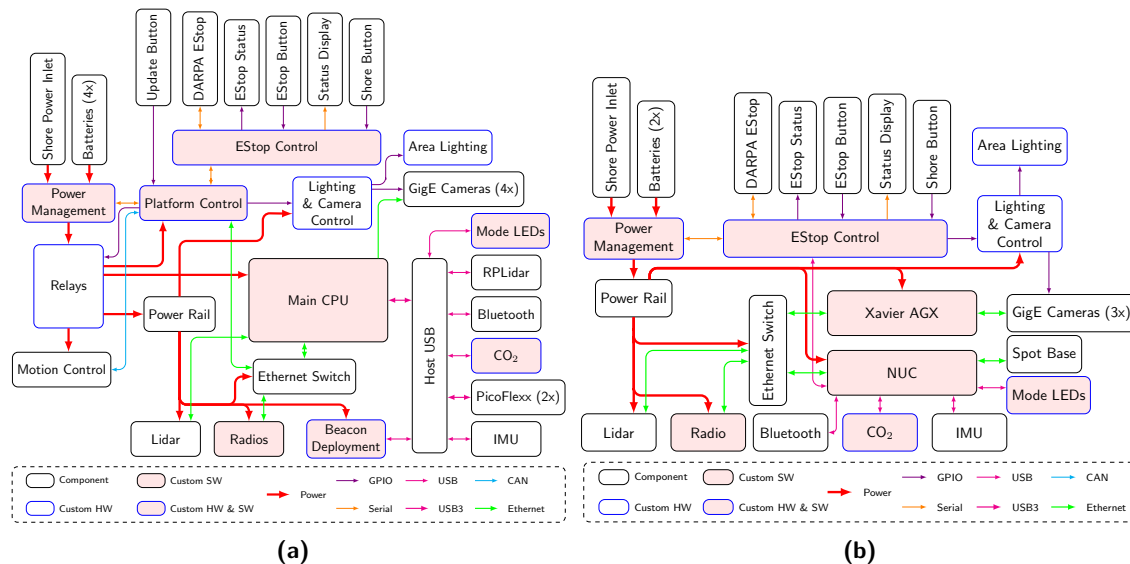
### 3. Platform Development

For the Final Event of the SubT Challenge, Team MARBLE deployed a heterogeneous fleet consisting of two Clearpath Husky A200s (H01, H02) and two Boston Dynamics Spots (D01, D02). Examples of each platform are shown in Figure 3. The Husky platforms are four-wheeled skid-steer ground vehicles capable of carrying heavy payloads, while the Spot quadrupedal “dog” platforms are agile, capable of climbing staircases and traversing uneven terrain. The Husky platform is robust and stable, with its generous payload budget allowing it to carry six communication beacons. The intended deployment strategy is to first deploy the Spot platforms to maximize exploration, and then the Husky platforms to establish the mesh communication network.


For processing power, each Husky is equipped with a 32-core AMD Ryzen CPU equipped with 128GB of RAM and 4TB of SSD storage, integrated into a complete platform as shown in Figure 4a. Dual NVIDIA GTX 1650 GPUs were used to accelerate object detection inference speed. The



**Figure 3.** The fleet is composed of two classes of robotic agents: (a) Clearpath Husky A200 and (b) Boston Dynamics Spot. Each platform carries a common sensor suite designed for exploration and object detection.



**Figure 4.** Power and signal routing diagrams of customized (a) Husky and (b) Spot platforms.



**Figure 5.** Final robot configurations, with (a) a deployment mechanism loaded with six communication beacons on Husky vehicles at the Final Event and (b) a modular perception suite installed on both Huskies and Spots.

primary computer on the Spot platforms is an AMD Ryzen 5800U with 64GB of RAM and 2TB of storage which is paired with a Jetson Xavier AGX to process the camera streams and perform artifact detection, following a similar integration as shown in Figure 4b. Many purpose-built components are common between platforms to reduce field maintenance efforts and a platform-specific code. Each system is outfitted with a custom power system, discussed in Section 3.2, which enables the ability to switch from a wired shore power supply to the onboard computer batteries. This leads to more efficient use of the onboard batteries, which are a limiting factor in the duration of field testing deployments.

### 3.1. Communication Beacons

Underground environments provide limited line-of-sight capabilities for wireless communications. As a result, Team MARBLE developed custom communication beacons to complement the custom multirobot coordination solution. This allows for robots to share information with the base station and other robots in the field. Each Husky platform is capable of carrying six beacons, each containing a single 2W Doodlelabs 802.11n radio as seen in Figure 5a. The autonomous beacon deployment mechanism relies on a latching solenoid release coupled with a novel passive system to gently lower each beacon to the ground to ensure maximum antenna height. Additional design details of the communication beacons can be found in Section 13.1 of the Appendix.

### 3.2. Power and Platform Control Systems

In order to support each platform’s sensor and processing needs, as well as meet DARPA equipment requirements for emergency stop systems, the systems integration efforts relied on several custom hardware and software components. Where feasible, these components are shared between the Husky, shown in Figure 3a, and the Spot platforms, shown in Figure 3b, reducing the development and validation efforts, as well as team member operational training requirements.

Of the custom capabilities developed, the power management subsystem deserves special mention. This component implements a hardware-interlocked, ideal diode system to permit downstream electronics to source power from either a wall-connected power supply, referred to as shore power, or onboard batteries. By switching to shore power, onboard systems can remain powered for development, testing, and analysis while the batteries are charged without carrying load. Further, the ideal diode component allows the battery packs to load share and charge independently. In contrast to a bus-tied battery system, this ideal diode design prevents high-energy charge equalization between packs and allowed each battery pack’s onboard management board to function



independently. The system also enables live monitoring of current consumption and battery voltage, as well as intelligent e-stop management which ensures the robot cannot exit the emergency stopped state while connected to shore power.

Emergency stop requirements dictated that each platform needed the ability to be stopped by a physical button, software, and via a DARPA deployed Xbee network. On the Husky control system the emergency stop system was integrated directly into the base controller. In contrast, the Spot platform’s emergency is tied into the available API to issue a “sit” command before terminating power to the motors which allowed the robot to be stopped gracefully.

Several important lessons learned emerged after three years of platform architecture development. We have highlighted the most critical lessons below and provide more details about the design and consequences of our platform compute systems in the Section 13.2 of the Appendix.

As part of our field testing campaign, we uncovered an issue where our USB-connected IMU was delayed in delivering measurements critical to localization performance. From an integration perspective, our IMU’s USB interface was implemented using a standard USB Communications Data Class Abstract Control Model (CDC-ACM) interface. Using CDC-ACM for IMU measurements was particularly problematic due to the way in which CDC-ACM uses *bulk* transport. USB has several methods of transferring data from device to host, including *interrupt*, *isochronous*, and *bulk* transport. CDC-ACM uses *bulk* transport, which does not include any guarantee for on-time delivery of data. As a consequence, during high CPU load, IMU measurements were occasionally delayed and resulted in localization error. In contrast, *interrupt* and *isochronous* transports are regularly serviced and can deliver on-time data. This problem could be solved in future deployments by either replacing the *bulk* interface with an *interrupt* interface or by using legacy serial interfaces such as RS-232 (not typically available on small form factor computing units). However, in practice, we found adjusting the IMU timing as described in Section 4, was sufficient and did not require engineering new firmware.

Another critical piece for reliable localization and consequently navigation is sensor synchronization. Synchronization is fundamentally necessary in order for our onboard sensors to communicate with their respective computers, and for those computers onboard individual agents to communicate with each other and the base station computer. The technical implementation of our solution is detailed in Section 13.3 of the Appendix.

## 4. Localization

One of the major challenges in the DARPA Subterranean Challenge is ensuring reliable localization across a diverse set of austere environments. Localization is a critical process for an autonomous system, as it provides pose information to downstream autonomy processes including volumetric mapping, path planning, artifact detection, and multiagent coordination. Section 4.1 details the simultaneous localization and mapping solution that was integrated into the autonomy stack, and Section 4.2 describes the process used to align all robots to the common DARPA reference frame.

### 4.1. Simultaneous Localization and Mapping

Simultaneous localization and mapping has relatively mature vision-based solutions (Leutenegger et al., 2015; Nobre et al., 2017; Qin et al., 2018), thanks to advances in feature extraction (Cheung and Hamarneh, 2009; Bay et al., 2008; Zhan et al., 2018). However, in mission-critical applications such as underground search and rescue, visual-inertial solutions are not reliable enough when faced with irregular lighting, specular highlights, and feature-poor scenes. Recent work has illuminated the possibility of leveraging thermal-based odometry estimation in visually degraded environments (Khattak et al., 2019; Wisth et al., 2021).

Because underground environments are typically rich in geometric features, lidar-based localization solutions are a compelling alternative. Some spaces though, such as a smooth tunnels and corridors, contain relatively few longitudinal features, and therefore pose limits to lidar-based

perception. Single-echo lidar also struggles in austere environments containing fog or smoke, though some recent work has focused on addressing these limitations (Shamsudin et al., 2016).

For the Final Event, Team MARBLE transitioned from Google Cartographer (Hess et al., 2016) to LIO-SAM (Shan et al., 2020), since its faster online loop closures during long-duration missions results in greater localization accuracy. Extensive testing was conducted in many different environments including parking garages, academic buildings, gold mines, and outdoor environments. The fast, lightweight loop closure performance can be attributed to performing scan-matching on a local level rather than a global level. LIO-SAM additionally performs IMU pre-integration to deskew point clouds, yielding better initialization for lidar odometry estimation. Because localization is the foundation to many autonomy modules, it was imperative to validate LIO-SAM’s performance onboard the Spot and Husky platforms during large-scale, long-duration missions. Some examples of such validation efforts are shown in Section 13.4 of the Appendix.

Several modifications are made to the system to improve localization accuracy and reliability. First, the IMU and lidar sensors are fastened to a 6061 aluminum sensor plate, with a mounting configuration that is common between Huskies and Spots. By specifying the relative transform between the two sensors to a high precision, the need for extrinsics calibration is reduced. Specifically, the mounting configuration consists of a tight-tolerance, precision-ground plate with a flatness tolerance of  $0.005''$ , which greatly improves the roll and pitch alignment between the two sensors. By using a high-quality MEMS IMU and such precise sensor mounting, the LIO-SAM parameter specifying how much to weight IMU roll, pitch, and yaw measurements relative to lidar odometry was increased by a factor of 100. Taken together, these modifications greatly reduce accumulated rotation and translation drift, enabling smooth autonomous operation across long missions.


Secondly, LIO-SAM requires sensor timestamps to be aligned and sensor rates to be consistent. In particular, if IMU message rates fluctuate too greatly, the IMU pre-integration factors (Forster et al., 2015) can fail and lead to LIO-SAM instability. To reduce sensitivity to fluctuating IMU sensor rates caused by USB transmission delays (described in Section 3.2), the IMU timestamp assignment is adjusted when messages are not received within 15% of the nominal rate. Additionally, the lidar sensor is synchronized with the onboard computer via PTP as discussed in Section 13.3 of the Appendix. These two timing solutions reduce the probability of erroneous measurements and greatly improve the stability of LIO-SAM.

## 4.2. Common Reference Frame Alignment

Accurate multirobot alignment is a core design decision for the MARBLE localization, mapping, and planning systems. Robots are required to share globally aligned map data for planning and navigation. In addition, it allows the human supervisor and robots to share global coordinates for artifact locations relative to the DARPA-provided world frame. In order to align with the DARPA frame, Apriltags (Malyuta et al., 2019; Brommer et al., 2018; Wang and Olson, 2016), retro-reflective targets, and Leica Total Station (LTS) reflectors are attached to a gate with relative transforms to the DARPA origin frame. The full setup of the gate setup is illustrated in Figure 6. The global frame was assumed to be aligned with gravity, but each team was responsible for aligning yaw, and XYZ-translation from their robots to the common DARPA frame. In the context of the SubT Challenge, the DARPA frame was purely used to align robots into the measured ground truth frame for artifact scoring and map accuracy analysis. However, in practice, an accurate initial alignment between robots results in more reliable multiagent coordination and global merging maps.

In order to align with the common reference frame, MARBLE primarily relied on the LTS reflectors. Based on conventional trigonometry and the assumption of needing to maintain less than 5m of error over the course of a 1km linear distance, it is determined that an initial alignment target required less than  $0.29^\circ$  of error.

To align the robot, 3 reflective prisms are attached to each robot, and their positions are scanned with an LTS. These points  $V_{LR}$  are then compared with a ground-truth set  $V_R$ , determined by the relative locations of the prisms to the robots tracking frame via CAD. These two sets of points are



**Figure 6.** Figure of the MARBLE gate alignment setup including the Leica Total Station (LTS), gate, and an example robot. Transforms from the LTS to the robot  $T_{RL}$ , from the LTS to the world frame  $T_{WL}$ , and the resulting transform from the world frame into the  $T_{WR}$ .

used to estimate the transform between the LTS recorded positions and the assumed positions by minimizing across the pose  $\mathcal{X}$  to solve:

$$\operatorname{argmin}_{\mathcal{X}} \sum \|V_{LR} - V_R T_{RL}\|. \quad (1)$$

The result is the robot’s position in the LTS frame  $T_{RL}$ . An additional calculation is used between scanned points of the gate  $V_{LG}$  and the provided coordinates  $V_W$  are used to solve for the gate’s position in the Leica frame  $T_{WL}$  using the equation:

$$\operatorname{argmin}_{\mathcal{X}} \sum \|V_{LG} - V_W T_{WL}\|. \quad (2)$$

Both minimization problems were solved using Horn’s absolute orientation method (Horn, 1987), a closed form solution to least squares alignment problems. Given these transforms, the robot’s position in the world frame was calculated by inverting the robot to LTS transform:

$$T_{WR} = T_{WL}(T_{RL})^{-1}. \quad (3)$$











To further reduce the impact of minor errors in either prism localization or low observability, these transforms are altered slightly by each robot. The LTS-predicted pitch and roll is substituted with an estimated pitch and roll from the lidar-inertial localization system, largely based on the initial measurements of the IMU.

After these adjustments, yaw estimates had the largest impact on our resulting transforms. Because yaw error has the potential to propagate to large translational discrepancies at far distances, it became imperative to modify our system. The solution involves increasing the lateral spacing of the prisms mounted on the robots, and is described in more detail in Section 13.5 of the Appendix.

## 5. Artifact Detection

A core component of the SubT Challenge is the detection and localization of objects that could potentially indicate human presence. Each artifact needed to be reported within a 5m radius of the ground truth location. To achieve this requirement a lidar-inertial based state estimator as described in Section 4 is used. Robots are put into a common reference frame based on survey-grade measurements from a Leica Total Station (LTS) and objects are projected using the mapping

**Table 1.** Sensing modalities, that Team MARBLE utilized (+) and did not utilize (−) for detecting the ten artifact classes. Blank entries indicate sensing modalities that are not useful for detecting specific artifact classes.

	Artifact Class	Visual	Thermal	Wireless	CO <sub>2</sub>
	Survivor	+	−		
	Cell Phone	−		+	
	Backpack	+			
	Drill	+			
	Fire Extinguisher	+			
	Gas				+
	Vent	+	−		
	Helmet	+			
	Rope	+			
	Cube	−		+	

framework described in Section 6. The available sensing modalities for various artifacts are discussed in Section 5.1, the visual detection system is described in Section 5.2, and the nonvisual detection system is explained in Section 5.3. The resulting performance of the artifact detection system at the Final Event is detailed in Section 10.4.

## 5.1. Sensing Modalities

Table 1 shows the classes of artifacts present at the Final Event along with the types of sensing modalities capable of detecting each artifact. Each robot in the fleet is equipped with RGB cameras, Bluetooth modules, and CO<sub>2</sub> sensors which enable the detection of all classes of artifacts using a minimal sensor suite. The visual detection system is not trained to detect either the cell phone, due to its small form factor, or the cube artifact which is detectable using Bluetooth. The cube artifact has rotating colors which pose significant challenges for visual detection methods.

## 5.2. Visual Detection

Visual object detection is a well-researched problem in computer vision and state of the art detectors are capable of identifying objects in both 2D and 3D. Common 2D detectors are typically based on Convolutional Neural Networks (CNN) (Zou et al., 2019), such as region proposal-based networks like Fast R-CNN (Girshick, 2015). Typically these networks require multiple passes over an image to classify an object and then detect where the object is in the image. In contrast, YOLO (Redmon et al., 2016) performs both classification and detection in a single regression making it a significantly faster detection: 0.5 frames per second (FPS) for Fast R-CNN and 45 FPS for YOLO. Object detectors operating in 3D typically use point clouds obtained from a lidar and until recently were limited to classification rather than full detection (Maturana and Scherer, 2015; Qi et al., 2017). Extensions to 3D classifiers such as Voxelnet (Zhou and Tuzel, 2018) and PointRCNN (Shi et al., 2019) are capable of performing object detection on powerful GPUs. These GPUs are impractical from both a size and power consumption standpoint for mobile robots. We selected the Yolo V3 (Redmon and Farhadi, 2018) model due to the fast and accurate nature of the YOLO (Redmon et al., 2016) family of networks.

Specifically, for classification and detection, the visual pipeline utilizes a YOLO V3 Tiny (Redmon et al., 2016) model with custom trained weights. The model is optimized for Nvidia TensorRT

acceleration and we infer images at a resolution of 608x608. The Husky platforms are able to perform inference at 60 FPS on a GTX 1650 based on Nvidia’s Turing architecture with 896 CUDA cores and 112 RT cores. The Spot platform uses a Nvidia Jetson Xavier AGX based on the Volta architecture with 512 CUDA Cores and 64 Tensor cores to perform inference at 40 FPS. These GPUs were chosen to balance performance against size and power constraints for on-board compute. The TensorRT YOLO detector outputs a message containing the detected artifacts as well as the coordinates of their bounding boxes.

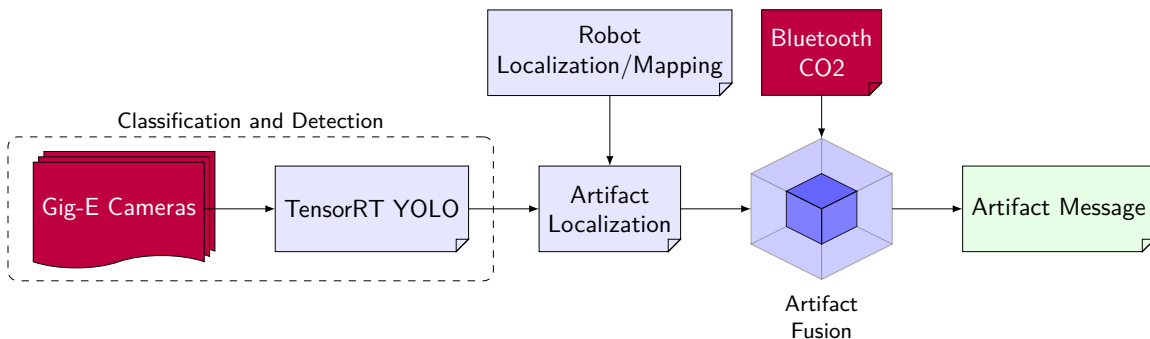
A systematic procedure targeted at low-light conditions is used to train the model. At each location, data was collected using three different brightness levels to minimize the impact of lighting conditions on the model’s performance. Specifically, images were taken from past circuit events as well as separate field exercises. Images with excessive motion blur were subsequently filtered out and the data was later augmented with images that contained false positive defections. The full details of our training procedure can be found in Section 13.6 of the Appendix.

Depth registration is performed using *marble\_mapping* as described in Section 6 which is generated by the Ouster 64-beam lidar. Utilizing an Octomap based framework allowed us to avoid implementing any additional filtering due to the probabilistic nature of the map. Additionally, the Octomap structure aggregates scans into the map with temporal memory. This important feature resolves the inconsistency between the 33.2° vertical field of view of the Ouster and the 68° vertical field of view of the cameras. At further distances, the agent is able to incrementally build out regions near ceilings and floors, overcoming the vertical blind spots of the Ouster. Essentially, this temporal memory allows us to decouple the depth measurement from the visual artifact detection. The biggest drawback of this approach is the potential for an additional 0.15m of error on each detection due to the voxel resolution. However, this error figure still falls within the design constraints of localizing an object to within 5m of its desired location.

After 3D coordinates are obtained via the Artifact Localization node, we run a weighted median filter in the world coordinate frame to de-noise the projected location within the Artifact Fusion node in Figure 7. Each localized artifact is considered to be part of the same measurement if it is the same class as a previous measurement and within 5m of that measurement. We then require five to 10 positive detection events and use the median position as the reported position to the human supervisor. The final detection is published in a custom ROS message which contains this position as well as a compressed version of a corresponding camera image and associated bounding box. The full overview of the artifact system can be seen in Figure 7.

### 5.3. Nonvisual Detection

Cell phone, cube, and gas reports are also fused using a weighted median filter. The Bluetooth and CO<sub>2</sub> detections are simply localized to the position of the robot at the time of detection. Bluetooth detections are also grouped together by unique SSIDs and gas detections within 10 m



**Figure 7.** Overview of the artifact detection system. Sensor inputs are shown in red and outputs are shown in green.

of another detection are assumed to have originated from the same source. The final positioning of these nonvisual artifacts relies on input from the human supervisor. Our human supervisor interface was designed to easily allow for movement of reported artifacts based on features observed in the map by the human operator. The details regarding the accuracy and success rate of these reports can be found in Section 10.4.


## 6. Mapping

Team MARBLE’s custom mapping package, *marble\_mapping* (Riley, 2021) is based on *Octomap* (Hornung et al., 2013) and is used to generate 3D occupancy grid representations of the world. The environment is subdivided into voxels, or cells which are marked as either occupied, free, or unknown using a probabilistic log-odds based model operating on sensor returns. The output of *marble\_mapping* is a direct input to the path planner and also provides depth measurements for visual artifact detection, as well as situational awareness for the human supervisor. The Octree (Meagher, 1982) structure of Octomap’s occupancy grids makes storing and transmitting maps more efficient than other representations such as point clouds; this efficiency is highly desirable when trying to transmit maps over low bandwidth mesh networks. *Marble\_mapping* extends *Octomap* by enabling map differences for low bandwidth transmission, map merging between multiple robots, and the addition of semantic information.

### 6.1. Difference-Based Map Merging

Despite the efficient encoding of the Octree data structure, regularly transmitting full volumetric maps of the explored space is impractical in bandwidth-constrained subterranean environments. Map differences are both a natural solution to reduce bandwidth, and have been shown to facilitate efficient data transfers (Sheng et al., 2004). In the *marble\_mapping* package, modifications to Octomap package were made to generate differences between different map sections, or “diff maps” shown in Figure 8. The implementation allows for diff maps, or smaller Octree structures, to be created at a predetermined rate, and contains all the mapping data for that time interval. The sum of an agent’s diff maps make up its “self” map and the differences can be transmitted to other agents. These differences are later merged into the robot’s “merged map” in the map merging process which is shown in Figure 8.

Merged maps generated from multiple agents are important both for a more complete view of the environment, and they also reduce redundant coverage in coordination strategies (Ko et al., 2003;



**Figure 8.** Sequential difference maps from top left to bottom right, with the final map on the far right constructed in real time for comparison. The diff maps can be merged to fully reconstruct the original map shown in the bottom right.

Simmons et al., 2000; Zlot et al., 2002). The *marble\_mapping* package enables map merging both for individual agents and on the base station which allows agents to intelligently act on the data and provides a holistic view for the human supervisor. The system does not re-align maps prior to merging, as it assumes agents are already in a common reference frame as described in Section 4.2. The lack of a re-alignment feature has the potential to cause one agent’s map to block pathways in a receiving agent’s map. To mitigate this, each agent prioritizes its own map by only appending cells from other maps into “unknown” areas. Areas that have already been “seen” by the agent are left untouched which prevents misaligned data from blocking free space. In cases where this mitigation procedure is not enough, such as narrow hallways, or a complete loss of localization by an agent, “bad” map diffs can be removed by the human supervisor using the base station GUI described in Section 9.

## 6.2. Semantic Mapping for Terrain-Aware Navigation

While the volumetric-based mapping produced by the Octomap framework provides the high-level structure of the environment, its resolution, set to a voxel size of 0.15m, is too coarse to capture details needed for high fidelity motion planning. In order to augment the existing *marble\_map* with terrain information, Team MARBLE evaluates the traversability of a given voxel using the normal and curvature values from raw point clouds. The planning solution is then able to utilize this semantic information to plan safe paths as described in Section 7. An additional label is attached to each voxel which enables the semantic labeling of staircases for the Spot platform.

Early approaches to evaluating the traversability of an environment include elevation based maps based on a 2D lidar (Ye and Borenstein, 2003) but are unable to take advantage of modern 3D sensors. The traversability classifier presented here is largely based on the Grid Map framework presented in (Fankhauser and Hutter, 2016), which evaluates the slope and roughness of point cloud regions to generate a multilayer surface map but only creates a 2D grid rather than a 3D volumetric map. Other fielded approaches in subterranean environments include “virtual surfaces” on occupancy maps (Hines et al., 2021) and Conditional-Value-at-Risk metrics, such as collision, step size, tip over, and slippage, which are incorporated into a dense 2.5D gridmap (Fan et al., 2021). These dense methods typically come in the form of high-resolution local maps, which enable more precise locomotion over varied terrain. An alternative approach presented in (Krüsi et al., 2017) computes paths with continuous curvature over raw point clouds. However, by computing semantic traversability information, our planning approach only required a low-resolution global map, greatly simplifying both mapping and planning systems and allows for sharing of semantic information between agents.


### 6.2.1. Traversability Classification & Map Integration

To estimate the traversability of a voxel, we segment the 3D point cloud produced by the Ouster lidar, and evaluate the unit normal vector  $\hat{\mathbf{n}}$  and curvature  $K$  of each point  $p$  at time step  $t$ . All calculations are performed with the aid of the *pcl* package (Rusu and Cousins, 2011) and a traversability value,  $\tau_{p,t}$ , is estimated for each point using Equation 4 where  $\hat{\mathbf{k}}$  is the gravity-aligned up vector,  $(1 - |\hat{\mathbf{n}} \cdot \hat{\mathbf{k}}|)^3$  is a measure of the slope of the terrain, and  $c_{\text{norm}}$  and  $c_{\text{curv}}$  are tunable parameters. The parameter values for the Final Event were set to  $c_{\text{norm}} = 40.0$ , and  $c_{\text{curv}} = 4.0$ , and  $\tau_{p,t} \in [0, 1]$ .

$$\tau_{p,t} = c_{\text{norm}}(1 - |\hat{\mathbf{n}}_{p,t} \cdot \hat{\mathbf{k}}|)^3 + c_{\text{curv}}K_{p,t} \in [0, 1] \quad (4)$$

Traversability is implemented in the Octomap framework using Equation 5 to estimate the traversability,  $\tau_{v,t}$ , of a given voxel,  $v$ , as a function of the voxel’s occupancy probability,  $P_{\text{occ},v,t}$ . The traversability estimate for the voxel is a linear combination its previous traversability estimate,  $\tau_{v,t-1}$ , and new estimate  $\tau_{p,t}$  for the points in the voxel. An example of this process is shown in Figure 9.

$$\tau_{v,t} = \tau_{v,t-1}P_{\text{occ},v,t} + \tau_{p \in \mathbb{V},t}(1 - P_{\text{occ},v,t}) \in [0, 1] \quad (5)$$



**Figure 9.** Traversability information (a) of section within the Edgar Experimental Mine in Idaho Springs, CO, USA, that contains railroad tracks. Raw traversability values of the lidar point clouds (top) are shown, where white is not traversable, and black is traversable. Resulting semantic map (bottom) illustrated nontraversable surfaces such as walls in white, traversable surfaces such as the ground in black, and semitraversable surfaces such as the railroad tracks in grey. Note that red voxels do not contain traversability data. An accompanying photo (b) of the section of mine is shown for reference.

### 6.2.2. Stair Classification & Map Integration

Semantic information on stairs is fused into the mapping framework using the open source *StairwayDetection* (Westfechtel et al., 2018) package and a binary Bayes filter (Thrun et al., 2005). Stair classification of point clouds via this approach consists of 4 major steps: (1) preanalysis, in which the point cloud is downsampled and filtered, normal and curvature is estimated for each point, and floor separation is performed; (2) segmentation via a region growing algorithm, which segments the point cloud into smooth regions; (3) plane extraction, in which the surfaces that make up the riser and tread regions of each stair step are extracted; and (4) recognition, where the tread and riser regions are connected and analyzed via a graph to determine whether they make up a valid set of stairs.

Stair detections are integrated into the map using a similar mechanism to the log-odds probability which determines occupancy in octomap. A binary Bayes filter (Thrun et al., 2005), shown in Equation 6, is used to estimate the probability that a given voxel is a part of a staircase. The extracted points from the stairway detector are modeled as measurements  $z$ , where  $P_{\text{stair}}(n|z_{\text{stair},t})$  is the probability that a voxel  $n$  is part of a staircase. The measurement through time step  $t$  is represented by  $z_{\text{stair},1:t}$  as shown in Equation 6a.  $\mathcal{L}_{\text{stair}}(n|z_{\text{stair},1:t})$  as shown in Equation 6b are the corresponding log-odds probabilities which are used for fast updates updates to the probabilistic estimate of each voxel. More details of the log-odds formulation are provided in (Hornung et al., 2013; Thrun et al., 2005). Our filter is tuned to prioritize true positive detections with the following parameters:  $P_{\text{stair}}(n) = 0.5$ ,  $\mathcal{L}_{\text{stair},\min} = -2.0$ ,  $\mathcal{L}_{\text{stair},\max} = 3.48$ ,  $\mathcal{L}_{\text{stair},\text{hit}} = 4.60$ , and  $\mathcal{L}_{\text{stair},\text{miss}} = -0.04$ .


$$P_{\text{stair}}(n|z_{\text{stair},1:t}) = \left[ 1 + \frac{1 - P_{\text{stair}}(n|z_{\text{stair},t})}{P_{\text{stair}}(n|z_{\text{stair},t})} \frac{1 - P_{\text{stair}}(n|z_{\text{stair},1:t-1})}{P_{\text{stair}}(n|z_{\text{stair},1:t-1})} \frac{P_{\text{stair}}(n)}{1 - P_{\text{stair}}(n)} \right]^{-1} \in (0, 1), \quad (6a)$$

$$\mathcal{L}_{\text{stair}}(n|z_{\text{stair},1:t}) = \mathcal{L}_{\text{stair}}(n|z_{\text{stair},1:t-1}) + \mathcal{L}_{\text{stair}}(n|z_{\text{stair},t}) \in [\mathcal{L}_{\text{stair},\min}, \mathcal{L}_{\text{stair},\max}], \quad (6b)$$

$$\text{where } \mathcal{L}_{\text{stair},\min} \in (-\infty, 0), \mathcal{L}_{\text{stair},\max} \in (0, \infty),$$

$$\mathcal{L}_{\text{stair}}(n|z_{\text{stair},t}) = \begin{cases} \mathcal{L}_{\text{stair},\text{hit}} > 0 \text{ on stairs} \\ \mathcal{L}_{\text{stair},\text{miss}} < 0 \text{ on nonstairs} \end{cases}.$$





**Figure 10.** Spot planning up a staircase using the estimated stair voxels in the Octomap shown in blue.

A raycast operation on the footprint of the vehicle is used to provide a binary signal indicating the robot is on stairs. Additionally, eigenvector decomposition is performed over each cluster of stair voxels to extract a straight path along the staircase. Figure 10 provides an example of the semantic classification of stair voxels, along with the corresponding path along the stairs. These triggers provide waypoints so that the local trajectory follower can navigate to the top of the staircase. It’s important to note that since stairs would generally be classified as nontraversable, a stair label takes precedence over a traversability label for the Spot platform, which is capable of walking up stairs. Additionally, this method requires sufficient lidar scans of the staircase, which is generally available when located at the bottom of a staircase but is not when facing the stairs leading down. As a result, detecting and navigating a descending staircase is not feasible with the current configuration, but could be with a wider field-of-view sensor or programmed forward pitching behavior of the Spot.

The *marble\_mapping* package enables the creation of difference-based Octomaps which allows for efficient transmission in underground environments. Furthermore the framework provides semantic and traversability information which the planner utilizes to ensure the robot is able to navigate safely. Details of the planner are described in Section 7.

## 7. Path Planning

Team MARBLE’s heterogeneous fleet relies on autonomous path planning onboard each agent to reduce the workload of the human supervisor. The path planner running onboard each agent generates safe and traversable paths that lead to unexplored areas. Paths are planned on the Octomap-based *marble\_mapping* framework described in Section 6. Team MARBLE used the same planner on all robots with the only difference being the collision-function depending on vehicle’s class. For instance, a wheeled robot cannot traverse stairs while a legged robot can. Existing methods discussed in Section 7.1 suffer computational costs that make it challenging to scale to large environments. Because the proposed planner is computationally efficient and minimally dependent on tuning gains, it performs well in large-scale environments. Our planner makes several significant contributions, such as light on-demand terrain assessment, which is discussed in Section 7.2, hierarchical solution-search that also incorporates position history-based multiagent coordination, which is discussed in Section 7.3, and handling of dynamic changes in the environment such as blocked passages, which is covered in Section 7.4.

### 7.1. Background

One of the widely-known methods (Yamauchi, 1997; Ahmad et al., 2021a) for autonomous exploration relies on explicitly detecting potential frontiers on an explored map, followed by a path

planned toward each cluster of frontiers. The method seeded significant developments in the area of autonomous robotic exploration since it was first proposed. However, this frontier-based method employs a computationally expensive optimization-based approach that plans paths to each frontier cluster, despite the fact that some may not be reachable.

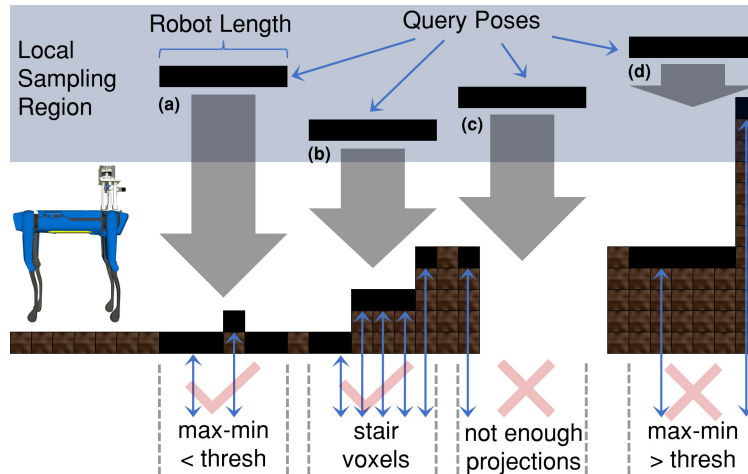
In recent decades, the planning community has witnessed significant advancements in more computationally efficient sampling-based approaches for path planning and exploration. One instance of such development is an exploration planner that uses Rapidly Exploring Random Trees (RRT) (LaValle, 2006) to sample an environment and chooses an *optimal* path from the set of sampled ones. The method samples the environment as a single batch, and therefore is not scalable to large-scale environments. A rectification of this limitation is recently proposed by (Dang et al., 2020) where a bifurcation approach is introduced for sampling and exploration. This approach implies that the environment is sampled only in the local neighborhood of a robot while simultaneously building a sparse graph that scopes the entirety of the explored map. The latter is essential to deal with local minima such as dead-ends and also to plan a path back home. Our autonomous exploration solution for the SubT Final Event relies on the principle of bifurcation with additional contributions in the terrain assessment, solution-search, dynamic obstacle avoidance and coordination.

The graph-based planners based on sampling and bifurcation approach use high resolution depth images to compute a 2.5D grid-based elevation map using the technique presented in (Fankhauser et al., 2018). This elevation map is further filtered to segment terrain characteristics such as slope, roughness and step (Wermelinger et al., 2016). The authors of such graph planners mention the scalability challenges with such computationally expensive approaches, which limit their terrain awareness to regions local to the agent. This further leads to challenges such as the planned path and the underlying graph being generated with an over-optimistic view of the terrain, consequently needing the robot to be backed up if it encounters impassible terrain.


## 7.2. Sample-and-Project Strategy

To rectify the terrain assessment challenges, a sample-and-project approach is followed, similar to the settling-based collision-check approach proposed in (Krüsi et al., 2017). We perform such checks on an Octomap with a resolution of 0.15m. At this resolution, all of team MARBLE’s robots were at least three voxels wide, providing a decent amount of robot footprint to project a robot’s pose on. The SubT Challenge rules highlight that the extremely narrow passages could be around 1m wide, with doorways as narrow as 36 inches. With this in mind, an Octomap voxel length of 0.15m was small enough to navigate narrow passages and large enough to be able to keep up with computational complexity of generating such a map in a large-scale environment. In case of a wheeled robot, each voxel in the Octomap is labelled with a roughness value which is obtained using high resolution point clouds as described in Section 6.2. However, on Spot legged robots dense roughness information is not required because of their onboard terrain assessment. Additionally, for Spots, encode semantic information about stairways into the map which overrides the default height parameters of the planner. With explicit labels, the planner is able to plan paths over built up staircases despite elevation changes the robot would not normally traverse over. More formally, for the legged robots capable of traversing stairs, each Octomap voxel maps to a label from the set {"occupied","unknown","free","stair"}, whereas in case of wheeled robots the label set is {"occupied","unknown","free","rough"}.

First, the environment is sampled in the local neighborhood of the robot using RRT\*, a variant of RRT with optimality considerations. Each tree sample is a robot position parameterized by the robot width and length. During sampling, the collision-checks are performed by vertically projecting a query sample to find the ground below it. Once the ground is found, the elevation change at the footprint of the sample is evaluated if there are enough projections on occupied voxels. In case of a wheeled robot, the average roughness information of the footprint voxels is also taken into consideration. For a legged robot, if a threshold amount of footprint voxels are labelled as "stairs," the sample is considered collision-free regardless of the elevation or roughness check. Expanding



**Figure 11.** A depiction of sample-and-project strategy for terrain checks on an OctoMap. The figure shows four different types of query poses. (a) and (b) depict collision-free samples whereas (c) and (d) are marked undercollision or nontraversable.



**Figure 12.** The figures highlight the planner solution search. The potential solutions for the planner include the paths leading toward the leaves of the RRT\* tree (blue) and the frontiers of the graph (green). (a) The first preference of the solution is the path that aligns best with the robot's exploration heading (bold blue path). In case this path has both sufficient volumetric gain and teammate separation, it is returned as a solution. (b) As a second preference, a thorough search is performed to find a local or global path that satisfies both constraints. (c) If none of the paths is found at both of the steps above then a path that satisfies the volumetric gain constraint is accepted as a potential solution.

an RRT\* requires checking path segments for collisions instead of isolated robot configurations. In order to check such a path segment, a set of robot configurations along the segment is checked for traversability. Figure 11 depicts the terrain assessment process.

### 7.3. Solution Search and Multirobot Coordination

At any replan iteration, a set of potential solutions include all of the local paths sampled using RRT\* and all of the global paths ending at the graph frontiers. The former set of solutions is represented by  $\mathcal{P}^L$  and latter solutions belong to the set  $\mathcal{P}^G$ . A set of all local paths that are leading the robot toward areas with greater than a defined threshold of volumetric gain and teammate separation are given as  $\mathcal{P}^{LV}$  and  $\mathcal{P}^{LS}$ , respectively. Similarly, a set of all global paths that are leading the robot toward areas with greater than a defined threshold of volumetric gain and teammate separation are given as  $\mathcal{P}^{GV}$  and  $\mathcal{P}^{GS}$ , respectively. These sets are highlighted in Figure 12. A typical approach to find an appropriate solution is to form an objective function with a combination of exploration objectives such as volumetric gain and exploration heading parameterized by the penalty gains. Volumetric gain calculation, however, is a computationally expensive operation that limits the amount of frontiers

a robot can process in a reasonable amount of time. Our approach relies on finding a *good enough* solution in terms of volumetric gain. To achieve multirobot coordination, the position histories of teammate robots on the network are used. If a path is leading a robot to a point such that the minimum distance of the point from the position histories of the teammate robots is more than the mapping range, it is guaranteed that new areas are being explored.

Following this intuition, the primary objective of Team MARBLE’s solution search method is not to optimize for the volumetric gain and the teammate position histories separation, but to accept a solution that has a satisfactory amount of volumetric gain and distance from the teammate position histories. The formal objective of the path planner is to output a solution that belongs to  $\mathcal{P}^{LV} \cup \mathcal{P}^{LS}$ ,  $\mathcal{P}^{GV} \cup \mathcal{P}^{GS}$ ,  $\mathcal{P}^{LV}$ , or  $\mathcal{P}^{GV}$  in the order of preference. The solution search process makes use of two different objective functions,

$$J^\alpha = c_0 \mathcal{D}^P(p^{\text{hist}}, p^{\text{cand}}) - c_1 |\theta^{\text{explore}} - \theta_{\text{mean}}^{\text{cand}}| - c_2 |h^{\text{explore}} - h_{\text{mean}}^{\text{cand}}|, \quad (7)$$

$$J^\beta = -c_1 |\theta^{\text{explore}} - \theta_{\text{mean}}^{\text{cand}}| - c_2 |h^{\text{explore}} - h_{\text{mean}}^{\text{cand}}| + c_3 \mathcal{G}^V(p^{\text{cand}}(1)) + c_4 \mathcal{D}^S(p^{\text{cand}}(1), p_1^{\text{hist}}, \dots, p_1^{\text{hist}}), \quad (8)$$

where  $J^\alpha$  is used to find a candidate path that aligns best with the current exploration heading of the robot and  $J^\beta$  is leveraged to perform a thorough search if required. The sets of points  $p^{\text{cand}}$  and  $p^{\text{hist}}$  represent a list of candidate solutions and the position history of a robot respectively. The exploration heading  $\theta^{\text{explore}}$  is calculated by averaging the most recent few points on  $p^{\text{hist}}$  of the robot. The mean heading and mean height of a candidate path are denoted by  $\theta_{\text{mean}}^{\text{cand}}$  and  $h_{\text{mean}}^{\text{cand}}$  respectively. The function  $\mathcal{D}^P$  accepts two paths as arguments and calculates the mean of minimum distance of all points along the first path with the second path. Furthermore, the function  $\mathcal{D}^S$  calculates the minimum distance of a candidate path from the position histories of all other teammate robots.

Algorithm 1 provides a deeper insight into the solution search steps. As a first step, a collision-free local path is found that best aligns with the direction of travel of the robot. This path is then checked if it has a satisfactory amount of volumetric gain and distance from the teammate position histories. If a good-enough solution is found at this step, the solution is returned and only a single volumetric gain function call is required. Therefore, we save significant computation time during most replan iterations. If a solution is not found at this first step then a more thorough search is performed, first through the sampled local paths and then through the global paths leading toward graph frontiers. This search is highlighted in Algorithm 1. The functions PLANLOCALLY() shown in Algorithm 2, and PLANGLOBALLY() shown in Algorithm 3, are responsible for outputting solutions that satisfy both volumetric gain and teammate separation constraints if possible, otherwise they output solutions that only satisfy the volumetric gain constraint. In the worst case, when neither constraint can be satisfied, paths with maximum teammate separation are returned as a contingency.

This attempt of finding a solution by breaking the potential solution space down into subsets instead of having one objective function to optimize over the entire space, helped us avoid extensive gain tuning. During testing and final event runs, we found our approach to be scalable for environments of various sizes without a need for tuning gains for different environment types. The details of the sampling-based path planner can be found in (Ahmad and Humbert, 2022). In this work, a simulation comparison of the proposed planner with an existing sampling-based planner (Dang et al., 2020) is presented, highlighting the improvement in scalability and computational efficiency.

## 7.4. Dynamic Replanning

Another challenge faced by the existing graph-based planners is that they rely on building a parallel graph representation of the environment. This representation does not naturally reflect changes in the environment, such as closed passages which were initially open at the time the graph is built. To handle this exception, the graph edges are labeled with a boolean representing its occupancy.

**Algorithm 1.** ScanPlan Solution Search.

---

```

1:  $p'_\alpha \leftarrow p^{\text{cand}} \in \mathcal{P}^L$  costing minimum  $J^\alpha$ 
2: if  $\mathcal{G}^V(p'_\alpha) \geq v_{\text{thresh}}^g$  and  $\mathcal{D}^S(p'_\alpha) \geq s_{\text{thresh}}^g$  then
3:   return  $p'_\alpha$ 
4: end if
5:  $p'_\beta \leftarrow \text{PLANLOCALLY}()$ 
6: if  $p'_\beta$  is non-empty and  $\mathcal{D}^S(p'_\beta) \geq s_{\text{thresh}}^g$  then
7:   return  $p'_\beta$ 
8: end if
9:  $p^l \leftarrow p'_\beta$ 
10:  $p^g \leftarrow \text{PLANGLOBALLY}()$ 
11: if  $p^g$  is non-empty and  $\mathcal{D}^S(p^g) \geq s_{\text{thresh}}^g$  then
12:   return  $p^g$ 
13: else if  $p^l$  is empty and  $p^g$  is empty then
14:   return  $p'_\alpha$ 
15: else if  $p^g$  is empty or
   ( $p^l$  is non-empty and  $\mathcal{D}^S(p^l) \geq \mathcal{D}^S(p^g)$ ) then
16:   return  $p^l$ 
17: else if  $p^l$  is empty or
   ( $p^g$  is non-empty and  $\mathcal{D}^S(p^g) \geq \mathcal{D}^S(p^l)$ ) then
18:   return  $p^g$ 
19: end if

```

---

**Algorithm 2.** PLANLOCALLY(). Returns a path in  $\mathcal{P}^{LV} \cap \mathcal{P}^{LS}$  or  $\mathcal{P}^{LV}$  in the order of preference.

---

```

1:  $J_{\min}^\beta \leftarrow +\text{inf}$ 
2:  $p^{\text{res}} \leftarrow \text{none}$ 
3: success  $\leftarrow$  false
4: for  $p_l \in \mathcal{P}^L$  do
5:   if  $\mathcal{G}^V(p^l) \geq v_{\text{thresh}}^g$  and  $\mathcal{D}^S(p^l) \geq s_{\text{thresh}}^d$  and  $\sim$  success then
6:      $J_{\min}^\beta \leftarrow J^\beta(p^l)$ ,  $p^{\text{res}} \leftarrow p^l$ , success  $\leftarrow$  true
7:   else if  $J^\beta(p^l) \geq J_{\min}^\beta$  then
8:     continue
9:   else if ( $\mathcal{G}^V(p^l) \geq v_{\text{thresh}}^g$  and  $\mathcal{D}^S(p^l) \geq s_{\text{thresh}}^d$ ) or ( $\mathcal{G}^V(p^l) \geq v_{\text{thresh}}^g$  and  $\sim$  success) then
10:     $J_{\min}^\beta \leftarrow J^\beta(p^l)$ ,  $p^{\text{res}} \leftarrow p^l$ 
11:   end if
12: end for
13: return  $p^{\text{res}}$ 

```

---

During exploration, the planned paths are constantly checked for collisions. If a planned path is under-collision, all edges in the local neighborhood of the robot are validated for collision and marked accordingly. Moreover, all occupied edges are checked for occupancy all the time when the planner finds some idle time which mostly happens when the vehicle is following a path. This enables the planner to take into account the cases where an occupied area is free again. In the case where the graph search is performed to plan a global path, the occupied edges are ignored.

## 8. Communication Systems

Effective communication with deployed systems from a fixed human operator is a crucial component of a complete robotic exploration system. While robots are capable of independent localization,

**Algorithm 3.** PLANGLOBALLY(). Returns a path in  $\mathcal{P}^{GV} \cap \mathcal{P}^{GS}$  or  $\mathcal{P}^{GV}$  in the order of preference.

---

```

1:  $\mathcal{P}^{GV} \leftarrow \emptyset, \mathcal{P}^{GVS} \leftarrow \emptyset$ 
2: for  $p^g \in \mathcal{P}^G$  do
3:   if  $G^V(p^g) \geq v_{\text{thresh}}^g$  and  $D^S(p^g) \geq s_{\text{thresh}}^d$  then
4:      $\mathcal{P}^{GVS} \leftarrow \mathcal{P}^{GVS} \cup \{p^g\}$ 
5:   else if  $G^V(p^g) \geq v_{\text{thresh}}^g$  then
6:      $\mathcal{P}^{GV} \leftarrow \mathcal{P}^{GV} \cup \{p^g\}$ 
7:   end if
8: end for
9: if  $\mathcal{P}^{GVS}$  is empty then
10:  return path  $p^g \in \mathcal{P}^{GV}$  with maximum  $D^S(p^g)$ 
11: end if
12:  $p_r^g \leftarrow$  path from  $\mathcal{P}^{GVS}$  leading to most recent frontier
13:  $p_c^g \leftarrow$  path from  $\mathcal{P}^{GVS}$  leading to closest frontier
14: if  $\text{PATHLENGTH}(p_r^g) \geq \text{PATHLENGTH}(p_c^g)$  then
15:  return  $p_c^g$ 
16: else
17:  return  $p_r^g$ 
18: end if

```

---


mapping, and artifact detection, the addition of a communication infrastructure is a force multiplier to enable human supervisory control, interrobot coordination, and timely artifact reporting. We developed a mesh network system to provide long-reach communications into underground environments which prioritizes reconnection times to maximize opportunities for data transfer.

## 8.1. Background

Previous work has developed several solutions to common problems encountered with deploying mesh networks, such as discovery and optimal routing. A wide variety of both closed-source and open-source solutions exist that include both hardware and software components. Mesh networking can largely be subdivided into three layers: physical, logical, and transport. We will detail several prominent open-source or commercially available options for each layer before describing our final solution.

From a logical layer standpoint, meshing layers lay between the physical transmission of frames over the medium and a higher-level protocol such as IP. For mobile robots operating in subterranean environments, a responsive mesh layer that minimizes lost link time is a major requirement due to the rapid movement of the robots. Further, to reduce integration effort, a mesh layer that operates at layer 2 of an OSI stack (for Standardization, 1996) is desirable to allow transparent use of higher-level protocols such as ARP and IP. Typically, meshing algorithms such as OLSR (Clausen et al., 2003) and AODV (Perkins et al., 2003) select a single best path for routing between nodes which hinders algorithmic performance in dynamic environments. A more recent example of a single-path logical meshing layer is Better Approach to Mobile Ad-hoc Networking-Advanced (*batman-adv*) (Seither et al., 2011), an open source implementation of a layer 2 mesh. In contrast to *batman-adv*, *meshmerize* (Pandi et al., 2019) provides multiple paths between nodes to ensure a reliable connection while still operating at layer 2; these multiple paths allow for a dramatic decrease in reconnect times when mesh topology changes. We relied on *meshmerize* as our layer 2 meshing solution in cooperation with Meshmerize GmBH.

Only transport layers designed for ROS were considered for ease of integration with the rest of the autonomy stack. In a traditional networked ROS architecture, a single computer runs a main node known as the *rosmaster* that coordinates the publish-subscribe mechanisms. When a node wishes to exchange data with another node via named *topics*, the master is consulted to determine



**Figure 13.** Several different types of ROS architectures. Red lines indicate data transfer, black lines indicate directory management, and blue lines are data paths that cross network segments. A basic, single-master ROS network node graph is shown in (a). A *fkie\_multimaster* multimaster ROS network node graph is shown in (b). In contrast to both (a) and (b), (c) shows how *udp\_mesh* creates a single virtual channel between nodes, shown in green, to implement data prioritization.

the computer to connect to, as in Figure 13a. A single *rosmaster* serves as a central directory of nodes and topics; when a subscription to a topic is requested, a list of publisher nodes is returned so that point-to-point TCP connections can be made directly between publisher and subscriber. These direct TCP connections break down when systems are linked over unreliable mesh networks which necessitates the need for an alternative transport mechanism.

One open source transport layer *multimaster\_fkie* (Juan and Cotarelo, 2015) solves the discovery and advertisement problems using multicast packets and specialized nodes on each machine with an architecture shown in Figure 13b. However *multimaster\_fkie* does nothing to establish prioritization of data flow. With the standard TCP transport provided by ROS, there is no centralized means of monitoring internode connections to arbitrate data priorities. Prioritization is crucial for monitoring the robot fleet in intermittent communication situations. Mission critical data such as artifacts needs to make it through to the human supervisor before other auxiliary data such as odometry and maps.

One alternative to *multimaster\_fkie*, *Pound*<sup>1</sup> (Tardioli et al., 2019), is specifically designed for use in unreliable mesh networks and implements many of the desired requirements. However, *Pound* relies on hardcoded topic names and fixed addressing information, which limit the flexibility of the system. Alternatively, *nimbro\_network* (Schwarz et al., 2016) implements a similar set of functions with regards to transport over wireless networks, but omits prioritization. Crucially, *nimbro\_network* still utilizes TCP for reliable interrobot communication, preventing adaptation of core TCP behavior (particularly retransmits) to unreliable mesh networks; UDP links are only used for nonguaranteed data delivery.

## 8.2. UDP Mesh

The main innovation in our system is our transport layer, *udp\_mesh* which allows for runtime reconfiguration, implements prioritization, and re-implements reliable communication over UDP to allow for more refined control over retransmits and fragmentation. Fundamentally, the *udp\_mesh* layer uses only unicast and broadcast UDP datagrams to implement higher-level services without requiring multicast support. In principle, multicasting would offer a performance benefit by reducing broadcast traffic. However, in a wireless mesh environment, these potential gains are offset by multicast group membership management overhead.

### 8.2.1. Discovery and Address Resolution

*Discovery* is the process of identifying nodes that are available for communication. We implement discovery through the use of a periodic heartbeat broadcast that advertises the node’s availability and provides name resolution information. In concept, this service is similar to the *mcast\_dns* service in Linux, where peers advertise their naming information to be able to address nodes by hostname

<sup>1</sup> <https://github.com/dantard/unizar-pound-ros-pkg>

instead of layer 2 MAC or layer 3 IP address. Nodes identified through discovery are added to the list of available nodes for communication as well as status reporting. This discovery heartbeat is also used as a lost-communications detector to prevent higher-level messages from queueing for unreachable nodes.

### 8.2.2. ROS Message Encapsulation

In the ROS ecosystem, messages are translated from a message definition language specification into internal representations appropriate to the implementing language<sup>2</sup>. This same language specification is used to serialize and deserialize messages; that is, to transform a ROS message into a buffer of bytes suitable for transmission over an arbitrary channel. *udp\_mesh* implements a generic message passing system such that the message to be transmitted is never deserialized, saving a significant amount of processing time in the case of complex, large message types such as images. Instead, a generic subscriber is used to acquire the serialized bytes for direct use to be transmitted to other nodes. On the receiver side, the transmitted byte stream is deserialized to instantiate the message in a format that other ROS ecosystem nodes can readily consume. These two functions abstract the transport of arbitrary messages over the *udp\_mesh* layer and remove any requirement to define a list of acceptable message types.

### 8.2.3. Point to Point Transport

In the *udp\_mesh* system, point-to-point transport is implemented via UDP datagrams. This envelope contains provisions for sequence tracking, fragmentation, and message reconstruction. When preparing a message for transmission, the byte buffer provided by the ROS encapsulation service is split into chunks that fit inside the underlying medium’s maximum transmit unit (MTU). We use the standard 802.11 framing with an MTU size of 1500 bytes, out of which 100 bytes are reserved for overhead, leaving 1400 bytes for payload.

In the implementation of our system, a configurable number of message fragments are permitted to be “in flight” at any given time, similar to TCP congestion window control. In order for the next fragment to be transmitted, the receiver must send an acknowledgment. During unit testing to determine an appropriate value for the number of in-flight fragments permitted, an initial increase yields improved throughput. However, after a certain point, throughput decreases as multiple packets are queued for transmission on the medium and start to destructively interfere. As a compromise determined via empirical testing, three packets are permitted to be in-flight between any two nodes at a time. With this configuration, our transport-layer throughput is approximately 20 Mbit/s of payload data, measured using raw images as representative high-density traffic over a wired gigabit Ethernet link.

Retransmits are automatically queued until either an acknowledgment is received or the host is marked offline due to nonreception of any heartbeat or acknowledgment messages. Once a host is marked offline, any attempts to send messages are discarded. Hosts may become online once again after receipt of a discovery message. On the receiver side, the message is kept in a temporary state while the fragments arrive. Should message fragments stop arriving, the partial message is purged after a timeout and the host is once again marked offline which indicates to higher levels that reliable transport is unavailable. In this case, the higher level is BOBCAT, which is discussed in Section 9.

### 8.2.4. Quality of Service

*Quality of Service* (QoS) is the notion that some traffic should be prioritized over other traffic for use of a limited communications channel, e.g., artifact reports need to arrive before mapping data. Fundamentally, TCPROS (the default transport used in ROS v1) is not capable of implementing a QoS scheme where a limited channel is shared between different topics (Figure 13c), as every node subscribing to a topic uses an individual TCP point-to-point link with no information about

<sup>2</sup> <http://wiki.ros.org/msg>



**Table 2.** Evolution of Team MARBLE’s communication system. As a note, ath9k and meshmerize are commercially available, B.A.T.M.A.N. and fkie\_multimaster are open-source software, and udp\_mesh, marble\_multi\_agent, and BOBCAT are custom packages developed for the SubT Challenge.

Event	Physical	Data Link	Transport	Application
Tunnel	ath9k	B.A.T.M.A.N.	fkie_multimaster	marble_multi_agent
Urban	ath9k	meshmerize	fkie_multimaster	marble_multi_agent
Final	ath9k	meshmerize	udp_mesh	BOBCAT

other links. This need to prioritize traffic was the driving rationale behind the development of the *udp\_mesh* layer. As part of the configuration of the layer, each topic to be transported includes a priority number. Internally, this priority number is used as a sorting key to order message fragments for transmission.

### 8.2.5. Point-to-Multipoint Transport

Although *udp\_mesh* is based around point-to-point message transfer, mission requirements sometimes necessitate system-wide messaging. For example, broadcast methods are used within the *udp\_mesh* layer to manage name resolution. To facilitate these type of messages originated at higher levels, a broadcast mechanism is provided by the transport layer. For messages that fit within a single MTU, a single, unacknowledged UDP broadcast is used to distribute the message. For larger messages, individual links to each node are used to send the broadcast as a series of unicast fragments using the same accounting and acknowledgments as the point-to-point mechanism.

## 8.3. Final Solution

The final communication solution used meshmerize as the logical layer with *udp\_mesh* as the transport layer. Both robots and beacons acted as nodes in the mesh with robots carrying 1W radios and beacons carrying 2W radios. Beacon drops are controlled by the methodology described in Section 9.3. Table 2 shows the evolution of our final networking solution from the Tunnel Circuit Event through the Final Event. Our meshing solution, including the *meshmerize* layer 2 software stack, was implemented on *ath9k*-compatible 802.11 hardware, while *udp\_mesh* was implemented on high-level compute units. Because of this split and radio hardware commonality, all of our radios ran essentially the same firmware image built off of an OpenWRT<sup>3</sup> base. Our beacons only participated in the mesh at layer 2, and as such did not contribute to any broadcast traffic associated with *udp\_mesh* services. By providing a reliable ROS-compatible mesh networking layer, higher-level autonomy and human interface via BOBCAT could be provisioned without knowledge of the underlying infrastructure.

## 9. Mission Management

While the combination of Team MARBLE’s large scale positioning system, mapping, and planning solutions provide a solid foundation for autonomy, higher level cognition and reasoning is required to take full advantage of the system. For Team MARBLE, this higher level reasoning consists of a flexible mission management solution which keeps the robots on task and allows for higher level instructions from a human supervisor. The core of the mission management solution is Behaviors, Objectives and Binary States for Coordinated Autonomous Tasks (BOBCAT) (Riley and Frew, 2021). BOBCAT controls the decision-making process for each individual agent while a separate process known as Multiagent Data Collaboration for Autonomous Teams (MADCAT) controls the data sharing and waypoint deconfliction between robots. In this section we highlight the design decisions, and algorithm details behind BOBCAT and MADCAT.

<sup>3</sup> <http://www.openwrt.org>

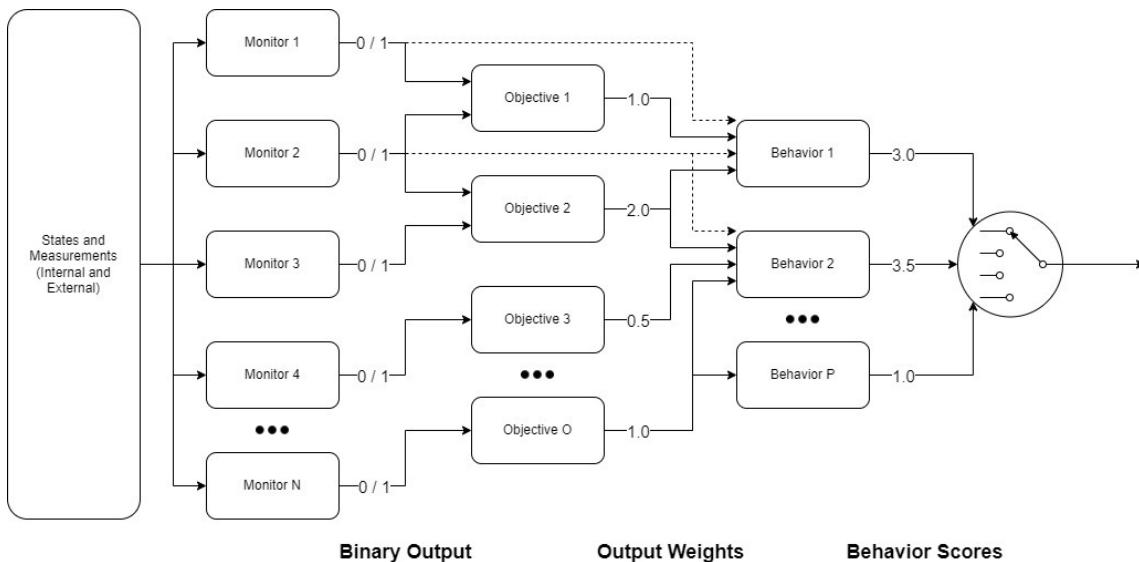
## 9.1. BOBCAT

BOBCAT simplifies the robot and environment states using *Monitors* such as communication status. The *Monitors* are combined with weighted goals which are called *Objectives* such as finding artifacts or extending communications. BOBCAT then selects the best *Behavior* such as exploring or deploying a beacon to fulfill and the most important *Objectives* to execute. A full list of implemented *Monitors*, *Objectives*, and *Behaviors* can be seen in Section 13.7 of the Appendix.

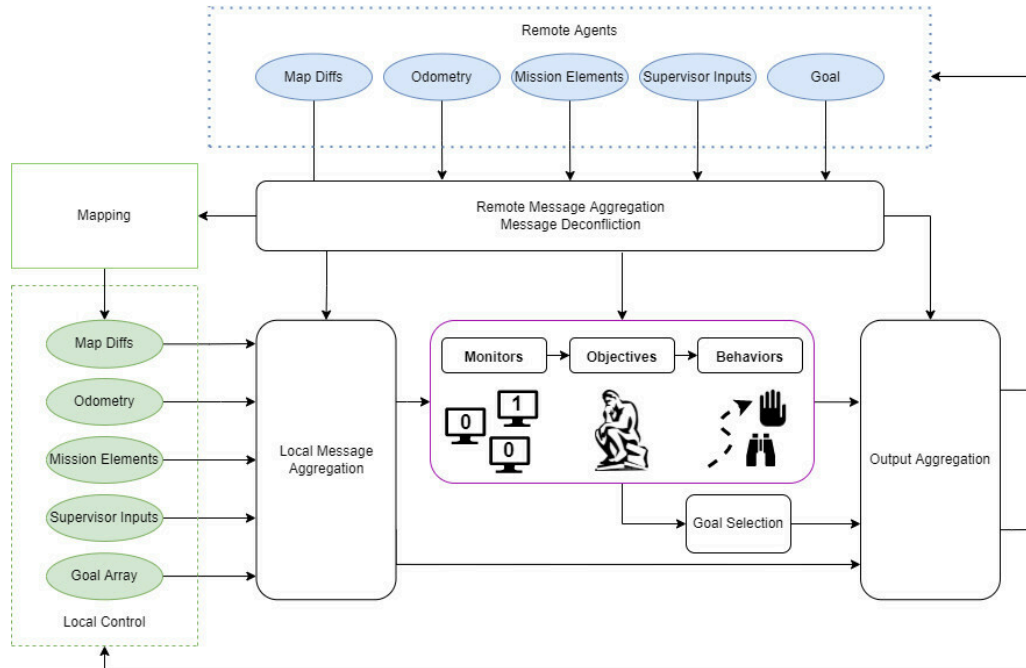
Formally, a BOBCAT is defined by the tuple  $\{\mathbf{x}, \mathbf{y}, \mathbf{w}, M, O, B, \pi_B\}$  where

- $\mathbf{x} \in X$  is the system state with state space  $X$ .
- $\mathbf{y} \in Y$  are the sensor measurements with measurement space  $Y$ .
- $\mathbf{w} \in W = \mathbb{R}_+^{|O|}$  is a vector of *input weights*. These weights are used by the respective *Objective* functions and represent the relative importance of the *Objective* to the overall mission
- $M$  is the set of *Monitor* functions of the form  $M_i : X \times Y \rightarrow \{0, 1\} \forall M_i \in M$ . *Monitor* functions  $M_i$  return a binary value based on the robot state and measurements.
- $O$  is the set of *Objective* functions of the form  $O_j : W \times \{0, 1\}^{|M|} \rightarrow \{0, W_j\} \forall O_j \in O$ . *Objective* functions  $O_j$  use the input weight  $W_j$  and a logical combination of *Monitor* outputs to return either the input weight or a 0, which indicates the current preference of the objective to be fulfilled.
- $B$  is the set of *Behavior* functions of the form  $B_k : \{0, 1\}^{|M|} \times \{0, W_j\}^{|O|} \rightarrow \mathbb{R}_{\geq 0} \times F_k \forall B_k \in B$ . *Behavior* functions  $B_k$  sum the outputs of the *Objectives* associated to that *Behavior*. *Monitor* outputs may be used to selectively inhibit specific *Objective* weights during evaluation steps. The *Behavior* function returns a real value that indicates the current utility score of the actions associated with that *Behavior*, and a pointer to an execution function. A *Behavior* may have a null execution function.
- $\pi_B$  is the policy for selecting the *Execution Behavior*  $B_E$  based on each of the *Behavior* utility scores.

*BOBCAT* can be represented graphically as in Figure 14. States and measurements from both the robot itself and external agents in a multiagent scenario feed the various *Monitors*. This represents what the robot “knows”, and provides a binary output to the rest of the system. The *Monitor*



**Figure 14.** Graphical overview of *BOBCAT*. Numbers represent binary outputs, output weights, and behavior scores, respectively. A full list of monitors, objectives and behaviors is provided in the Appendix in Tables 8, 9, and 10, respectively.



**Figure 15.** Overview of the Multiagent Data Collaboration for Autonomous Teams (*MADCAT*) framework.

output lines in Figure 14 and other figures represent the cases where the *Monitor* is associated with the respective *Objective* or *Behavior*.

## 9.2. MADCAT

The MADCAT framework depicted in Figure 15 provides the multiagent data sharing capabilities required for the mission. The framework includes transmission of relevant coordination data and maps, as well as map merging functionality and decision making for each agent. *MADCAT* uses *BOBCAT* to accomplish the high-level mission management for individual agents with additional higher-level direction provided by the human supervisor.

### 9.2.1. Messages

*MADCAT* sends most messages by broadcast with no acknowledgement required, and therefore does not require the sender to needlessly wait. This allows any agent who receives the message to act accordingly without a requirement to respond. This is helpful in the event the sender leaves communications range shortly after the broadcast. An exception to this policy is made for high bandwidth data such as maps and images, because the receiver can not act on incomplete data. Bandwidth is not strictly managed, but instead uses a “best-available” strategy consistent with the prioritizations assigned to differing message classes, e.g., telemetry, supervisor commands, maps, and FPV.

An agent’s pertinent local messages are concatenated into a single message in order to limit the number of messages broadcast over the communications channels. Messages are re-built and broadcast every second. Only the most recent message is needed for time-varying data such as odometry or the current goalpoint and any older messages are discarded. Other data such as artifact reports and relay locations are appended to a growing list, so any message a remote agent received has all of this type of data. Larger data that could grow to become impractical to transmit repeatedly, such as maps and images, use a point-to-point handshake transmission. Messages are deconflicted using sequence numbers, to allow agents to share the messages of other agents but ensure only the latest data is used by the receiver.

Each agent's broadcast message contains not only its own local data, but that of any neighbor agents as well. This allows downstream agents who can communicate with agent A but not agent B to still receive relatively current information from agent B.

### 9.2.2. Artifact Report Management

Agents keep track of both their own detected artifacts using the procedure described in Section 5 as well any artifacts they have received from other agents. The framework aggregates all of the artifact reports and the *Artifact* monitor determines if the agent needs to return to communications to report the new information to the base station. The base station further parses these messages for display, selection, and transmission to the scoring system. More details of this display can be found in Section 9.5. Images are sent using a point-to-point request system over a low priority channel to reduce bandwidth requirements.

The *BOBCAT Artifact* monitor which is triggered by 3 unreported artifacts or 5 minutes of exploration with a pending artifact, ultimately determines whether the robot should return to communications for unreported artifacts. The raw artifact reports are always used in this determination, but transmission of images is configurable. By default, and as configured during the Final Event Prize Run, artifact images not received by the base station are considered unreported artifacts, and will force the robot to return to communications until they are fully transmitted.

## 9.3. Beacon Deployment

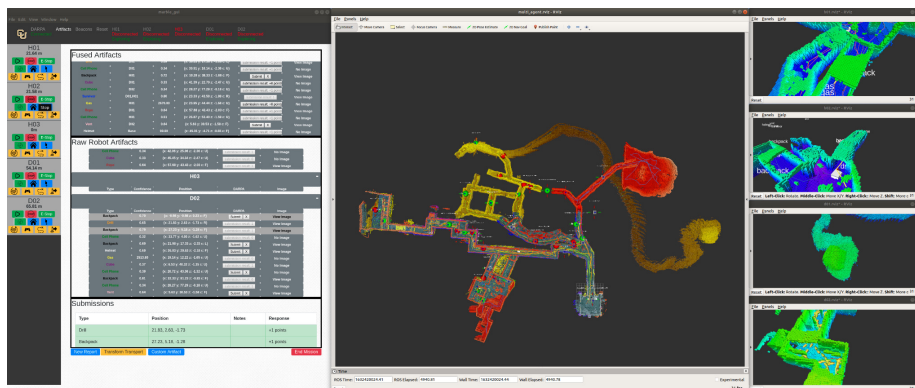
The framework is responsible for identifying locations to deploy communications relays to extend the communications reach into the environment. It uses a combination of communications status, distance and turn detection to identify potential locations. The human supervisor is also able to command drops based on a robot's location on the map.

## 9.4. Goal Selection

Some behaviors, particularly *Explore*, require a goal selection step once that behavior has been chosen to execute. The goals either come from the global planner described in Section 7 or from human supervisor input. If two agents goals are found to be conflicting, BOBCAT requests a new goal from the planner which provides a path to the next closet goal point.

## 9.5. Human Supervisor Interface

The human supervisor interacts with BOBCAT using a custom GUI shown in Figure 16. This interface allows the human supervisor to set a goal point for the robot using an Interactive Marker.



**Figure 16.** Example of the human supervisor interface showing the end of Team MARBLE's Prize Run at the Final Event.

MADCAT then passes this goal to the robot through the communications network if a connection is available. The human supervisor can also remotely control robots using an Xbox controller when communication systems allow.

First-person view (FPV) allows the human supervisor to see the environment from the robot's perspective in semireal time, instead of just through the map representation and infrequent artifact images. In addition to increased situational awareness it allows the human supervisor to identify visual artifacts that may be missed by the on-board artifact detection system, or identify them more rapidly. Finally, FPV helps the human supervisor during manual teleoperation in the event it is needed to direct the robot.

Flexibility of our *udp\_mesh* communication architecture made it possible to rapidly adapt to mission specific constraints. During the Final Event, we observed extra bandwidth in the communication system and decided to add FPV to our Spot robots to enhance their exploration potential. Compressed images from the Spot forward facing cameras were transmitted as 1 Hz low-priority messages. On the base station computer, these images were displayed live in RViz, and saved locally so the human supervisor can review that at any time. The additional scoring potential of FPV is highlighted in Section 10.4.

Several features of the human supervisor GUI were designed to help reduce operator workload. First, another artifact fusion process runs on the Base Station computer, to aid the operator in tracking artifacts reported by multiple vehicles. If reports of the same type are within 3 m of prior reports, they are fused to the mean position. Redundant artifact reports that have already been seen by another robot appear in a light gray color. In contrast, new reports flash with large white text to bring attention to the human supervisor. When submitting artifact reports, the human supervisor can select from individual or fused reports. If an artifact is successfully scored, the submission is locked out to prevent resubmitting. If it does not result in a score, the operator can utilize additional map, trajectory, and FPV information to improve the estimated position. New reports can be submitted by shifting fused artifacts icons in the map or specifying a manual position.

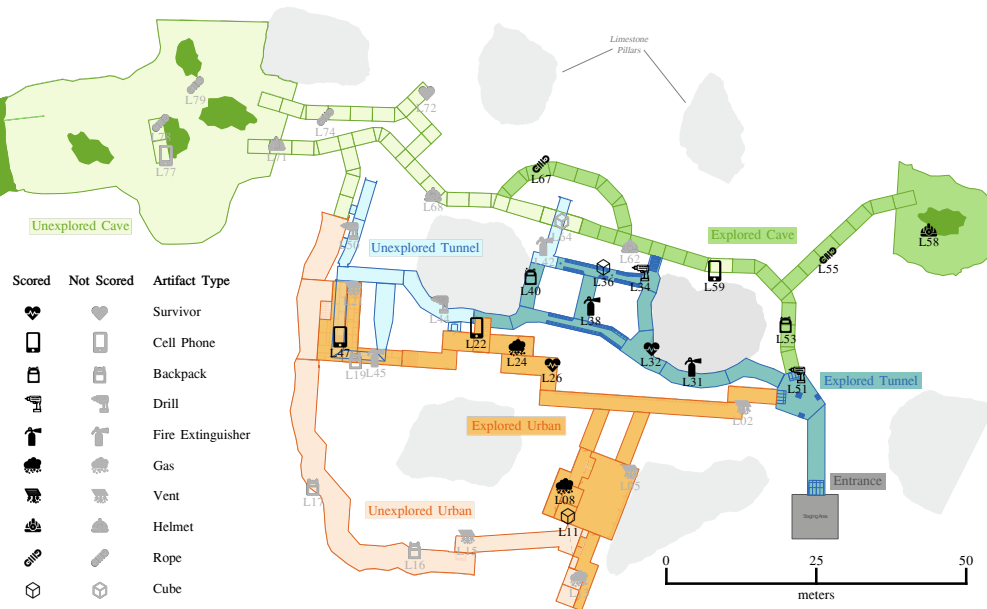
## 10. Final Event Results

The SubT Final Event Prize Run on September 23, 2021 provided Team MARBLE an excellent opportunity to evaluate the performance of our complete supervised autonomy solution, and we share our results in this section. First, an overview is provided in Section 10.1, which includes an outline of the mission objectives, a description of the previously unknown course, as well as a high-level summary of the results. Localization and mapping results are presented in Section 10.2. Further analysis of the planner and resulting exploration effort is described in Section 10.3. Artifact detection results are thoroughly analyzed in Section 10.4. The communication environment was friendlier than expected, and in Section 10.5, we discuss how we capitalized on that opportunity. Mission management results are detailed in Section 10.6, including the five instances where the human supervisor manually intervened, as well as the seven artifacts that were scored via FPV imagery. Together, this section elucidates how our systems worked together to score 18 artifacts, while also discussing the areas that limited even higher performance. Data from Team MARBLE's deployment during the Final Event Prize Run is publicly available and discussed in Section 10.7.

### 10.1. Overview

The mission objectives were to accurately report as many of the 40 artifacts in the course as possible during the mission. There are three hard constraints: the mission is 60 minutes long, there are a total of 40 attempts to report artifacts, and only one human, the human supervisor, is permitted to supervise the mission, manipulate robotic agents, and submit artifact reports.

The final course was custom constructed as illustrated in Figure 17, which breaks the course out into distinct tunnel, urban, and cave environments. The course contains numerous hazards



**Figure 17.** Course map of the 60-minute Final Event Prize Run designed by DARPA, highlighting the 18 of 40 artifacts scored by Team MARBLE, along with the areas of the tunnel, urban, and cave sections were explored by agents.

and challenges: rough terrain, railroad tracks, slippery surfaces, ramps, stairs, large drop-offs, rocky cliffs, narrow hallways, low-to-the-ground corridors, wide-open caverns, fog, standing water, dynamic obstacles, trap doors, and a degraded communications environment. These challenges are discussed further in Section 10.3.

Here, we provide a brief high-level summary of the artifacts scored and the extent of the environment explored. Team MARBLE scored 18 of the 40 artifacts and explored roughly half of the environment. For reference, the top-scoring team scored 23 artifacts, and the performance for all teams is listed in Section 13.9 of the Appendix. The location and class of all 40 artifacts can be visualized in the context of the course map shown in Figure 17. This map also indicates which regions of the course that Team MARBLE explored as well as the 18 scored artifacts. The artifacts that Team MARBLE scored are also listed in Table 3, ordered chronologically from mission start to mission end. Each of the 18 artifacts in Table 3 correspond by ID to the scored artifacts in Figure 17. Further analysis of artifact detection results are detailed in Section 10.4.

## 10.2. Localization and Mapping


A secondary objective of the 60-minute Final Event Prize Run was to rapidly map the environment and transmit the real-time map back to DARPA every ten seconds. The map takes the form of a point cloud, or a collection of three-dimensional points that represent occupied space in the environment, e.g., floors, walls, ceilings. Figure 18 provides a comparison of the DARPA-generated ground truth map in black against the final map submitted by Team MARBLE, split into inliers in green and outliers in red.

DARPA has generously collected, processed, and shared this map data with participating teams (Schang et al., 2021). Inliers are defined as points within 1m of ground truth map points, and outliers are defined as points outside 1m. Map coverage is a metric representing the ratio of the environment explored, and is defined as

$$\text{map coverage} = \frac{\text{ground truth points within 1m of an inlier point}}{\text{total ground truth points}}. \quad (9)$$

**Table 3.** List of all artifacts scored by Team MARBLE during the 60-minute Final Event Prize Run, along with the corresponding mission time when artifacts were reported and scored, artifact type, unique DARPA-assigned ID numbers, and Euclidean distance error between the reported and ground truth location of the artifact.

Score	Time [mm:ss]	Type	ID	Error [m]
1	01:08	Drill	L51	0.57
2	01:23	Backpack	L53	2.23
3	06:23	Rope	L55	0.84
4	12:03	Survivor	L26	0.62
5	16:35	Survivor	L32	1.40
6	17:23	Gas	L08	1.80
7	28:08	Fire Extinguisher	L31	1.31
8	35:51	Drill	L34	1.43
9	36:53	Fire Extinguisher	L38	2.82
10	37:08	Cube	L36	3.94
11	37:58	Backpack	L40	1.40
12	38:47	Rope	L67	2.87
13	47:53	Cube	L11	1.83
14	50:33	Cell Phone	L22	4.06
15	50:45	Cell Phone	L47	4.00
16	51:48	Cell Phone	L59	2.15
17	52:45	Gas	L24	2.55
18	56:33	Helmet	L58	1.74




**Figure 18.** Final point cloud map submitted by Team MARBLE to DARPA staff during the Final Event Prize Run.


Map error or deviation is a metric representing the ratio of the submitted map that is inaccurate relative to the ground truth map, and is defined as

$$\text{map deviation} = \frac{\text{outlier points}}{\text{total submitted points}}. \quad (10)$$

Figure 19 shows that map coverage steadily increases throughout the mission, with some periods of rapid exploration, and by the end of the mission, nearly 50% of the environment has been mapped. Map error on the other hand, increases modestly throughout the mission due to localization drift.



**Figure 19.** Map coverage, error, and cumulative score throughout Team MARBLE’s deployment during the DARPA SubT Final Event Prize Run.



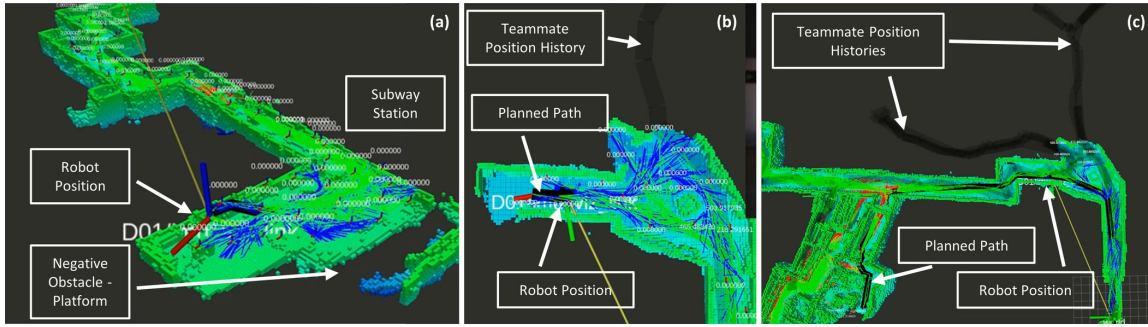
**Figure 20.** The final maps built onboard (a) D01, (b) H01, (c) D02, and (d) H02 at the DARPA SubT Final Event Prize Run.

However, it increases significantly due to a localization failure on D01, which is discussed further in Section 10.3.4. This failure generated erroneous sections of map, which appear as a long winding corridor in Figure 18 and Figure 20.


### 10.3. Planning

During the entire 60-minute mission and across a diversity of environments, none of the robots were teleoperated due to a planner failure or to improve the volumetric gain, which is seen as





**Figure 21.** The figure shows three snapshots from the SubT Final Event Prize Round run. The blue lines in the figure represents the locally sampled RRT\* tree and the green lines represent the global graph. (a) The robot can be seen to avoid sampling over the negative obstacles i.e., the edge of the subway platform, owing to the settling-based collision checks. (b) The agent planned away from the position history of a teammate robot that was launched before it, demonstrating effective multiagent coordination. (c) This snapshot shows an instance where teleoperation was initialized on one of the robots. The robot can be seen to plan a path that had sufficient volumetric gain leading the robot toward a frontier in the urban area that had not been seen by any other robots. However, the human supervisor decided to teleoperate the robot through the initial section of the tunnel area and then let it autonomously explore the tunnel.



**Figure 22.** Volumetric gain explored by each robot across the 60-minute Final Event Prize Run. Planner-Off, denoted by blue lines, represents instances when the planner was paused to allow the autonomous mission-management system to take over when robots are in close proximity, as well as the several interventions when the human supervisor manually teleoperated robots. Go-Home, denoted by black lines, represents instances when the planner began returning home to reconnect to the network and report new artifact reports and map data.

a successful demonstration of the planning reliability and flexibility. The human supervisor only intervened to complement autonomy with human-level cognition and intelligence, as discussed further in Section 10.6.

As an overview, Figure 21 presents three scenarios from the Final Event Prize Run, which demonstrate negative obstacle avoidance, multiagent coordination based on teammate position histories, and teleoperation initiation by the human supervisor. Below, Section 10.3.1 highlights the exploration performance, Section 10.3.2 shares examples of agents avoiding treacherous terrain, and Section 10.3.3 demonstrates the planner adapting to dynamic changes in the environment. Section 10.3.4 discusses system limitations that prevented the agents from exploring the entire course.

### 10.3.1. Exploration

The volume of unseen area explored by each agent across the mission is illustrated in Figure 22. Unlike the statistics of the global map submitted to DARPA in Figure 18, Figure 22 presents agent-specific exploration information stored onboard each robot. D02 was the largest contributor to overall exploration, partly because it was launched first and had the most time to explore.


D01 played a complementary role by exploring most of the tunnel environment, which remained unexplored by D02. Because the subset of the environment safely traversable by the wheel robots was relatively more constrained, it naturally led to less exploration from the Huskies across the mission.

Besides showing the volume explored by each of the robots, Figure 22 also shows time periods during which the default planner exploration behavior was paused for higher-level mission management and teleoperation commands. These interjections by the autonomous mission management system were primarily triggered when agents approached each other, and resulted in lower-priority agents pausing and higher-priority agents resuming their task. During these encounters, agents appear as fast-approaching dynamic objects, and this simple procedure was employed rather than incorporating reactive obstacle avoidance into the planner.


### 10.3.2. Treacherous Terrain

The path planning solution successfully kept each agent safe from collision throughout the entire mission. The Spot robots, which explored more challenging features in the environment, fully avoided negative obstacles, such as shear drop-offs and rocky slopes, and traversed up and down stairs.


The Spot robots successfully traversed up and down the small set of stairs leading up to the subway platform, as shown in Figure 23. However, when stairs were first encountered from above, agents did not plan down them due to the limited  $\pm 16.6^\circ$  vertical field of view of onboard Ouster OS-1-64 lidar sensors, as shown in Figure 24a. In addition, the Spots can only safely walk down stairs backwards, and therefore additional logic would be required to autonomously traverse those stairs.




**Figure 23.** Instance of (a) D01 planning up stairs with the green edges of the graph, pink planned path, and associated semantic map with blue voxels representing stairs, along with (b) FPV imagery.



**Figure 24.** Instance of (a) D02 thoroughly exploring the subway platform without planning over the edge, with green edges of the graph, pink planned path, and associated map, along with (b) FPV imagery.



**Figure 25.** Instance of (a) D02 autonomously exploring the entire cavern, traversing the safer surfaces and avoiding treacherous areas, such as the one shown in (b).



**Figure 26.** Early in the mission, D01 walked by the corridor with the trap door, as shown by (b) the left camera (1:47). The agent later returns to the corridor, (a) walks under the trap door (10:34), and soon after sees it has closed, as shown by (d) the right camera (11:01). After seeing the trap door close, (c) the updated map and graph show (red) edges as untraversable (11:11). Had the agent moved closer to the trap door, and fully been within mapping range, all edges would have updated as untraversable.

Both Spot robots thoroughly explored the subway platform, approaching the edge, but never stepping and falling over, as demonstrated by D02 in Figure 24. Additionally, the same Spot robot, D02, explored the entire cavern autonomously without entering treacherous terrain, as shown in Figure 25.

### 10.3.3. Dynamic Environment

The planner has the ability to adapt to dynamic environments, such as closing or opening of doors, falling rubble, as well as other situations also lead to dynamic changes in the map, including other nearby mobile agents, and localization and mapping error.

During the Final Event Prize Run, D01 traveled through a side branch of the urban environment, triggering a trap door. Figure 26 shows the planner adapting to the dynamic environment by re-assigning previously traversable edges as untraversable.

In addition, there were several instances of temporary localization and mapping error, which caused erroneous new map data to change previously traversable edges of the planning graph to untraversable. In each case, the planner adapted to the new scenario, and continuously operated throughout the localization drift and loop closure correction. An example of this is included in Section 13.10 of the Appendix.

When agents pass each other, they appear as fast-moving dynamic obstacles, and cannot re-plan around one another fast enough. Therefore, an agent-based prioritization scheme prevents both collision and deadlock, by enforcing one agent to wait while the other passes. Examples can be found in Section 13.10 of the Appendix.

### 10.3.4. Limitations

However, several limitations did prevent agents from exploring roughly half of the course. In total, there were three types of bottlenecks: constrained corridors, slippery surfaces, and downward sets of

stairs, each exposing unique limitations within the autonomy system. None of these are limitations of the planner itself, but rather limitations of mapping, mobility, and perception.

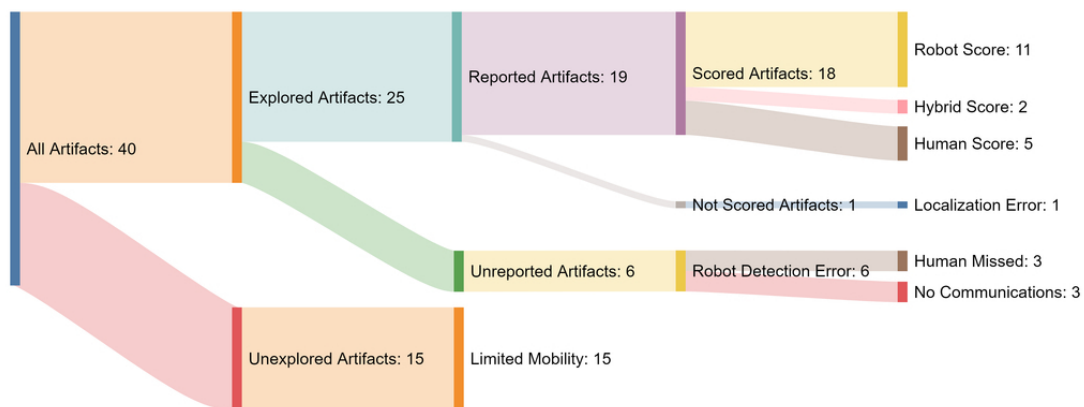
The planner did not plan through all constrained spaces because the selected planning parameters for agent width and height did not allow the graph to propagate through especially short and narrow spaces. These parameters were intentionally chosen to be conservative to prevent the agent from moving along an unsafe trajectory. Utilizing a higher-resolution local map and planner could result in a more agile robot that could safely traverse those spaces. Additionally, autonomously transitioning into a crouching gait could improve the Spots ability to traverse spaces with low ceilings. Examples are included in Section 13.11 of the Appendix.

The second limitation is slippery surfaces, and led to D01 slipping and falling in the cave section. Some of these rocky surfaces were intentionally designed to be slippery, and plenty of humans walking through the course after the event also slipped and fell. After D01 fell, it also experienced a localization failure. Methods to recovery the system from an event such as this one would involve implementing a fall detection algorithm, as well as autonomous self-righting and localization reset logic.

#### 10.4. Artifact Detection

In this section, we present performance results of the artifact detection and reporting system. During the 60-minute Final Event Prize Run, Team MARBLE scored a total of 18 artifacts out of the 40 artifacts in the environment. Figure 27 presents a flow diagram that summarizes how our team scored 18 artifacts and the limitations that resulted in the remaining 22 from being scored. Of all 40 artifacts, our agents explored enough of the environment that they were in the vicinity of 25 artifacts, leaving 15 unexplored due to mobility challenges discussed in Section 10.3. Team MARBLE reported 19 of the 25 artifacts that were explored, and successfully scored 18 of those 19 reported. A map of the area explored and scored artifacts is shown in Figure 28. Details of these 25 explored artifacts, of which 18 were scored, one was missed, and six were unreported, are shared in Section 13.13 of the Appendix.

Of the 18 artifacts that Team MARBLE successfully scored, 11 were scored by autonomous robot reports, two were scored by the human supervisor modifying the position of autonomous reports, and five were scored by the human supervisor via robot FPV imagery. One artifact was reported but did not score due to localization error in excess of 5m. There were six artifacts that agents saw, but did not report due to errors in the autonomous artifact detection system. The human supervisor received information regarding three of these six artifacts but missed them due to high workload demands during the mission. The other three artifacts were located in areas that prevented agents from communicating back to the human supervisor.



**Figure 27.** Flow diagram illustrating how Team MARBLE scored 18 of the 40 artifacts in the course, and the limitations preventing the remaining 22 artifacts from being scored.

**Table 4.** Artifact statistics for Team MARBLE during the Final Event Prize Run. A total of 18 artifacts were *scored*. Only one artifact was reported but *not scored*, which was due to localization error greater than 5m. Agents were in the vicinity of six artifacts, but they went *unreported* due to autonomous artifact detection failure and in some cases, also missed by the human supervisor due to excess workload. The remaining 15 *unexplored* artifacts were never seen by agents because they were located in parts of the course that were never reached.

Artifact Type	Scored	Not Scored	Unreported	Unexplored	Total
Survivor	2	0	0	1	3
Cell Phone	3	0	0	1	4
Backpack	2	0	0	3	5
Drill	2	0	0	2	4
Fire Extinguisher	2	0	1	1	4
Gas	2	0	1	0	3
Vent	0	0	3	1	4
Helmet	1	0	1	3	5
Rope	2	0	0	3	5
Cube	2	1	0	0	3
<b>Total</b>	<b>18</b>	<b>1</b>	<b>6</b>	<b>15</b>	<b>40</b>



**Figure 28.** Locations of ground truth artifacts in white, reports that scored in green, and reports that did not score in red, overlaid on ground truth map of the course.

#### 10.4.1. Visual Detection

The focus of this section is to quantify the performance of the visual artifact detection system. Of the 18 artifacts that Team MARBLE scored, 11 of them were visual artifacts, as shown by Table 4. First, we focus on the six artifacts (L51, L53, L26, L34, L40, L67) that were successfully reported by the agents' autonomous artifact detection systems in the course. All six artifacts were accurately localized to within 5m. In fact, the largest error for a visual artifact was 2.87m (L67). This eliminated the need for the human supervisor to spend time trying to correctly localize artifacts.

**Table 5.** Artifact report statistics for Team MARBLE during the Final Event Prize Run. A total of 34 reports, or attempts, were made throughout the mission. A total of 18 attempts resulted in *scores*, 11 attempts were *misses* and did not result in a score due to localization error greater than 5 m, and five attempts were *false* attempts in that they were false positives and no artifact of that class was in the vicinity.

Attempt Type	Scored	Missed	False	Total
Survivor	2	0	0	2
Cell Phone	3	5	0	8
Backpack	2	0	1	3
Drill	2	0	0	2
Fire Extinguisher	2	0	1	3
Gas	2	1	3	6
Vent	0	0	0	0
Helmet	1	0	0	1
Rope	2	0	0	2
Cube	2	5	0	7
<b>Total</b>	<b>18</b>	<b>11</b>	<b>5</b>	<b>34</b>

The autonomous artifact detection system filters raw frame-to-frame detections onboard the agent, with the aim to reduce the number of false positive and redundant artifact reports. Because the human supervisor has limited bandwidth, unnecessary distractions detract from the completing other mission-related tasks. In the process of scoring these six visual artifacts, the human supervisor had to process 21 artifact reports from the automated artifact detection systems onboard agents. Of these, 11 were true positives, six were approved by the human supervisor and successfully reported, and five were ignored because they were redundant reports that were previously scored. The other 10 reports were also ignored by the human supervisor because they were false positives.

In total, there were only four false reports that the human supervisor submitted. One was caused by human error, the other three were caused by erroneous CO<sub>2</sub> reports, as detailed in Section 13.16 of the Appendix.

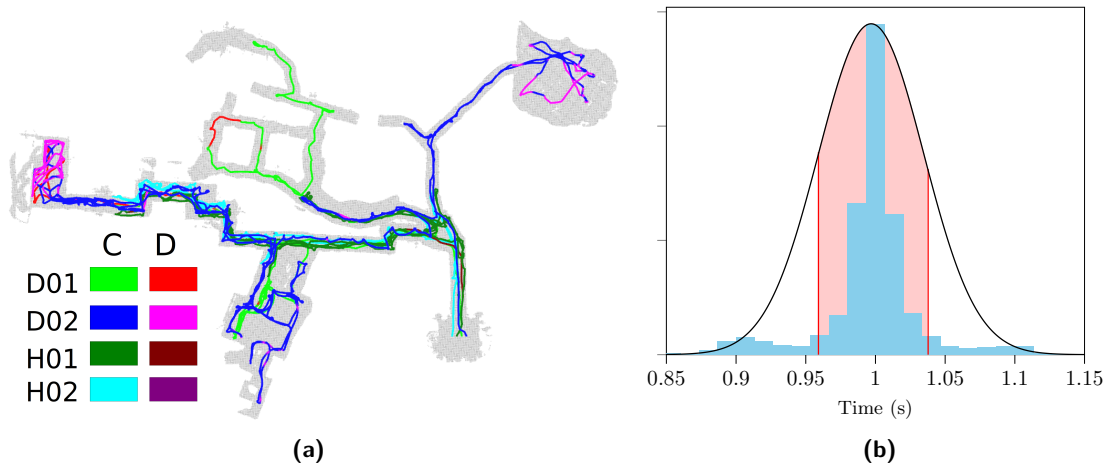
Agents in the course failed to autonomously report the other five artifacts (L55, L32, L31, L38, L58), but did transmit FPV imagery back to the human supervisor, who manually reported and scored them. The fact that the autonomous artifact detection system did not detect five of the 11 visual artifacts it saw, indicates that certain reliability limitations exist. Team MARBLE acknowledged this limitation and relied on the human supervisor and FPV system to fill in that void, which is further in Section 10.6.

#### 10.4.2. Nonvisual Detection

Team MARBLE scored seven nonvisual artifacts, i.e., cell phone, cube, and gas, but as shown in Table 5, required submitting more reports due to difficulty around accurately localizing the source. The main limitation is that the detection scheme relies on threshold-based logic for RF and CO<sub>2</sub> levels, and when triggered, simply reports the current location of the agent. The thresholds were intentionally set low to increase the probability of detection when agents pass by the vicinity of these nonvisual artifacts. Of the seven nonvisual artifacts scored, five of them (L08, L36, L22, L47, L24) were scored via the autonomous robot reports, with an average error of 3.27m. The remaining two artifacts (L11, L59) were scored by the human supervisor manually adjusting the reported artifact locations.

### 10.5. Communications

The performance of the communication system was evaluated in the Final Event Prize Run using both qualitative and quantitative measures. Subjectively, the human operator was able to employ



**Figure 29.** Performance results of the communication systems, including (a) the map of the Team MARBLE's deployment during the Final Event Prize Run, overlaid with locations of robot connection (C) to the network in blue-green, and locations of robot disconnection (D) from the network in red-magenta, as well as (b) the distribution of intermessage arrival times for D01 with a nominal publishing rate of 1 Hz, overlaid with  $\mathcal{N}(1.00, 0.04)$ . The highlighted red region represents the first  $\sigma$  value which contains 68% of the message times.

live FPV video from the robots, a capability that directly contributed to team's third place finish. The robots were in communication with the base station over the majority of the explored regions of the course, as shown by the blue-green hues in Figure 29a, allowing the human supervisor to monitor and intervene as needed. Overall, 125.2 MB of data was transferred through the communications network, including all map segments, telemetry, artifact reports, and other data products. Of that, FPV video comprised 51.6 MB. The latency of the mesh networking solution was evaluated using intermessage arrival times of a heartbeat message sent from a long-ranging robot to the base station. This mission management message, transmitted regularly from the robot, was part of our protocol scheme and is analyzed as a message of convenience. A distribution of the intermessage arrival times of these heartbeat messages from D01 to the base station is shown in Figure 29b.


Since these messages originate from the robot at 1 Hz, an ideal system would observe all intermessage arrival times to be one second in duration. As our mission management system on the robot does not run with hard realtime constraints, the 1 Hz publish rate is an estimate that includes noise due to process load, etc. On the base station, there was an approximately normal ( $\mathcal{N}(1.00, 0.04)$ ) distribution of arrival times. Since messages may experience delays, the immediately following message may exhibit an intermessage time of less than one second, leading to the symmetry apparent in Figure 29b. Our key observation of this plot is that the bulk of messages arrive within five percent of their expected times across a distance of hundreds of meters and multiple mesh hops, validating the performance of the entire mesh networking solution.

## 10.6. Mission Management

Overall, the mission management system was able to keep the robots on task with minimal human supervisor intervention. However, when needed, the interventions were crucial towards both the exploration capabilities of the system and the Final Event performance. Figure 30 presents a detailed timeline of the four agents in the field, as well as the human supervisor. The four robot launches and five robot interventions were the only times when the human supervisor used teleoperation. There were only two other types of instructions agents received from the human supervisor. One was commanding H02 to drop two communication beacons. The other was commanding D01 to return home three times, each occurring while the human supervisor was also teleoperating D01 through the fog.

**Table 6.** List of the four robot launch (RL) sequences executed by the human supervisor during the 60-minute Final Event Prize Run. Extensive process streamlining and repeated practice deployments resulted in quick and repeatable launch sequences. The course entry column represents the mission time at which the agent crossed into the course.

Launch	Agent	Duration [s]	Teleoperation Window [mm:ss – mm:ss]	Course Entry [mm:ss]
RL1	D02	75	–00:37 – 00:38	00:02
RL2	D01	55	01:29 – 02:24	01:44
RL3	H02	35	07:15 – 07:50	07:25
RL4	H01	50	10:57 – 11:32	11:08
	<b>Mean</b>	<b>41</b>		



**Figure 30.** Mission management timeline for all robots and human supervisor during the 60-minute Final Event Prize Run. When not manually teleoperating a robot, the human supervisor was monitoring the mission, which includes watching live FPV streams from D01 and D02, reviewing incoming artifact reports from agents, reporting artifacts to DARPA, and in the last five minutes of the mission, checking archived FPV images from the Spots for previously missed artifacts.

### 10.6.1. Robot Launches

The human supervisor was under immense pressure to optimally balance many competing tasks during the 60-minute Final Event Prize Run. With such limited time, the primary objective at the beginning of the mission is to launch robots into the environment as fast as possible. Through extensive practice and full-scale comprehensive field deployments, discussed further in Section 11.2, Team MARBLE launched all four robots, with a mean launch time of 41 seconds, as shown in Table 6.

Two agents experienced failures late in the mission, reducing overall fleet utilization rate from 92% to 73%. H02 experienced a hardware failure (36:48), and post-event inspection revealed better vibration isolation of the computing system would reduce the likelihood of such a failure in the future. D01 experienced a mobility failure (39:46), in which the Spot slipped on a slick rock and fell over. The agent then experienced a localization instability due to the large induced velocity. To recover from such an incident in the future, Team MARBLE could implement an autonomous self-righting maneuver and localization reset logic.

Despite the fact that some artifacts were not detected and reported by the autonomous onboard artifact detection system, the human supervisor filled in the void. The human supervisor reported and scored five artifacts that were seen via robot FPV imagery, but not autonomously detected. Of



**Table 7.** List of the five robot interventions (RI) executed by the human supervisor during the 60-minute Final Event Prize Run. Our concept of operations relies on autonomous multiagent exploration, and does not necessitate manual waypoints or teleoperation from the human supervisor. Therefore, agents were completely autonomous, except for the human supervisor input during these five instances of teleoperation. The interventions goals varied, but were all specific scenarios where human intervention would augment autonomous agent capabilities in a mission-relevant manner.

Intervention	Agent	Duration [s]	Window [mm:ss]	Goal	Success	Points
RI1	D02	240	03:02 - 07:02	Enter narrow cave corridor	+	6
RI2	D01	598	21:54 - 31:52	Enter foggy tunnel area	+	2
RI3	H02	24	35:04 - 35:28	Avoid course exit	+	0
RI4	D02	200	39:24 - 42:44	Walk down stairs	-	0
RI5	H01	21	49:40 - 50:01	Avoid course exit	+	0
	<b>Mean</b>	<b>217</b>				<b>1.6</b>



**Figure 31.** Imagery during robot intervention 2 (RI2). when human supervisor was pushing D01 through the foggy area in the tunnel environment, eventually leading to six additional points. Shown is a comparison of (a) full-resolution FPV imagery onboard the agent and (b) compressed low-resolution FPV imagery transmitted to the human supervisor.

these, the human supervisor saw four (L55, L32, L31, L38) from live FPV streams, while one (L58) was found while reviewing archived FPV images near the end of the mission. More details of these artifact reports are presented in Section 13.17 of the Appendix.

### 10.6.2. Robot Interventions

In total, the human supervisor intervened in the autonomous fleet five times during the 60-minute mission. Table 7 presents the duration of each intervention, with a mean length of 217 seconds, as well as the reason for intervening. Each of these interventions was in the form of manual teleoperation, in which FPV streams were available at the base station, and velocity commands from the human supervisor was transmitted to the remote agent in the field. Of the five manual teleoperation interventions, only two had significant impact on the mission. The first (RI1) was navigating D02 through the narrow cave corridor early in the mission (3:02 - 7:02), which led the agent to the small cavern and two artifact scores. The second (RI2) was navigating D01 through fog in the tunnel section midway through the mission (21:54 - 31:52), leading to six artifact scores. Figure 31 provides a side-by-side comparison of the full-resolution FPV imagery processed onboard by the autonomous artifact detection system and the compressed FPV imagery transmitted to the human supervisor.

Two interventions (RI3, RI5) commanding agents back into the course, was out of a sheer abundance of caution. The planning algorithm is configured so that the staging area is treated

as explored, so agents should not attempt to explore it. One intervention failed (RI4), in which the human supervisor attempted to teleoperate D02 down the stairs by the subway platform. This attempt failed because communication to and from the robot was intermittent.

The main takeaways from these results is that our agents are highly autonomous, leaving the human supervisor to focus on mission monitoring and targeting strategic, high-value intervention opportunities. The mission management system enabled convenient transition between autonomy and manual operation, while the communication system enabled visibility and control over the agents in the field.

## 10.7. Open-Source Data

During the final run at the Final Event, our team collected a significant amount of data related to autonomous subterranean exploration in the form of ROS “rosvags”. Our datasets are split up by agent, and each set contains a rosvag of the system inputs, mostly consisting of raw sensor data, and another rosvag of the outputs used for visualization and performance monitoring. The complete collection of data recorded at the Final Event can be found at <https://doi.org/10.25810/a8ym-5z52>.

## 11. Lessons Learned

Underground exploration of previously unknown environments, especially in a time and resource-constrained search-and-rescue context, requires a highly adaptable human-robot team. The lessons learned presented in the following sections enhance our proposed system’s flexibility across mobility, communications, human-robot teaming, and multiagent coordination.

### 11.1. Platform Mobility

Systems with heterogeneous platforms allow for specialization by each platform for both specific environments, and roles which benefit the entire mission. Specifically, in Team MARBLE’s case, the addition of Spot platforms into an exploration role enabled rapid multistory expeditions. This capability was further augmented with the utility of higher-payload, wheeled Huskies, carrying communication beacons. When deployed, these beacons allowed the Spot platforms to report artifacts without the need to traverse “home,” leading to much more efficient exploration. No single platform is as performant as the combination of platforms operating in different roles based on each one of their strengths.

### 11.2. Testing and Validation

Significant testing and validation was conducted for every element of the Team MARBLE system. These tests taken over a variety of environments, shown in Figure 32, reduced mission-to-mission variability, and increased system-wide adaptability. Through these tests, Team MARBLE locked in well-tested solutions with minimal unexplained errors. This reduced changes on the system as we approached the Final Event, given that new solutions had to be verified against similarly stringent tests. To our knowledge our team made the fewest hardware and software based adjustments from the two Preliminary Runs to the Prize Run to achieve our score, instead relying on our testing having eliminated all significant errors outside of a so-called “Poisson distributed fatal error.” These types of errors were mission ending to any individual robot but difficult to predict or adapt to in a meaningful sense. These can be seen in our Prize Run where D01 fell over on rough steep terrain, and where H02 suffered a hardware error.

A secondary goal of these full-scale deployments was to reduce human-based variability in performance. They helped both the human supervisor and pit crew prepare for the stressors of operating and responding quickly to complex interactions between robots, the environment, and potential failure modes. At no point was a single test considered sufficient for validating a solution as sufficient;



**Figure 32.** Team MARBLE conducted comprehensive field deployments at various sites including (a) the Edgar Experimental Mine operated by Colorado School of Mines, located in Idaho Springs, CO. which tested terrain, distance travelled, and communications; (b) the Engineering Center complex also located on University of Colorado Boulder Main Campus which tested multistory navigation and repeated features from urban environments; and (c) the Folsom Parking Garage located on University of Colorado Boulder Main Campus in Boulder, CO. which tested planning in open spaces and vertical localization across multiple stories.

instead all solutions released demonstrated repeatability across the same and varied environments. For reference, all full scale deployments are tabulated in Section 13.18 of the Appendix.

### 11.3. System Adaptability

The challenges posed by operating a system in an unknown environment necessitate a high level of system adaptability. Predicting every capability that a system will need for a given mission, such as search and rescue, is impractical. Instead, having a highly flexible system capable of adapting to unknown situations is crucial. Having already addressed the mobility considerations in Section 11.1, we also found software adaptability key to our success. For instance, the flexible communication network enabled the ability to pass on FPV to the human supervisor with a simple configuration change. While our system was designed with autonomy in mind and the capability was not previously planned for, the adaptation proved invaluable. Minimal human supervision directly impacted the final score and exploration capabilities of the system as a whole.

The notion of adaptability extends to other software architecture as well. The artifact detection system was adjusted during the final competition to include SSID information for Bluetooth artifacts. These artifacts could then be more accurately merged between agents and by the human supervisor, despite inaccurate positioning from wireless signals.

In some cases, adaptability is explicitly accounted for in our design, most notably in the mission management system. BOBCAT has many parameters which allowed the human supervisor to adjust exploration activity, including how long a robot can explore before reporting detected artifacts, whether a robot should find multiple artifacts before reporting, and how the system should adjust to time constraints in the mission.

### 11.4. Autonomy and Human Robot Interaction

Team MARBLE emphasized autonomous system design as our primary goal. Since only one person, the human supervisor, was permitted to interact with agents while they were deployed, having robots controlling their own paths and decision making was key to reducing the cognitive load required to manage mission objectives. Despite the focus on autonomy, the human supervisor is still necessary to maximize system performance. As Team MARBLE approached the final competition, we found targeted areas where direct human supervisor control improved results. This influenced decisions regarding the number of deployed platforms, how artifacts were passed from agents to the human supervisor, and how we recovered from anomalous robot behavior.

To maximize the human supervisor's ability to track the data streams present across the fleet, only four robots were deployed. While a larger fleet might have enabled a rapid exploration of

the environment, it had potential to reduce the human supervisor’s ability to meaningfully address problems arising from any specific robot.

The artifact detection system was designed to filter against false positives before passing information to the human supervisor. Artifacts had to consistently be detected in multiple frames, and be a sufficient distance from existing artifact estimates. Where possible, artifacts were returned with additional information including SSID for Bluetooth artifacts and images for visual artifacts. This data allowed the human supervisor to sort remaining false positives quickly without being distracting from the core mission objectives.

Specifically, the human supervisor spent most of their operational load solving problems the robots were incapable of correcting through their autonomy stack. For example, during the Final Event Prize Run, the human supervisor used the first person vision to navigate through dense fog, which the robot was incapable of planning through autonomously. After this intervention, the human supervisor allowed the robot to return to autonomous operations, where it explored several new areas, and reported a total of six new artifacts. The human supervisor was able to opportunistically intervene like this, only because the other agents were operating autonomously without supervision.

### 11.5. Retrospective

It’s important, after a three-year effort of this size and scope, to examine some of the bigger picture questions. What would we do differently next time? What did we wish we knew at the start, that we know now? The answer to the first two questions is that we would have spent more time at the onset to scope out short-term and long-term development goals. We excelled at the Tunnel Circuit Event, placing fourth amongst a large group of competitors. This was mostly due to meeting well-scoped short-term goals within the one year cycle. However, during the six-month development period for the Urban Circuit Event, we focused on long-term goals that we ultimately only partially validated, leading to a disappointing performance. In hindsight, we should have focused on making our already capable platforms more capable, rather than spreading our resources thin across many thrusts. However, what our team excelled at after this experience, was pivoting and adopting two new strategies. First, we leaned into student-led project management, which led to greater team coordination, a key component of rapid and effective system development. Secondly, we focused on long-term development goals, and given a year and a half, had the time to adequately develop, test, and validate each subsystem, each fully autonomous robot, as well as our entire fleet in numerous search and rescue missions. Together, these two powerful changes allowed a small, lean team, produce a highly functional autonomy solution that could perform under pressure.

## 12. Conclusions

In this paper, we showcase our flexible autonomy solution for exploring unknown subterranean environments. The highly performant autonomy solution directly lead to a third place finish at the DARPA SubT Final Event which was focused on search and rescue. Moreover, the specific innovations presented in graph planning, flexible communications, and mission management are directly applicable to other multirobot teaming applications under limited human supervision.

Specifically, the combination of legged and wheeled robots allowed for heterogeneous teaming enabling both rapid exploration and a robust communication network. The deployed mesh network which is described in Section 8 enables flexible configuration and prioritization of data sent both to other robots, and a human supervisor for review. A powerful graph planning framework, as described in Section 7, paired with semantically encoded Octomaps enabled safe, rapid exploration. The final system was able to explore a variety of underground environments, including gold mines and subway stations, with minimal human input.













Human input was reserved for specific situations where higher-level reasoning had the potential to improve the mission. The flexible mission management system described in Section 9 enables safe transitions between human input and the underlying autonomy system. One of the most important

lessons learned from the developed system is that focusing on autonomy is core for human robot teaming. Robots need to be able to make decisions on their own, enabling the limited human resource to only act in critical situations.

## Acknowledgments

This work was supported through the DARPA Subterranean Challenge, cooperative agreement number HR0011-18-2-0043. A special thanks to Andrew Beathard, Nikolaas Bender, Cesar Galan, Nicole Gunderson, Davis Landry, Greg Lund, Cole Radetich, Ben Rautio, Zoe Turin for assisting with the design, testing, and deployments of our platforms and algorithms. Thank you to the Colorado School of Mines and the Edgar Experimental Mine for allowing us to conduct mock deployments in the mine. Special thanks also to Simon Wunderlich and his team at Meshmerize GmbH for their mesh networking support. Finally, thank you to all of the DARPA staff who have planned and executed absolutely incredible Subterranean Challenge system track circuit events.

## ORCID

Harel Biggie  <https://orcid.org/0000-0002-6281-1218>  
 Eugene R. Rush  <https://orcid.org/0000-0001-9451-224X>  
 Danny G. Riley  <https://orcid.org/0000-0003-2518-3088>  
 Shakeeb Ahmad  <https://orcid.org/0000-0002-6251-2146>  
 Michael T. Ohradzansky  <https://orcid.org/0000-0002-6403-9137>  
 Kyle Harlow  <https://orcid.org/0000-0002-9092-3655>  
 Michael J. Miles  <https://orcid.org/0000-0003-2894-2449>  
 Daniel Torres  <https://orcid.org/0000-0001-8674-0381>  
 Steve McGuire  <https://orcid.org/0000-0003-4650-7950>  
 Eric W. Frew  <https://orcid.org/0000-0003-3686-089X>  
 Christoffer Heckman  <https://orcid.org/0000-0002-9651-6866>  
 J. Sean Humbert  <https://orcid.org/0000-0002-0863-875X>

## References

- Agha, A., Otsu, K., Morrell, B., Fan, D. D., Thakker, R., Santamaria-Navarro, A., Kim, S.-K., Bouman, A., Lei, X., Edlund, J. A., Ginting, M. F., Ebadi, K., Anderson, M. O., Pailevanian, T., Terry, E., Wolf, M. T., Tagliabue, A., Vaquero, T. S., Palieri, M., Tepsuporn, S., Chang, Y., Kalantari, A., Chavez, F., Lopez, B. T., Funabiki, N., Miles, G., Touma, T., Buscicchio, A., Tordesillas, J., Alatur, N., Nash, J., Walsh, W., Jung, S., Lee, H., Kanellakis, C., Mayo, J., Harper, S., Kaufmann, M., Dixit, A., Correa, G., Lee, C.-A., Gao, J. L., Merewether, G. B., Maldonado-Contreras, J., Salhotra, G., da Silva, M. S., Ramtoula, B., Fakoorian, S. A., Hatteland, A., Kim, T., Bartlett, T., Stephens, A., Kim, L., Bergh, C. F., Heiden, E., Lew, T., Cauligi, A., Heywood, T., Kramer, A., Leopold, H. A., Choi, C. S., Daftry, S., Toupet, O., Wee, I., Thakur, A., Feras, M., Beltrame, G., Nikolakopoulos, G., Shim, D. H., Carlone, L., and Burdick, J. W. (2021). Nebula: Quest for robotic autonomy in challenging environments; team costar at the darpa subterranean challenge. *ArXiv*, abs/2103.11470.
- Ahmad, S. and Humbert, J. S. (2022). Efficient sampling-based planning for subterranean exploration. In *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE.
- Ahmad, S., Mills, A. B., Rush, E. R., Frew, E. W., and Humbert, J. S. (2021a). 3d reactive control and frontier-based exploration for unstructured environments. In *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 2289–2296. IEEE.
- Ahmad, S., Sunberg, Z. N., and Humbert, J. S. (2021b). Apf-pf: Probabilistic depth perception for 3d reactive obstacle avoidance. In *2021 American Control Conference (ACC)*, pages 32–39.
- Bay, H., Ess, A., Tuytelaars, T., and Van Gool, L. (2008). Speeded-up robust features (SURF). *Comput. Vis. Image Underst.*, 110(3):346–359.
- Brommer, C., Malyuta, D., Hentzen, D., and Brockers, R. (2018). Long-duration autonomy for small rotorcraft UAS including recharging. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, page arXiv:1810.05683. IEEE.

- Cheung, W. and Hamarneh, G. (2009). n0-sift: n -dimensional scale invariant feature transform. *IEEE Transactions on Image Processing*, 18(9):2012–2021.
- Clausen, T., Jacquet, P., Adjih, C., Laouiti, A., Minet, P., Muhlethaler, P., Qayyum, A., and Viennot, L. (2003). Optimized link state routing protocol (OLSR).
- Coulter, R. C. (1992). Implementation of the pure pursuit path tracking algorithm. Technical report, The Robotics Institute, Carnegie Mellon University.
- Dang, T., Khattak, S., Mascarich, F., and Alexis, K. (2019). Explore locally, plan globally: A path planning framework for autonomous robotic exploration in subterranean environments. In *2019 19th International Conference on Advanced Robotics (ICAR)*, pages 9–16. IEEE.
- Dang, T., Tranzatto, M., Khattak, S., Mascarich, F., Alexis, K., and Hutter, M. (2020). Graph-based subterranean exploration path planning using aerial and legged robots. *Journal of Field Robotics*, 37(8):1363–1388.
- DARPA (2022). DARPA Subterranean Challenge. <https://www.darpa.mil/program/darpa-subterranean-challenge>.
- Ebadi, K., Chang, Y., Palieri, M., Stephens, A., Hatteland, A., Heiden, E., Thakur, A., Funabiki, N., Morrell, B., Wood, S., Carlone, L., and Agha-Mohammadi, A. A. (2020). LAMP: Large-Scale Autonomous Mapping and Positioning for Exploration of Perceptually-Degraded Subterranean Environments. *IEEE International Conference on Robotics and Automation*, pages 80–86.
- Fan, D. D., Otsu, K., Kubo, Y., Dixit, A., Burdick, J., and Agha-Mohammadi, A.-A. (2021). STEP: Stochastic traversability evaluation and planning for safe off-road navigation. *arXiv preprint arXiv:2103.02828*.
- Fankhauser, P., Bloesch, M., and Hutter, M. (2018). Probabilistic terrain mapping for mobile robots with uncertain localization. *IEEE Robotics and Automation Letters*, 3(4):3019–3026.
- Fankhauser, P. and Hutter, M. (2016). A universal grid map library: Implementation and use case for rough terrain navigation. In *Robot Operating System (ROS)*, pages 99–120. Springer.
- for Standardization, I. O. (1996). Information technology – Open Systems Interconnection – Basic Reference Model: The Basic Model (OSII) approach. Standard, International Organization for Standardization, Geneva, CH.
- Forster, C., Carlone, L., Dellaert, F., and Scaramuzza, D. (2015). Imu preintegration on manifold for efficient visual-inertial maximum-a-posteriori estimation. In *Robotics: Science and Systems*. Georgia Institute of Technology.
- Girshick, R. (2015). Fast R-CNN. In *Proceedings of the IEEE international conference on computer vision*, pages 1440–1448.
- Hess, W., Kohler, D., Rapp, H., and Andor, D. (2016). Real-time loop closure in 2D LIDAR SLAM. In *2016 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1271–1278.
- Hines, T., Stepanas, K., Talbot, F., Sa, I., Lewis, J., Hernandez, E., Kottege, N., and Hudson, N. (2021). Virtual surfaces and attitude aware planning and behaviours for negative obstacle navigation. *IEEE Robotics and Automation Letters*, 6(2):4048–4055.
- Horn, B. K. (1987). Closed-form solution of absolute orientation using unit quaternions. *Josa a*, 4(4):629–642.
- Hornung, A., Wurm, K. M., Bennewitz, M., Stachniss, C., and Burgard, W. (2013). Octomap: An efficient probabilistic 3d mapping framework based on octrees. *Autonomous robots*, 34(3):189–206.
- Huang, Y.-W., Lu, C.-L., Chen, K.-L., Ser, P.-S., Huang, J.-T., Shen, Y.-C., Chen, P.-W., Chang, P.-K., Lee, S.-C., and Wang, H.-C. (2019). Duckiefloat: a Collision-Tolerant Resource-Constrained Blimp for Long-Term Autonomy in Subterranean Environments.
- Hudson, N., Talbot, F., Cox, M., Williams, J., Hines, T., Pitt, A., Wood, B., Frousheger, D., Surdo, K. L., Molnar, T., Steindl, R., Wildie, M., Sa, I., Kottege, N., Stepanas, K., Hernandez, E., Catt, G., Docherty, W., Tidd, B., Tam, B., Murrell, S., Bessell, M., Hanson, L., Tychsen-Smith, L., Suzuki, H., Overs, L., Kendoul, F., Wagner, G., Palmer, D., Milani, P., O’Brien, M., Jiang, S., Chen, S., and Arkin, R. C. (2021). Heterogeneous ground and air platforms, homogeneous sensing: Team csiro data61’s approach to the darpa subterranean challenge.
- Juan, S. H. and Cotarelo, F. H. (2015). Multimaster ROS systems. *Institut de Robotics and Industrial Informatics*, pages 1–18.
- Khattak, S., Nguyen, H., Mascarich, F., Dang, T., and Alexis, K. (2020). Complementary Multi-Modal Sensor Fusion for Resilient Robot Pose Estimation in Subterranean Environments. In *2020 International Conference on Unmanned Aircraft Systems (ICUAS)*, pages 1024–1029, Athens, Greece. IEEE.

- Khattak, S., Papachristos, C., and Alexis, K. (2019). Keyframe-based Direct Thermal-Inertial Odometry. In *2019 International Conference on Robotics and Automation (ICRA)*, pages 3563–3569, Montreal, QC, Canada. IEEE.
- Ko, J., Stewart, B., Fox, D., Konolige, K., and Limketkai, B. (2003). A Practical, Decision-theoretic Approach to Multi-robot Mapping and Exploration. In *Proceedings 2003 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2003)*, volume 4, pages 3232–3238, Las Vegas, NV.
- Krüsi, P., Furgale, P., Bosse, M., and Siegwart, R. (2017). Driving on point clouds: Motion planning, trajectory optimization, and terrain assessment in generic nonplanar environments. *Journal of Field Robotics*, 34(5):940–984.
- Lajoie, P. Y., Ramtoula, B., Chang, Y., Carlone, L., and Beltrame, G. (2020). DOOR-SLAM: Distributed, Online, and Outlier Resilient SLAM for Robotic Teams. *IEEE Robotics and Automation Letters*, 5(2):1656–1663.
- LaValle, S. M. (2006). *Planning algorithms*. Cambridge university press.
- Leutenegger, S., Lynen, S., Bosse, M., Siegwart, R., and Furgale, P. (2015). Keyframe-based visual-inertial odometry using nonlinear optimization. *The International Journal of Robotics Research*, 34(3):314–334.
- Malyuta, D., Brommer, C., Hentzen, D., Stastny, T., Siegwart, R., and Brockers, R. (2019). Long-duration fully autonomous operation of rotorcraft unmanned aerial systems for remote-sensing data acquisition. *Journal of Field Robotics*, page arXiv:1908.06381.
- Maturana, D. and Scherer, S. (2015). Voxnet: A 3d convolutional neural network for real-time object recognition. In *2015 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 922–928. IEEE.
- Meagher, D. (1982). Geometric modeling using octree encoding. *Computer graphics and image processing*, 19(2):129–147.
- Miller, I. D., Cohen, A., Kulkarni, A., Laney, J., Taylor, C. J., Kumar, V., Cladera, F., Cowley, A., Shivakumar, S. S., Lee, E. S., Jarin-Lipschitz, L., Bhat, A., Rodrigues, N., and Zhou, A. (2020). Mine Tunnel Exploration Using Multiple Quadrupedal Robots. *IEEE Robotics and Automation Letters*, 5(2):2840–2847.
- Nobre, F., Heckman, C. R., and Sibley, G. T. (2017). Multi-sensor slam with online self-calibration and change detection. In Kulić, D., Nakamura, Y., Khatib, O., and Venture, G., editors, *2016 International Symposium on Experimental Robotics*, pages 764–774, Cham. Springer International Publishing.
- Ohradzansky, M. T., Mills, A. B., Rush, E. R., Riley, D. G., Frew, E. W., and Sean Humbert, J. (2020). Reactive control and metric-topological planning for exploration. In *2020 IEEE International Conference on Robotics and Automation (ICRA)*, pages 4073–4079.
- Ohradzansky, M. T., Rush, E. R., Riley, D. G., Mills, A. B., Ahmad, S., McGuire, S., Biggie, H., Harlow, K., Miles, M. J., Frew, E. W., Heckman, C., and Humbert, J. S. (2021). Multi-agent autonomy: Advancements and challenges in subterranean exploration.
- Otsu, K., Tepsuporn, S., Thakker, R., Vaquero, T. S., Edlund, J. A., Walsh, W., Miles, G., Heywood, T., Wolf, M. T., and Agha-Mohammadi, A.-A. (2020). Supervised autonomy for communication-degraded subterranean exploration by a robot team. In *2020 IEEE Aerospace Conference*, pages 1–9.
- Pandi, S., Gabriel, F., Zhdanenko, O., Wunderlich, S., and Fitzek, F. H. (2019). MESHMERIZE: An interactive demo of resilient mesh networks in drones. In *2019 16th IEEE Annual Consumer Communications & Networking Conference (CCNC)*, pages 1–2. IEEE.
- Papachristos, C., Khattak, S., Mascarich, F., Dang, T., and Alexis, K. (2019). Autonomous aerial robotic exploration of subterranean environments relying on morphology-aware path planning. In *2019 International Conference on Unmanned Aircraft Systems (ICUAS)*, pages 299–305.
- Perkins, C., Belding-Royer, E., and Das, S. (2003). RFC3561: Ad hoc on-demand distance vector (AODV) routing.
- Qi, C. R., Su, H., Mo, K., and Guibas, L. J. (2017). Pointnet: Deep learning on point sets for 3d classification and segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 652–660.
- Qin, T., Li, P., and Shen, S. (2018). VINS-Mono: A Robust and Versatile Monocular Visual-Inertial State Estimator. *IEEE Transactions on Robotics*, 34(4):1004–1020.
- Redmon, J., Divvala, S., Girshick, R., and Farhadi, A. (2016). You only look once: Unified, real-time object detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 779–788.

- Redmon, J. and Farhadi, A. (2018). Yolov3: An incremental improvement. *arXiv preprint arXiv:1804.02767*.
- Riley, D. G. (2021). MARBLE Mapping.
- Riley, D. G. and Frew, E. W. (2021). Assessment of Coordinated Heterogeneous Exploration of Complex Environments. In *2021 IEEE Conference on Control Technology and Applications (CCTA)*, San Diego, CA.
- Riley, D. G. and Frew, E. W. (2022). BOBCAT: Behaviors, Objectives and Binary states for Coordinated Autonomous Tasks. *IEEE Robotics and Automation Letters*. Under review.
- Rouček, T., Pecka, M., Čížek, P., Petříček, T., Bayer, J., Šalanský, V., Heřt, D., Petrлік, M., Báča, T., Spurný, V., Pomerleau, F., Kubelka, V., Faigl, J., Zimmermann, K., Saska, M., Svoboda, T., and Krajník, T. (2020). Darpa subterranean challenge: Multi-robotic exploration of underground environments. In Mazal, J., Fagiolini, A., and Vasik, P., editors, *Modelling and Simulation for Autonomous Systems*, pages 274–290, Cham. Springer International Publishing.
- Rusu, R. and Cousins, S. (2011). 3d is here: Point cloud library (pcl). In *Robotics and Automation (ICRA), 2011 IEEE International Conference on*, pages 1–4.
- Santamaria-Navarro, A., Thakker, R., Fan, D. D., Morrell, B., and Agha-mohammadi, A.-a. (2020). Towards Resilient Autonomous Navigation of Drones. *arXiv:2008.09679 [cs]*. arXiv: 2008.09679.
- Schang, A., Rogers, J., and Maio, A. (2021). Map analysis (version 1.0) [source code].
- Scherer, S., Agrawal, V., Best, G., Cao, C., Cujic, K., Darnley, R., and et al. (2021). Resilient and modular subterranean exploration with a team of roving and flying robots. *Field Robotics*.
- Schwarz, M., Beul, M., Droschel, D., Schüller, S., Periyasamy, A. S., Lenz, C., Schreiber, M., and Behnke, S. (2016). Supervised autonomy for exploration and mobile manipulation in rough terrain with a centaur-like robot. *Frontiers in Robotics and AI*, 3:57.
- Seither, D., König, A., and Hollick, M. (2011). Routing performance of wireless mesh networks: A practical evaluation of BATMAN advanced. In *2011 IEEE 36th Conference on Local Computer Networks*, pages 897–904. IEEE.
- Sethian, J. A. (1999). *Level set methods and fast marching methods: evolving interfaces in computational geometry, fluid mechanics, computer vision, and materials science*, volume 3. Cambridge University Press.
- Shamsudin, A. U., Ohno, K., Westfechtel, T., Takahiro, S., Okada, Y., and Tadokoro, S. (2016). Fog removal using laser beam penetration, laser intensity, and geometrical features for 3D measurements in fog-filled room. *Advanced Robotics*, 30(11-12):729–743.
- Shan, T., Englot, B., Meyers, D., Wang, W., Ratti, C., and Rus, D. (2020). LIO-SAM: Tightly-coupled Lidar Inertial Odometry via Smoothing and Mapping. *arXiv:2007.00258 [cs]*. arXiv: 2007.00258.
- Sheng, W., Yang, Q., Ci, S., and Xi, N. (2004). Multi-robot area exploration with limited-range communications. In *2004 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, volume 2, pages 1414–1419, Sendai, Japan.
- Shi, S., Wang, X., and Li, H. (2019). Pointcnn: 3d object proposal generation and detection from point cloud. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 770–779.
- Simmons, R., Apfelbaum, D., Burgard, W., and Fox, D. (2000). Coordination for multi-robot exploration and mapping. In *Proceedings of the AAAI National Conference on Artificial Intelligence*, pages 852–858, Austin, TX.
- 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. QC 20181210.
- Thrun, S., Burgard, W., and Fox, D. (2005). *Probabilistic Robotics (Intelligent Robotics and Autonomous Agents)*. The MIT Press.
- Tranzatto, M., Mascari, F., Bernreiter, L., Godinho, C., Camurri, M., Khattak, S., Dang, T., Reijgwart, V., Loeje, J., Wisth, D., Zimmermann, S., Nguyen, H., Fehr, M., Solanka, L., Buchanan, R., Bjelonic, M., Khedekar, N., Valceschini, M., Jenelten, F., Dharmadhikari, M., Homberger, T., De Petris, P., Wellhausen, L., Kulkarni, M., Miki, T., Hirsch, S., Montenegro, M., Papachristos, C., Tresoldi, F., Carius, J., Valsecchi, G., Lee, J., Meyer, K., Wu, X., Nieto, J., Smith, A., Hutter, M., Siegwart, R., Mueller, M., Fallon, M., and Alexis, K. (2022a). Cerberus: Autonomous legged and aerial robotic exploration in the tunnel and urban circuits of the darpa subterranean challenge.
- Tranzatto, M., Miki, T., Dharmadhikari, M., Bernreiter, L., Kulkarni, M., Mascari, F., Andersson, O., Khattak, S., Hutter, M., Siegwart, R., and Alexis, K. (2022b). Cerberus in the darpa subterranean challenge. *Science Robotics*, 7(66):eabp9742.



- Wang, J. and Olson, E. (2016). AprilTag 2: Efficient and robust fiducial detection. In *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 4193–4198. IEEE.
- Wermelinger, M., Fankhauser, P., Diethelm, R., Krüsi, P., Siegwart, R., and Hutter, M. (2016). Navigation planning for legged robots in challenging terrain. In *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 1184–1189. IEEE.
- Westfechtel, T., Ohno, K., Mertsching, B., Hamada, R., Nickchen, D., Kojima, S., and Tadokoro, S. (2018). Robust stairway-detection and localization method for mobile robots using a graph-based model and competing initializations. *The International Journal of Robotics Research*, 37(12):1463–1483.
- Williams, J., Jiang, S., O’Brien, M., Wagner, G., Hernandez, E., Cox, M., Pitt, A., Arkin, R., and Hudson, N. (2020). Online 3d frontier-based ugv and uav exploration using direct point cloud visibility. In *2020 IEEE International Conference on Multisensor Fusion and Integration for Intelligent Systems (MFI)*, pages 263–270.
- Wisth, D., Camurri, M., Das, S., and Fallon, M. (2021). Unified Multi-Modal Landmark Tracking for Tightly Coupled Lidar-Visual-Inertial Odometry. *IEEE Robotics and Automation Letters*, 6(2):1004–1011.
- Yamauchi, B. (1997). A frontier-based approach for autonomous exploration. In *Proceedings 1997 IEEE International Symposium on Computational Intelligence in Robotics and Automation CIRA '97: Towards New Computational Principles for Robotics and Automation'*, pages 146–151. IEEE.
- Ye, C. and Borenstein, J. (2003). A new terrain mapping method for mobile robots obstacle negotiation. In *Unmanned ground vehicle technology V*, volume 5083, pages 52–62. International Society for Optics and Photonics.
- Zhan, H., Garg, R., Saroj Weerasekera, C., Li, K., Agarwal, H., and Reid, I. (2018). Unsupervised learning of monocular depth estimation and visual odometry with deep feature reconstruction. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Zhou, Y. and Tuzel, O. (2018). Voxelnet: End-to-end learning for point cloud based 3d object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4490–4499.
- Zlot, R., Stentz, A. T., Dias, M. B., and Thayer, S. (2002). Multi-robot exploration controlled by a market economy. In *Proceedings of the IEEE International Conference on Robotics and Automation*, volume 3, pages 3016–3023, Washington, DC.
- Zou, Z., Shi, Z., Guo, Y., and Ye, J. (2019). Object detection in 20 years: A survey. *arXiv preprint arXiv:1905.05055*.

**How to cite this article:** Biggie, H., Rush, E. R., Riley, D. G., Ahmad, S., Ohradzansky, M. T., Harlow, K., Miles, M. J., Torres, D., McGuire, S., Frew, E. W., Heckman, C., & Humbert, J. S. (2023). Flexible supervised autonomy for exploration in subterranean environments. *Field Robotics*, 3, 125–189.

**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.

## 13. Appendix


### 13.1. Communication Beacon Design

Each beacon consists of two 3D printed nylon-carbon fiber infused internal brackets serving as structural support components. Figure 33a shows the top of each beacon that contains a charging port, power button and LED power status indicator between the antennas. The two stabilizers seen towards the front provide longitudinal stability for the beacon once deployed on the ground. These complement the steel counterweights on the rear of the beacon. The deployment mechanism shown in Figure 33b uses solenoids to hold the beacons in place. When the solenoid is released, the beacon falls at a consistent rate using a constant force spring.

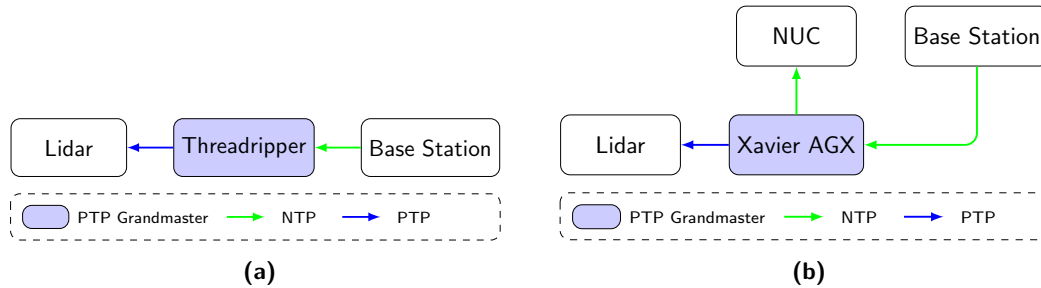
### 13.2. Platform Compute Design

When designing our system, we sought to balance the ease of development with a single, monolithic compute unit versus potential integration challenges of a distributed system. Early on, we decided that, given uncertain compute requirements, we should attempt to pack as much compute as possible into the Husky platforms. This decision had a wide array of collateral consequences, including power system requirements, cooling requirements, and mechanical considerations. Further, the decision to not utilize an industrial-style motherboard, but rather a consumer gaming motherboard, created integration difficulties that might have otherwise been avoided. For example, the power requirements of our Ryzen Threadripper-based system were roughly 425 W at peak consumption of both CPU and GPUs, at the limit of potential commercially available DC/DC ATX power supplies. Using discrete GPUs mounted in PCI Express slots presented a mechanical challenge, particularly for shock and vibration mounting, which was discovered as a failure mode late in the design process. In contrast, the distributed architecture developed for our Spot platforms utilized significantly less power and space while delivering similar performance. Future system designs are more likely to follow a distributed approach, rather than a monolithic approach, to ease mechanical and electrical integration efforts at the expense of only minimal added software effort.

At a deeper level of interrobot module communications, we underestimated the challenges of differing ground potentials. After shorting out several serial links between components and having unreliable USB communication, we realized that several ground loops were responsible. By adding serial optocouplers and converting to Ethernet-based platform control, these ground loops were



**Figure 33.** On the left (a) a perspective view of the beacon design as well as (b) a side view of the beacon attached to the deployment mechanism.



**Figure 34.** Block diagrams showing Precision Time Protocol (PTP) distribution between system components on the Threadripper (a) and Xavier+NUC (b).

eliminated, resulting in highly reliable platforms. Future development would rely exclusively on differential signalling such as controller-area network or Ethernet for intermodule communication.

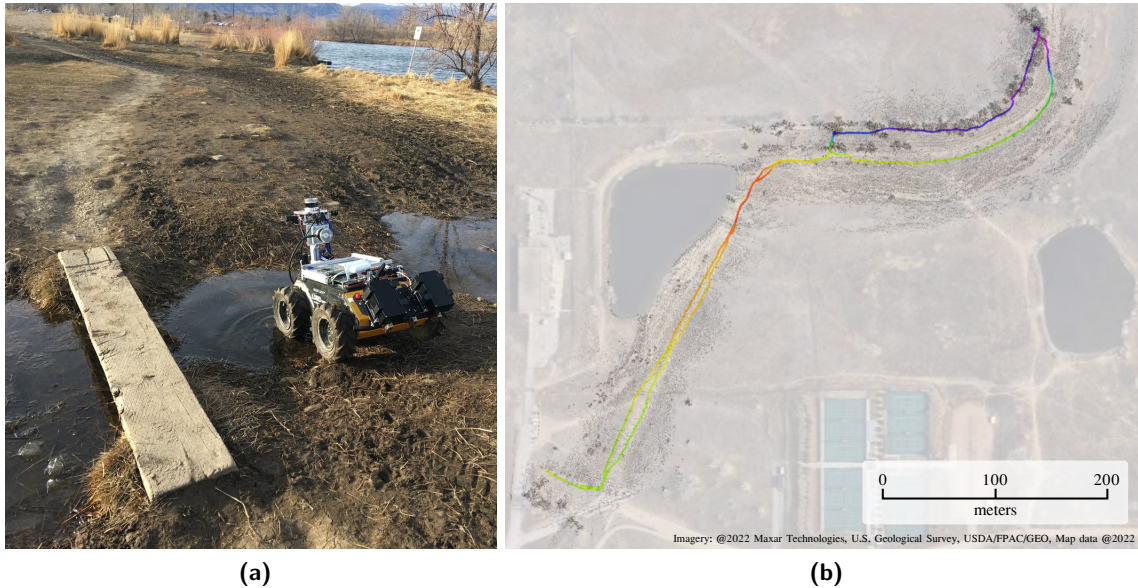
### 13.3. Sensor Synchronization Design

To effectively share sensor data between robots, sensor and system timing has to be considered. Our lidar solution could utilize IEEE 1588v2 timing (also known as the Precision Time Protocol v2 or PTP), but our wireless mesh networking solution could not support IEEE-1588v2. Therefore, we implement Network Time Protocol (NTP) between robots and PTP within the same robot. On startup, each robot attempts to synchronize with the Base Station using NTP over the mesh network. This synchronization step is critical for multi-robot operations and coordination to provide a consistent time basis across all nodes. In testing, the relative time drift (a few milliseconds) over the course of a run (1 hour) was not significant enough to cause problems. If the Base Station is unreachable (say for single-robot testing), the robot falls back to its own battery-backed realtime clock as a time source. In either case, after attempted synchronization, no further attempts are made to match times with any other robot or the Base Station. In testing, we observed that clock slews resulting from attempted time synchronizations as robots entered and left communications range had a negative impact on localization performance.

In the Threadripper monolithic architecture used on the Husky platforms, PTP on the secondary Ethernet interface functioned perfectly, allowing the Threadripper to become the grandmaster and the lidar to follow along, as shown in Figure 34a. However, in the distributed Xavier+NUC architecture (shown in Figure 34b), designing an effective PTP interface encountered significant challenges.

Fundamentally, PTP requires hardware support in order to function by performing sensitive timing operations as close to transport medium as possible. The network hardware options on the Spot included a Realtek r8125, an Nvidia platform SoC module, and a quad-port Intel i210 card. Realtek r8125 support for PTP was not functional, as verified by the *phc\_ctl* utility; this relatively new chipset relies on an out-of-tree kernel driver for Linux at the time of our development. The Nvidia platform module appeared to support PTP when interrogated by *phc\_ctl*, but on further investigation through network traffic inspection, the platform module was not inserting the correct information into outbound Ethernet traffic. Our final solution is based on using a spare port on the Intel i210 card, which had robust, verifiable PTP support. As the NUC lacks a port with viable PTP support, we fall back on NTP as a synchronization method, relying on the Xavier as the robot’s grandmaster time source.

As an aid to the community, we offer the following verification steps to assist in debugging PTP issues. First, verify that there is hardware support via *ethtool -T <iface>* to verify kernel-level hardware PTP support. Second, use *phc\_ctl <iface> cmp* to verify that the Ethernet hardware clock is synchronized to the Linux system clock. Finally, ensure that the hardware timestamps encoded in the PTP network traffic match the local system by capturing network traffic from a PTP-enabled grandmaster Ethernet interface. These steps can verify that the PTP software stack is broadcasting the system time via the Ethernet hardware to downstream consumers.



**Figure 35.** On the left (a) photo taken at CU South Campus of a Husky robot during a large-scale, long-duration localization test. On the right (b), is an LIO-SAM point cloud map, denoted by small black dots, and LIO-SAM robot trajectory denoted by larger dots colored by elevation, overlaid with Google Maps satellite imagery of the area. The Husky was manually controlled, beginning in the CU South parking lot, continuing along a dirt path and up a hill, looping back, and ending back at the parking lot.

### 13.4. Large-Scale Localization Validation

Validation testing of LIO-SAM onboard Spot and Husky platforms was imperative for ensuring sufficient accuracy, speed, and stability for long-duration and large-scale missions. Some examples include an outdoor test at CU South Campus, as shown in Figure 35, and as well as test that begins at the CU Engineering Center building and treks across campus to the bottom of a three-story underground parking garage, as shown in Figure 36.

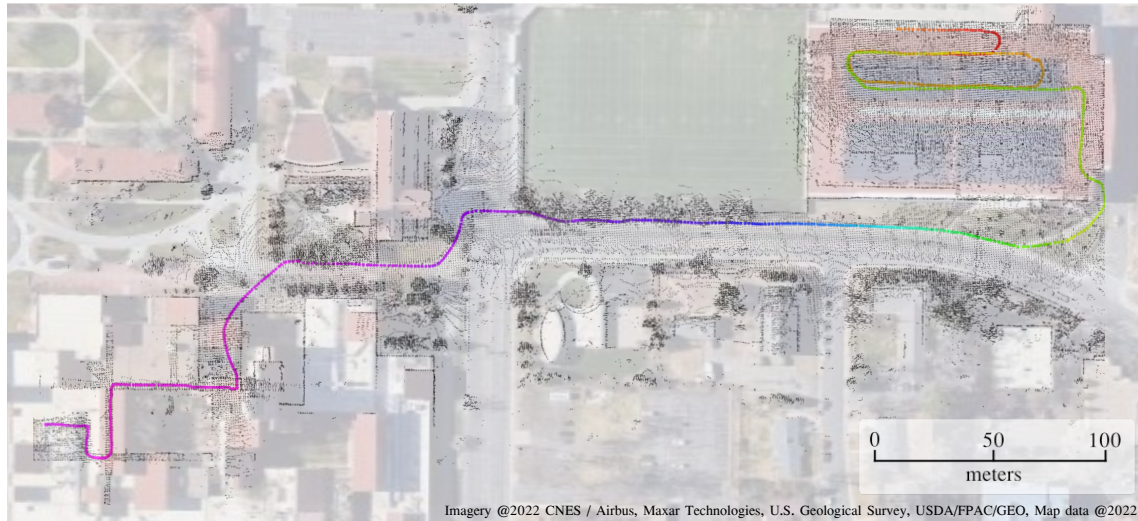
### 13.5. Common Reference Frame Alignment Optimization

To further reduce the yaw error of the common reference frame alignment, the lateral spacing of the robot prisms was increased.

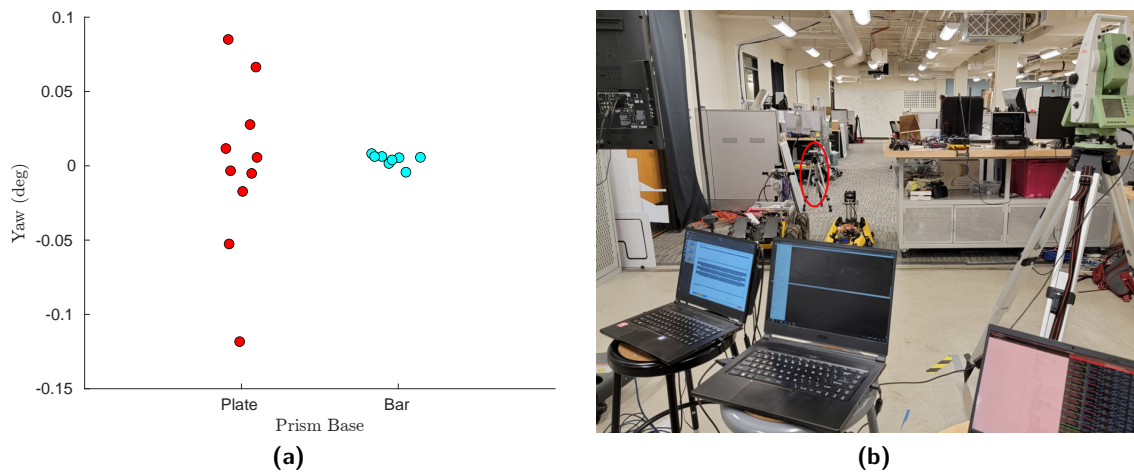
In our initial testing, we found that the roughly 120-mm plates attached to the robots along with the prisms 1.5-mm centering error lead to consistent variance in the resulting yaw of approximately  $0.7^\circ$ . This aligned with a calculated value of  $\arcsin(1.5/120) = 0.72^\circ$ . In order to reduce the impact of this error, a mechanical bar holding two of the prisms at a distance of 655 mm, is added to each robot. The resulting angle error after adding the bar was  $\arcsin(1.5/655) = 0.13^\circ$ . This was consistent with external testing. For a rough comparison see Figure 37a, where a test was conducted using a mock robot plate and the prism separating bar. The difference from the average of each test shows a higher precision for the tests conducted with the prism bar in place. An example setup for these test is shown in Figure 37b.

### 13.6. Artifact Detection Training Procedure

A systematic procedure targeted at low-light conditions is used to train the model. At each location, data was collected using three different brightness levels to minimize the impact of lighting conditions on the model's performance. Specifically, images were taken from past circuit events as well as



**Figure 36.** Overlay of LIO-SAM point cloud map, denoted by small black dots, and LIO-SAM robot trajectory denoted by larger dots colored by elevation, with Google Maps during a large-scale, long-duration localization test at CU Main Campus. This test was conducted on a Spot robot, that was manually controlled, beginning in the CU Engineering Center building on the bottom left corner of the map, continuing across campus, ending at the bottom of the three-story underground Folsom Parking Garage on the top right corner of the map.



**Figure 37.** A comparison of transforms generated by LTS using the standard robot sensor plate, and after attaching prisms to a bar placed on the plate instead (a). The difference from the average, and precision of the bar is higher than without the bar. Setup for a prism test with a mock gate highlighted in red (b). The mock gate was designed to have the same dimension as a robot sensor plate.

separate field exercises. Remote data collection sessions took place inside the dark and rocky Edgar Experimental Mine in Idaho Springs, Colorado. Local data collection took place on University of Colorado Boulder campus, primarily within the outdoor courtyard of our Engineering Center building, and during evening hours when there was no natural illumination. The data was collected with three onboard illumination levels: 0%, 50%, and 100%. The cameras, FLIR Blackfly PGE-05S2C-CS GigE Vision cameras, were mounted in cardinal directions on the robots as shown in Figure 5b.

**Table 8.** MARBLE *Monitors*, with a description of robot state requirements for an output of 1.

Monitor	Criteria
ExploreToGoal	Received command to explore to a specific goalpoint, either submitted by the human supervisor or generated by a node other than the global planner.
iExplore iGoToGoal iStop iDeployBeacon iGoHome	The associated input command has been sent by the human supervisor from the GUI or joystick to execute a specific behavior.
NearbyRobot	This and any other robot's paths are within 2m of each other.
Beacon	A communications beacon is available and other criteria has been met to deploy it.
ReverseDrop	A communications beacon is available and communications have been lost with the base station for 10 seconds.
Comms	Any message has been received from the base station in the last 3 seconds.
Artifact	There are at least 3 unreported artifacts, or it has been at least 5 minutes since the first unreported artifact was detected.

**Table 9.** MARBLE *Objectives*. Evaluation functions calculate the weight for each objective. Components prefixed by “iw” represent the input weight of the objective, while monitor names represent the binary output of the monitor.

Objective	Evaluation Function	Description
FindArtifacts	$iwFindArtifacts$	Find, identify and localize artifacts. Always active.
Input	$iwInput * OR(Input\ Monitors)$	Allow supervisor to override autonomy when necessary.
BeSafe	$iwBeSafe * NearbyRobot$	Safety of the robot, particularly collision with other robots.
ExtendComms	$iwExtendComms * (Beacon \parallel ReverseDrop)$	Extend communications as far into the environment as possible to reduce any delay in reports and minimize robot travel back and forth.
MaintainComms	$iwMaintainComms * !Comms$	Communicate with the base station, either by staying in or returning to communications.
ReportArtifacts	$iwReportArtifacts * Artifact$	Report artifact types and locations to base station.

Photos in which the artifacts suffered excessive motion blur and occlusions, determined by the ability of the human reviewer to detect the artifact, were removed from the data set. After the AD pipeline was trained on this initial data collection effort, we found that it did not generalize well to new environments. Therefore, we augmented the dataset with additional imagery collected from a greater diversity of backgrounds, including a nearby university loading dock. The dataset was later augmented with data from areas where false positives were frequently identified in order to reduce the identification of these false positives.








### 13.7. BOBCAT Components

*BOBCAT* calculates *Objective* weights and *Behavior* scores to select an *Execution Behavior* whenever one or more *Monitor* outputs or *Objective* input weights change. *Behavior* execution functions may be either blocking or nonblocking. They should be nonblocking to the maximum extent possible, to increase reactivity and allow *Behaviors* to change at any time. *Behaviors* that need to block may be required for certain actions that must be completed before the robot can do something different. If a *Behavior* is blocking, *BOBCAT* will delay evaluation until the actions have completed. Tables 8, 9, and 10 provide an exhaustive list of the *Monitors*, *Objective*, and *Behaviors*, respectively.

**Table 10.** MARBLE Behaviors. Evaluation functions calculate the score for each behavior for use by the policy  $\pi_B$ . Components prefixed by “ow” represent the output weight of the objective, while monitor names represent the binary output of the monitor.

Behavior	Evaluation Function	Actions
Explore	$owFindArtifacts * !ExploreToGoal + owInput * iExplore$	Request global planner to plan to unexplored areas.
GoToGoal	$owFindArtifacts * ExploreToGoal + owInput * iGoToGoal$	Request global planner to plan a path to a specific goalpoint.
Stop	$owBeSafe + owInput * iStop$	Controller stops using plan generated by global planner. Causes robot to stop autonomous movement, but will still accept manual movement by human supervisor.
DeployBeacon	$owExtendComms * !ReverseDrop + owInput * iDeployBeacon$	Initiate beacon deployment maneuver, which positions robot, stops, and deploys a beacon.
GoHome	$owReportArtifacts + owMaintainComms + owExtendComms * ReverseDrop + owInput * iGoHome$	Request global planner to plan a path to the starting point.

**Table 11.** List of data sets utilized to train the artifact detection system. Also listed is the frames per second (FPS) of the imagery, the total number of images, the number of labels for each artifact, and the total number of artifact labels.

Dataset Name	FPS	Images								Labels
<b>Courtyard</b>	10	54523	4837	8299	2682	3190	6219	3928	3654	32809
<b>Edgar Mine</b>	10	3603	254	159	148	142	122	187	84	1096
<b>NIOSH Mine</b>	1	13437	118	45	106	173	0	0	0	442
<b>Satsop Nuclear</b>	15	33953	0	630	0	0	302	0	0	932
<b>Loading Dock</b>	10	6923	184	234	143	143	76	329	116	1225
<b>Total</b>		<b>112439</b>	<b>5393</b>	<b>9367</b>	<b>3079</b>	<b>3648</b>	<b>6719</b>	<b>4444</b>	<b>3854</b>	<b>36504</b>

### 13.8. Artifact Detection Training Data

Table 11 presents a summary of the training data Team MARBLE collected and annotated. Our main data collection effort was conducted in the Engineering Center Courtyard at University of Colorado Boulder, during late evening hours when there was no natural illumination. This location was chosen because we believed it be more representative of a subterranean environment, with the large amount of concrete and low illumination levels. Another smaller dataset was collected at the Edgar Mine in order to introduce data from tunnel-level environments. Real-world imagery from the Tunnel Event at the NIOSH Mine and the Urban Event at the Satsop Nuclear Power Plant was incorporated as well. To round out training, a final data set at the loading dock at the Engineering Center was added.


At the Final Event, our artifact detection system did not correctly detect any vents. It also falsely detected many white walls as vents. One reason for this poor performance is likely because the vent we trained on was different than the vent at the Final Event. To support the vent, white sides were added, making it appear as a box, as shown in Figure 38.

### 13.9. Competition Results

Table 12 lists the eight teams that competed in the DARPA Subterranean Challenge Final Event, and the number of artifacts each team found during the 60-minute Prize Run, out of a maximum possible score of 40.

**Table 12.** List of final scores of all teams that participated in the 60-minute Final Event Prize Run.

Score	Team	Funding
23	CERBERUS	DARPA
23	CSIRO Data61	DARPA
18	MARBLE	DARPA
17	Explorer	DARPA
13	CoSTAR	DARPA
7	CTU-CRAS-NORLAB	DARPA
2	Coordinated Robotics	Self
2	Robotika	Self

**Figure 38.** Team MARBLE trained on a vent with white-walled sides.**Figure 39.** Robot-robot interactions


### 13.10. Planning in a Dynamic Environment

The planning algorithm presented in this paper has the capability of adapting to dynamic environments, i.e., closing or opening of doors, falling rubble, etc. Other situations also lead to dynamic changes in the map, including other nearly mobile agents, as well as localization and mapping error. The paper presents the planner response to a trap door and examples of planner responses to robot-robot encounters as well as localization and mapping error.

Robots often came within close proximity of other robots during the mission. Some examples of this are shown in Figure 39. Rather than resolve these interactions through the planner, which operates on a slower timescale than fast-approaching robots, the autonomous mission management system specifies the higher-priority agent continue while the lower-priority agent wait. The planner marks previously traversable edges as untraversable, and later updates them as traversable again once the other agents leaves the vicinity.

Figure 40 presents an instance when the planner on D01 adapted to localization and mapping drift. The erroneous new map data caused many of the previously traversable edges of the graph to change to untraversable. The planner adapted to the new scenario, helping the agent return to the main corridor, at which point a loop closure occurred, correcting the agent's pose. The planner operated continuously throughout this localization drift and loop closure correction.





**Figure 40.** Instance of D01 experiencing localization drift and erroneous mapping, causing planner to mark traversable (green) edges in the graph as no longer traversable (red). The progression of events is first (a) no localization error (13:03), (b) initial localization and mapping error (17:12), (c, d) continued localization and mapping error (13:53, 17:16), (e) loop closure correcting the robot pose (17:34), teleporting the agent away from its planned path (pink), and (f) continuing on the mission (18:23).

### 13.11. Planning in Constrained Spaces

The planner parameters were configured such that the agents would operate safely and not attempt to traverse highly confined spaces. One limitation to this approach, is that agents could not autonomously plan and traverse some areas, such as the utility corridor with low ceilings and narrow cave section, as shown in Figures 41 and 42, respectively. The human supervisor was able to teleoperate the Spot through the cave section, and it autonomously traversed it when later exiting the cavern.

### 13.12. Artifact Reports

All true positive and false positive artifact detections can be seen in Figures 43 and 44 respectively.




Figure 41. Mobility limitations.




Figure 42. Mobility limitations.





Figure 43. All true positive reports of visual artifacts from autonomous artifact detection systems onboard remote agents D02 [(a) and (b)], D01 [(c)–(g)], H02 (h), and H01 [(i)–(k)], in order from mission start to mission end.



**Figure 44.** All false positive reports of visual artifacts from autonomous artifact detection systems onboard remote agents D02 [(a)–(d)], D01 [(e.)–(h)], H02 (i), and H01 (j), in order from mission start to mission end.

### 13.13. Twenty-Five Explored Artifacts

This section focuses on all 25 artifacts that Team MARBLE agents were in the vicinity of during the Final Event Prize Run. To summarize, 18 of these artifacts were scored, one was not scored, and six were unreported. These three categories of artifacts are discussed in Sections 13.13.1, 13.14, and 13.15, respectively.

Each artifact has at least one attempt associated with it. The eighteen scored artifacts all end with a scored attempt (SA), and some will have multiple missed attempts (MA) before reaching a scored attempt. The one missed artifact was not scored, so it only has missed attempts associated with it. The six unreported artifacts have no attempts associated with them.

#### 13.13.1. Eighteen Scored Artifacts

**L51 Drill (SA1):** This drill was the first artifact scored, within 1 minute and 8 seconds of the mission start. D02 reported and scored the drill as it passed through the first junction of the course, splitting it into tunnel, urban, and cave corridors (SA1). The human supervisor also saw the drill via D02 live FPV view and would have attempted to score it had D02 not automatically reported it.

**L53 Backpack (SA2):** D02 quickly continued into the cave corridor, reporting and scoring the backpack (SA2). If D02 did not autonomously score the backpack, the human supervisor may have scored it via D02 live FPV. Had that failed too, H02 and H01 accurately reported the backpack later in the mission and would have scored it.

**L55 Rope (SA3):** Because D02 had difficulty traversing the narrow cave corridor, the human supervisor manually teleoperated the agent through this section of the course, and in the process, saw the rope via D02 live FPV and reported it (SA3). D02 did not automatically detect the rope, likely due to poor lighting conditions.

**L26 Survivor (SA4):** While D02 immediately explored the cave section, D01 began exploring the urban section and autonomously reported the survivor and scored it (SA4). Later in the mission, H01 accurately reported L26 and would have scored had D01 missed it at the beginning of the mission.

**L32 Survivor (SA5):** The HS saw the survivor through D02 live FPV stream, submitted a manual report and scored the artifact (SA5). Later in the mission, accurate autonomous reports from D01 and H01 would have scored the artifact, had the HS had not already scored it.

**L08 Gas (SA6):** D01 autonomously reported and scored the gas artifact as it traversed the urban environment (SA6).

**L31 Fire Extinguisher (SA7):** During the middle of the mission, the HS teleoperated D01 through a foggy section of the tunnel environment. In this process, the HS saw the fire extinguisher through D01 live FPV stream, and scored the artifact via manual report (SA7). After this manual intervention, D01 went onto help score five more artifacts, L34, L38, L36, L40, and L67. The L31 fire extinguisher was also seen at the end of the mission when the HS was reviewing the D02 archived FPV images.

**L34 Drill (SA8):** After the HS teleoperated D01 through the fog, D01 autonomously reported and scored the drill (SA8). The HS also saw the drill through D01 live FPV stream, and would have reported the artifact manually had D01 not already scored it.

**L38 Fire Extinguisher (SA9):** After teleoperating D01 through the fog, the HS saw the fire extinguisher through D01 live FPV stream and manually scored the artifact (SA9).

**L36 Cube (SA10):** After the HS teleoperated D01 through the fog, D01 autonomously reported and scored the cube (SA10).

**L40 Backpack (SA11):** After the HS teleoperated D01 through the fog, D01 autonomously reported and scored the backpack (SA11).

**L67 Rope (SA12):** After the HS teleoperated D01 through the fog, D01 autonomously reported and scored the rope. The HS also saw the artifact through D01 live FPV stream, and would have manually reported the artifact had D01 not already scored it (SA12).

**L11 Cube (MA5, MA6, SA13):** D01 and D02 both autonomously reported the cube, but were both missed attempts (MA5, MA6), with corresponding errors of 10.73 and 9.42 m. The HS then used the location of the missed attempts to manually submit an adjusted location, scoring the cube (SA13) with an error of 1.83 m.

**L22 Cell Phone (MA2, MA4, SA14):** D01 autonomously reported the cell phone early in the mission, but was a missed attempt (MA2) with an error of 19.47 m. H02 also autonomously reported the cell phone, but the report was a missed attempt (MA4) with an error of 7.77 m. Near the end of the mission, H01 accurately localized the cell phone and scored the artifact (SA14), with an error of 4.06 m.

**L47 Cell Phone (SA15):** D02 autonomously reported the cell phone along the subway platform and scored it (SA15).

**L59 Cell Phone (MA1, MA3, MA10, SA16):** Located in the left branch of the cave section, the cell phone was autonomously reported by D02, but was a missed attempt (MA1) with an error of 13.69 m. The HS then manually submitted the cell phone with an adjusted location, but this was also a missed attempt (MA3) with an error of 7.74 m. Later in the mission, a third missed attempt occurred (MA10), as D01 autonomously reported the cell phone with an error of 9.97 m. Immediately after, the HS used the reported locations from D02 and D01 to submit another manual report with an adjusted location, and scored the artifact (SA16), with an error of 2.15 m.

**L24 Gas (MA11, SA17):** The gas was autonomously reported by H01, but with an error of 5.38 m, resulted in a missed attempt (MA11). Soon after, H02 traversed the same space and autonomously reported the gas independently of H01. This report better estimated the position of the artifact, with an error of 2.55 m, and scored (SA17) the gas artifact.

**L58 Helmet (SA18):** The HS reviewed the D02 archived FPV images near the end of the mission, and saw the helmet in the cavern. The report was accurate to 1.74 m and scored Team


MARBLE'S 18th and final point of the mission (SA18). This point was made possible by the HS teleoperation through the narrow cave corridor.

### 13.14. One Missed Artifact


**L64 Cube (MA7, MA8, MA9):** This cube artifact was located atop a steep slope found along the main corridor of the cave section. D01 autonomously reported the cube but was a missed attempt with an error of 8.24 m (MA7). Based on that experience, the HS manually submitted two reports at adjacent locations, both of which were missed attempts (MA8, MA9), with errors of 12.23 m and 20.15 m. This artifact was never scored during the mission. Figure 45b shows the failed score attempts circled in red, and the actual position circled in green.

### 13.15. Six Unreported Artifacts

**L02 Vent** L02 vent was not detected by the on-board artifact detection system, but was available in the human supervisor's FPV feed. At approximately 20 minutes into the mission, D01 (Spot) was returning home due to a localization error noticed by the human supervisor. While it was returning, the human supervisor turned attention to other robots, and did not scan the FPV feed for some time. The robot stopped for 33 seconds and transmitted a series of images similar to Figure 46. Then



**Figure 45.** Helmet FPV image transmitted to human supervisor, but missed due to workload (a). L64 Cube location, circled in green, and incorrect submissions, circled in red (b).



**Figure 46.** Vent seen by D01 while stopped and transmitted to human supervisor, but missed due to human supervisor workload.



**Figure 47.** Vent detected and reported by D02 late in mission but missed by the human supervisor due to workload.

the robot moved forward and was stuck on the corner underneath the vent, due to the localization error, and this is when the Supervisor returned attention to the robot. The FPV images were stored for later review by the Supervisor, but due to workload, these images were never reviewed during the mission, and thus the artifact remained unscored.

**L05 Vent** Approximately 7 minutes prior to the end of the mission, D02 reported a vent and transmitted the image in Figure 47 to the base station. Unfortunately, due to workload and poor notification design in the GUI, the human supervisor never noticed the report, and thus it was not submitted. Although the detection system identified the bucket as a vent, the Supervisor could easily see the actual vent above it and the reported position was 1.24 m from the actual position.

**L62 Helmet** L62 helmet located in the cave section near the tunnel intersection was seen via FPV, as in Figure 45a but not detected by the robot. This was transmitted by D01 and also available for review by the human supervisor, but due to workload the images were not reviewed during the mission.

**L13 Gas** Approximately 10 minutes prior to the end of the mission, D02 passed approximately within 1m of this gas artifact, but provided no reports. Further analysis shows the only gas detection D02 had was a false report in an area with no CO<sub>2</sub> nearby. The only other robot to go near this gas was D01, but it only went near a doorway leading to the area where the gas was located, and had already reported and scored another gas artifact near that location, so if it did detect L13 it would have assumed it was the same as the previous artifact.

**L21 Vent** This vent was in the subway platform area of the environment. Both Spot robots viewed the vent with various cameras, but never detected it, likely due to the white wall background. Additionally, communications to the base station were limited in this area, and none of the FPV images relayed to the human supervisor had the vent in view. Interestingly, D01 did provide a false report of a vent in this area, and transmitted an image seen in Figure 48, which the human supervisor discarded as a false report. According to the truth data, the location reported was 3.93 m from the actual vent, and so would have scored if submitted. However, visual analysis indicates the vent position in the ground truth file appears 2 m off, which would not have scored.

**L42 Fire Extinguisher** This fire extinguisher was seen only with the right-facing camera on D01 for only a few frames, as seen in Figure 49, which was not enough to trigger a detection and report. The robot was outside of communications, so even if the front camera had seen it, it would not have been available for the human supervisor.



**Figure 48.** False vent reported close to actual vent by D01.



**Figure 49.** Missed fire extinguisher only seen for a few frames by the right camera of D01.

### 13.16. Two False Artifacts

This section focuses on two false artifacts that Team MARBLE reported, but were actually the result of false detections. Each false artifact has at least one false attempt (FA) associated with it. These false attempts are listed in Table 13.

**Gas (FA1, FA2, FA5):** Gas was autonomously reported by H02 early in the mission (FA1). The HS modified the location and manually reported again (FA2), but did not score. Later in the mission, D02 reported gas again in a very similar location as D01. This increased the confidence in the HS that gas was in the area, so the HS modified the location of D02 report, but this too did not score (FA4). It remains unknown why both agents detected elevated levels of  $\text{CO}_2$ , but it is likely that a source in that vicinity existed, even if it was not an gas artifact.



































**Fire Extinguisher (FA3):** Due to an unknown cause, the human supervisor inadvertently submitted the same report twice, which was a fire extinguisher at xyz-coordinates of (20.60, 20.59, -2.64). Because the fire extinguisher was scored by the last report, this report did not score, but simply wasted a report.

**Backpack (FA4):** The image was not clear, but the HS attempted to report it, and it did not score.

### 13.17. Tabulated Reports

All reports that Team MARBLE submitted to DARPA are listed in Table 13.

**Table 13.** List of all artifact reports submitted by Team MARBLE during the 60-minute Final Event Prize Run. Bolded entries represent reports that resulted in a score. There are three types of reports: a *Scored Attempt (SA)*, a *Missed Attempt (MA)* due to error exceeding 5m, and a *False Attempt (FA)* due to a false positive detection. Listed next are artifact ID, artifact type, error, cumulative score, and time since mission start. The scorer is the agent that submitted the report and scored (or attempted to score), and the assister is the agent(s) that provided information that aided the scorer in scoring (or attempting to score). The reporting of these artifacts was completely autonomous, save artifacts scored by the HS as well as those with a (<sup>†</sup>), denoting artifacts that were seen as a result of the HS temporarily teleoperating the agent into new areas of the course.

Report	ID	Type	Error [m]	Score	Time [mm:ss]	Scorer	Assister
<b>SA1</b>	<b>L51</b>		<b>0.57</b>	<b>1</b>	<b>01:08</b>	<b>D02</b>	
<b>SA2</b>	<b>L53</b>		<b>2.23</b>	<b>2</b>	<b>01:23</b>	<b>D02</b>	
<b>SA3</b>	<b>L55</b>		<b>0.84</b>	<b>3</b>	<b>06:23</b>	<b>HS<sup>†</sup></b>	<b>D02 (live FPV stream)</b>
<b>SA4</b>	<b>L26</b>		<b>0.62</b>	<b>4</b>	<b>12:03</b>	<b>D01</b>	
MA1	L59		13.69		14:09	D02	
MA2	L22		19.47		14:13	D01	
MA3	L59		7.74		14:57	HS	D02 (MA1)
<b>SA5</b>	<b>L32</b>		<b>1.40</b>	<b>5</b>	<b>16:35</b>	<b>HS</b>	<b>D02 (live FPV stream)</b>
<b>SA6</b>	<b>L08</b>		<b>1.80</b>	<b>6</b>	<b>17:23</b>	<b>D01</b>	
FA1	—		—		17:33	H02	
FA2	—		—		17:59	HS	H02 (FA1)
MA4	L22		7.77		18:38	H02	
MA5	L11		10.73		19:49	D01	
<b>SA7</b>	<b>L31</b>		<b>1.31</b>	<b>7</b>	<b>28:08</b>	<b>HS<sup>†</sup></b>	<b>D01 (live FPV stream)</b>
FA3	—		—		33:38	HS	
<b>SA8</b>	<b>L34</b>		<b>1.43</b>	<b>8</b>	<b>35:51</b>	<b>D01<sup>†</sup></b>	
<b>SA9</b>	<b>L38</b>		<b>2.82</b>	<b>9</b>	<b>36:53</b>	<b>HS<sup>†</sup></b>	<b>D01 (live FPV stream)</b>
<b>SA10</b>	<b>L36</b>		<b>3.94</b>	<b>10</b>	<b>37:08</b>	<b>D01<sup>†</sup></b>	
<b>SA11</b>	<b>L40</b>		<b>1.40</b>	<b>11</b>	<b>37:58</b>	<b>D01<sup>†</sup></b>	
<b>SA12</b>	<b>L67</b>		<b>2.87</b>	<b>12</b>	<b>38:47</b>	<b>D01<sup>†</sup></b>	
FA4	—		—		45:18	D02	
FA5	—		—		46:44	HS	H02 (FA1), HS (FA2) & D02
MA6	L11		9.42		47:31	D02	
<b>SA13</b>	<b>L11</b>		<b>1.83</b>	<b>13</b>	<b>47:53</b>	<b>HS</b>	<b>D01 (MA5), D02 (MA6)</b>
MA7	L64		8.24		48:12	D01	
MA8	L64		12.23		48:42	HS	D01 (MA7)
MA9	L64		20.15		49:09	HS	D01 (MA7) & HS (MA8)
<b>SA14</b>	<b>L22</b>		<b>4.06</b>	<b>14</b>	<b>50:33</b>	<b>H01</b>	<b>D01 (MA2), H02 (MA4), D02</b>
<b>SA15</b>	<b>L47</b>		<b>4.00</b>	<b>15</b>	<b>50:45</b>	<b>D02</b>	
MA10	L59		9.07		50:57	D01	
<b>SA16</b>	<b>L59</b>		<b>3.15</b>	<b>16</b>	<b>51:48</b>	<b>HS</b>	<b>D02 (MA1), HS (MA3), D01 (MA10)</b>
MA11	L24		5.38		52:20	H01	
<b>SA17</b>	<b>L24</b>		<b>2.55</b>	<b>17</b>	<b>52:45</b>	<b>H02</b>	
<b>SA18</b>	<b>L58</b>		<b>1.74</b>	<b>18</b>	<b>56:33</b>	<b>HS<sup>†</sup></b>	<b>D02 (archived FPV images)</b>



**Table 14.** List of all full-scale field deployments. For the Final Event, Team MARBLE gave greater resources to system performance validation and human-robot teaming practice. This realized itself as more frequent and diverse field deployments. For the circuit events, just two to three weeks were spent on field deployments, whereas for the Final Event, the team devoted two months. Instead of practicing in just one or two environments, the team was asked to perform in five unique environments. Selecting locations that were “onsite” University of Colorado Boulder campus enabled the team to be more nimble and operationally efficient. The (\*) denotes that the Cave Event was a self-managed mock event.

Date	Deployment	Environment	Onsite	Location
<i>Apr 7, 2019</i>	<b>STIX Event</b>	<i>Edgar Experimental Mine</i>		<i>Idaho Springs, CO</i>
Jul 29, 2019	Pre-Tunnel 1	Edgar Experimental Mine		Idaho Springs, CO
Aug 1, 2019	Pre-Tunnel 2	Edgar Experimental Mine		Idaho Springs, CO
Aug 6, 2019	Pre-Tunnel 3	Edgar Experimental Mine		Idaho Springs, CO
Aug 9, 2019	Pre-Tunnel 4	Edgar Experimental Mine		Idaho Springs, CO
Aug 12, 2019	Pre-Tunnel 5	Edgar Experimental Mine		Idaho Springs, CO
<i>Aug 17, 2019</i>	<b>Tunnel Event</b>	<i>NIOSH Exp. &amp; Safety Research Mines</i>		<i>Pittsburgh, PA</i>
Feb 4, 2020	Pre-Urban 1	Geotech Warehouse		Denver, CO
Feb 8, 2020	Pre-Urban 2	Geotech Warehouse		Denver, CO
Feb 12, 2020	Pre-Urban 3	Geotech Warehouse		Denver, CO
<i>Feb 21, 2020</i>	<b>Urban Event</b>	<i>Satsop Nuclear Power Plant</i>		<i>Elma, WA</i>
Aug 4, 2020	Pre-Cave 1	Edgar Experimental Mine		Idaho Springs, CO
Sep 17, 2020	Pre-Cave 2	Eng. Center (L1)	X	Boulder, CO
Sep 19, 2020	Pre-Cave 3	Eng. Center (L1)	X	Boulder, CO
<i>Sep 21, 2020</i>	<b>Cave Event*</b>	<i>Edgar Experimental Mine</i>		<i>Idaho Springs, CO</i>
Jul 13, 2021	Pre-Final 1	Folsom Parking Garage	X	Boulder, CO
Jul 14, 2021	Pre-Final 2	Folsom Parking Garage	X	Boulder, CO
Jul 15, 2021	Pre-Final 3	Edgar Experimental Mine		Idaho Springs, CO
Aug 13, 2021	Pre-Final 4	Folsom Parking Garage	X	Boulder, CO
Aug 17, 2021	Pre-Final 5	Eng. Center (LL & Courtyard)	X	Boulder, CO
Aug 19, 2021	Pre-Final 6	Sust., Energy, and Env. Community	X	Boulder, CO
Aug 24, 2021	Pre-Final 7	Eng. Center (LL & Courtyard)	X	Boulder, CO
Aug 26, 2021	Pre-Final 8	Edgar Experimental Mine		Idaho Springs, CO
Sep 1, 2021	Pre-Final 9	Edgar Experimental Mine		Idaho Springs, CO
Sep 8, 2021	Pre-Final 10	Eng. Center (LL & Courtyard)	X	Boulder, CO
Sep 10, 2021	Pre-Final 11	Eng. Center (LL & Courtyard)		Boulder, CO
Sep 12, 2021	Pre-Final 12	Eng. Center (L2) + Rustandy	X	Boulder, CO
Sep 14, 2021	Pre-Final 13	Eng. Center (L2) + Rustandy	X	Boulder, CO
<i>Sep 21, 2021</i>	<b>Final Event</b>	<i>Louisville, Megacavern</i>		<i>Louisville, KY</i>

### 13.18. Field Deployments

Table 14 lists the dates and locations of all full-scale deployments, within the context of the events at the DARPA SubT Challenge.