

Recent Developments in Subterranean Robotics

Aaron Morris, Dave Ferguson,
Zachary Omohundro, David Bradley,
David Silver, Chris Baker, Scott Thayer,
Chuck Whittaker, and William Whittaker
*Field Robotics Center
Carnegie Mellon University
Pittsburgh, Pennsylvania 15213*

Received 31 August 2005; accepted 1 November 2005

Robotic systems exhibit remarkable capability for exploring and mapping subterranean voids. Information about subterranean spaces has immense value for civil, security, and commercial applications where problems, such as encroachment, collapse, flooding and subsidence can occur. Contemporary method for underground mapping, such as human surveys and geophysical techniques, can provide estimates of void location, but cannot achieve the coverage, quality, or economy of robotic approaches. This article presents the challenges, mechanisms, sensing, and software of subterranean robots. Results obtained from operations in active, abandoned, and submerged subterranean spaces will also be shown. © 2006 Wiley Periodicals, Inc.

1. INTRODUCTION

Subterranean spaces, such as caves, sewers, and mines (see Figure 1), create problems for operations both above and below the surface. Subsidence, ground water contamination, and mine floods, such as those of Quecreek,¹ Zapadnaya,² and Daxing³ are just a few consequences of poor or inaccurate documentation of subterranean voids. In the United States

alone, “tens of thousands, perhaps even hundreds of thousands, of abandoned mines are believed to exist. Not even the Department of Surface Mines knows the exact number, because federal recording of mining claims was not required until 1976” (Belwood & Waugh, 1991). With continued above-ground and below-ground operations encroaching upon such spaces, the need for improved documentation of subterranean voids is clear.

Knowledge of subterranean spaces has tremendous value across a number of applications. In addition to preventing and mitigating the aforementioned problems, void documentation plays key roles in everything from the design, planning, and permitting of

¹June 2002: The Quecreek Coal Mine in Pennsylvania floods leaving 9 miners trapped for 3 days.

²October 2003: The Zapadnaya Coal Mine in southern Russia floods leaving 46 miners trapped for 6 days.

³August 2005: The Daxing Coal Mine in southern China floods killing 123 miners.




Figure 1. Examples of various subterranean environments.

mines, sewer lines, and surface structures, to the active development and environmental upkeep of these establishments. Documentation of older voids, in the form of maps and surveys (see Figure 2), is often inaccurate, illegible, or nonexistent. Correcting this situation necessitates technologies that can obtain reliable information and motivates the development of subterranean robotics.

The current methods available for mapping and surveying underground spaces⁴ can be grouped into three categories:

- **Direct Observation** places people into a void for first-hand inspection.
- **Borehole Observation** probes for subterranean voids by drilling a lattice of boreholes and deploying down-hole cameras to observe void existence [see Figure 11(b) later in

this paper] (Wells, 1999; Morris, Kurth, Huber, Whittaker & Thayer, 2003).

- **Remote Sensing** employs nonintrusive, geophysical procedures that utilize electromagnetic waves and assumed soil composition to locate voids without drilling. Techniques include airborne multi-spectral imaging (Ellis & Dodd, 2000), resistivity (Reynolds, 1997), ground-penetrating radar, and seismic refraction/reflection (Kearey, 2002).

Of these methods, direct observation is unmatched in terms of the fidelity and sampling density of data acquired. Borehole observation and remote sensing can procure satisfactory results in certain circumstances; however, these methods only supply indirect evidence of void and rely heavily on inference (Morris *et al.*, 2003). Direct observation achieves irrefutable first-hand perspective on the state and geometry of underground voids, and requires little to no inference or interpretation.

Subterranean robots offer a unique solution to

⁴For information on current mapping and mining methodologies, see the Canadian Institute of Mining, Metallurgy, and Petroleum Journal at www.cim.org.

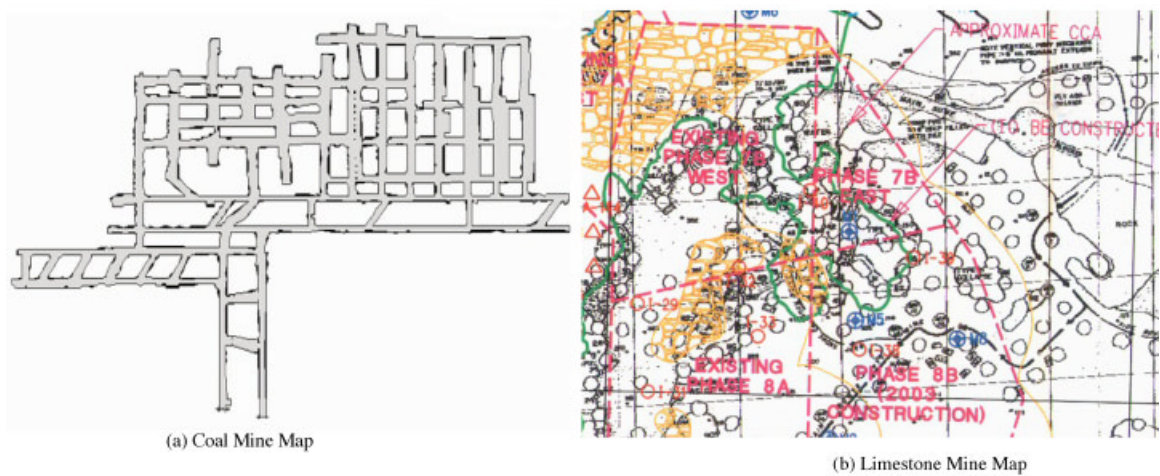


Figure 2. Example maps for various subterranean environments.

void mapping and profiling. Robots can establish a presence in distant subterranean spaces without risk to humans. Robots function in the harshest of environments: both mechanisms and sensing can be tailored to contend with water, smoke, darkness, explosive gas, and countless other hostile conditions. Robots can be compactly constructed to reach remote and space-constrained destinations that are inaccessible to humans. Furthermore, robots can document and recall all sensory input throughout the duration of a void inspection. Data gathered by such means is rich in content and excellent for analysis of void conditions.

Despite the compelling case for subterranean robots, relatively little attention has been paid to the subterranean realm by the robotics community. As a field, robotics has extensively researched applications on land, in water, in air, and in space; yet only a few researchers have concentrated on developing robots for underground operations. Of those that have, most have addressed automating underground mining machines and not mapping voids.

As an example, robotic systems such as autonomous guidance vehicles (AGVs)⁵ and CSIRO's load-haul-dump (LHD) truck (Scheding, Dissanayake, Nebot & Durrant-Whyte, 1997; Roberts *et al.*, 2000), have been designed to navigate active mine corridors for the purpose of hauling extracted minerals and

equipment. AGV systems are guided by a support infrastructure embedded in the corridor walls during construction of the mine. The AGVs navigate by following a predetermined series of infrastructure landmarks. The LHD system also moves unmanned through a mine, but requires no modification to the environment for localization. Navigation is accomplished using a topological map⁶ and information gathered from prior runs to guide the vehicle from corridor to corridor (Roberts, Duff & Corke, 2002; Duff & Roberts, 2003; Duff, Roberts & Corke, 2003).

Mineral extraction platforms, such as continuous miners,⁷ are also being automated for the purpose of safe and efficient coal mining. These systems require no embedded support infrastructure. The precise position and orientation measurements necessary for alignment during mining are obtained from onboard sensors that read off-board survey lasers (Stentz *et al.*, 1999).

The first robot to capture map data from a mine environment was "Terregator," a six-wheeled multi-purpose research vehicle (Champeny-Bares, Copper-smith & Dowling, 1991). Terregator achieved semiautonomous navigation using a combination of sonar and laser scanners for position estimation and ob-

⁵See Automated Mining Systems (<http://www.robominer.com/guidance.html>).

⁶A topological map is a graph representation where the nodes and edges of said graph correspond to distinct locations in the environment (Kuipers & Byan, 1991).

⁷See (<http://www.frc.ri.cmu.edu/~axs/miner.html>).




Figure 3. Different means of access to abandoned mines place very different constraints on robot design.

stacle avoidance. These sensors also enabled Terregator to acquire high-resolution scans of coal mine interiors.

In nonmaintained or natural environments, SANDIA's cave robot, RATLER,⁸ and RedZone's sewer robots⁹ navigate within cluttered caves and sewage-filled pipes, respectively, for the purpose of inspection. These systems are teleoperated via communication links, such as tethers or radio signals. Due to the difficulties encountered in nonmaintained environments, these robots do not utilize autonomy as part of operations.

The objective of subterranean robotics is to operate robots in caves, tunnels, sewers, mines, and cavernlike spaces for exploration, mapping, and characterization. This emerging field of study is poised to define itself among the many fields of robotics given the unique challenges and potential benefits of operating robots within underground regions.

2. CHARACTERISTICS OF SUBTERRANEAN SPACES

Nonmaintained or natural subterranean spaces are among the most challenging environments for robotic operations. These spaces are characteristically high in uncertainty, offer little a priori information,

and restrict communication. In response, robots operating in these conditions must exhibit a degree of autonomy, reliability and robustness, beyond what is typically seen in indoor or outdoor systems. This section describes the characteristics common to most subterranean spaces that have significant influence over robot design and autonomy.

2.1. Limited Ingress

The two standard means of accessing subterranean voids are openings and boreholes (see Figure 3). Openings, like the portal of a mine or entrance to a cave, are primary entryways into subterranean spaces and are generally much larger than boreholes. Openings allow the most flexibility in robot design, but are frequently sealed or blocked. Alternatively, boreholes are manmade holes drilled from the surface. Boreholes, especially those with small diameters, are easy to drill and grant entrance to otherwise unreachable areas, but restrict the cross section and payload of a robot.

2.2. Constrained Volumes

Underground voids vary greatly in size: Height and width vary anywhere from a few centimeters to several tens of meters with void lengths spanning kilometers. The height and width dimensions (or cross

⁸See (http://www.ri.cmu.edu/projects/project_361.html).

⁹See (<http://www.redzone.com/index.cfm>).



Figure 4. (a) “Yellow boy,” an iron oxide and sulfate deposit from acidic mine water. (b) An abandoned haulage-way. Abandoned for only a year, a roof beam collapsed bringing with it cabling, pipes, rock and roof bolts.

section) constrain robot size while the length (or extent) of a void influences the required robot mobility and sensor range.

2.3. Water

Most underground voids contain some water [see Figure 4(a)], whether the depth is as shallow as a small puddle or as deep as complete submersion. Voids where the nominal water depth does not exceed a few centimeters are termed “dry,” voids that are fully submerged are termed “wet,” and voids with water depths in between these extremes are termed “mixed.” These categories impact the mobility and sensing configuration of a robot (e.g., dry spaces favor wheels and lasers and wet spaces favor fins and sonar).

2.4. Gases

Dry or mixed environments have the capacity to house pockets of explosive or corrosive gas. In abandoned mines and sewers, for instance, methane gas is commonly encountered and may ignite if exposed to spark or open flame. Electric motors and open electronics are possible sources of ignition; therefore, explosive threats must be avoided through early detection by environmental sensing.

2.5. Permissions

When working in hazardous spaces, such as those that threaten explosion, safety is as much a concern

for the robot as for the establishments near or above the void [see Figure 5(b)]. Mandated safety regulations, which vary by jurisdiction, dictate the permissible electronics, construction materials, and functionality of a subterranean robot.

2.6. Debris Content


Dense obstacle distributions, such as the aggregate buildup of a sewer line or the fallen beams, rock, and forgotten mining artifacts of an abandoned mine, require a robot to get around, go over, or push through a myriad of obstacles [see Figure 4(b)]. These conditions impose perception, mobility, and navigation challenges upon a robot.

2.7. Lack of Illumination

Subterranean voids are generally dark. Robots must therefore provide their own illumination or employ active sensors such as lasers or sonar.

2.8. Communication

Communication with a subterranean robot is established by either tether or radio; however, these technologies are limited underground. Tethers hinder maneuverability and easily catch or become damaged in the debris of subterranean voids. High-frequency waves cannot penetrate rock, limiting radio communication to areas with line of sight between the transceivers. In general, a global posi-



(a) Mine Map



(b) Mine Fire Damage

Figure 5. (a) This map demonstrates the highly cyclic and symmetric structure of room-and-pillar mines. (b) An example of surface damage caused by a mine fire.

tioning system (GPS), remote commands, or distress signals are not feasible underground. Subterranean robots must, therefore, either: (1) Remain within signal range of the command station, or (2) possess reliable autonomy with robust failure detection and recovery mechanisms.

2.9. Void Structure

The more cyclic and uniform the structure of a subterranean void, the greater the challenge a robot has localizing within the void. Highly cyclic spaces, like that of room-and-pillar mines and caves [see Figure 5(a)], can quickly disorient a robot. As such, robust localization is a critical requirement for mobile subterranean robots.

3. ROBOTIC PLATFORMS FOR SUBTERRANEAN SPACES

The spectrum of conditions encountered in subterranean spaces calls for a variety of robotic systems in response. Depending upon the target environment, a robot may need to scale rock, tread mud, sink, or swim to collect information. No single robot design could be applicable to every conceivable subterranean space; therefore, specialization is essential in subterranean robot design.

The ensuing robot descriptions will be classified along three dimensions: Mobility, void accessibility, and submersion. Mobility covers actuation characteristics; void accessibility covers ingress capabilities; and submersion covers the ability to operate under water. Figure 6 shows a diagram of this classification space.

Each of the following robotic systems represents a particular region in this space of subterranean robotic configurations. These regions consist of *static* and *mobile* along the mobility axis, *borehole* and *portal* along the void accessibility axis, and *dry* and *wet* along the submersion axis. In total, eight configurations encompass the subterranean robots that are either in use or under development. Two of these regions, Static-Portal Dry and Static-Portal Wet, embody the simple sensing platforms utilized on all subterranean robotic systems, and are therefore not discussed explicitly.

3.1. Borehole Deployable Lasers: Static-Borehole Dry

Borehole deployable lasers (BDLs) are simple yet effective tools for acquiring data from subterranean spaces (see Figure 7). BDLs consist of a primary laser-range measurement sensor, a set of actuators that aim the sensor, inertial sensors that provide absolute device pose, and cameras or side-mounted




Figure 6. Varying constraints in terms of void accessibility, mobility, and submersion result in drastically different robot designs.

proximity sensors that assist during deployment. These devices are compact (currently less than 15.2 cm or 6 in. in diameter) to allow for deployment into boreholes or pipes. BDLs are characteristically tethered and offer linear mobility. Deployment is achieved by lowering/raising or pushing/pulling the device to a fixed position.

BDLs are controlled remotely by an operator who monitors and modifies scan progress. Direct operator control allows for quick operational changes and in-field analysis. The straightforward mechanical designs of BDLs decrease the likelihood of failure and permit in-field repairs should an incident occur. In addition, experience has shown these systems to be extremely manageable, requiring only small teams, typically two people, to perform field operations.

Figure 8(a) shows a standard deployment and scanning scenario. This procedure entails:

1. An existing or newly created borehole is selected for deployment. Newly created boreholes are usually lined with pipes to reduce

device contact with mud, to prevent borehole collapse, and to ensure the safe recovery of the device.

2. The device is lowered via a winch-and-pulley system from the surface into the void space. Cameras or side-mounted proximity sensors are used to identify entry into the void breach (i.e., the point where a borehole opens into the subterranean void).
3. Once in the void, the range measurement sensor is actuated from a stationary vantage point to collect range data from all void surfaces. At this point, the device operator may alter the speed of the sensor actuators to increase or decrease the density of scan data.
4. Upon completion of a scan, the device may be raised or lowered to other vantages where additional scans are taken. Scan depths are documented from the surface with respect to the top of the borehole.
5. Lastly, the device is recovered and survey measurements (i.e., GPS data on borehole location, relative measurement from other



Figure 7. Borehole deployable lasers developed to scan subterranean voids.

landmarks, etc.) are recorded. These measurements are used in conjunction with log data for analysis during postprocessing.

The data recovered from BDL scans serve a wide variety of uses. Raw range data taken directly from the range sensor is converted into a three-dimensional (3D) point cloud, as seen in Figures 9(a) and 13(b). Point clouds provide insight into the geometrical characteristics of a void. When a mesh is applied to a point cloud as in Figure 8(b), these geometric shapes can be interpreted as structural features, which provide clues to the history and integrity of the void.

Another property calculated from high-fidelity 3D range data is an estimate of volume. Volume can provide an assessment for the amount of backfill material required to fill a subterranean void or the amount of buildup that has collected and must be removed. Volume can even indicate the structural status of a void by providing a metric to calculate the likelihood of collapse and to what extent the collapse damage will affect the surface.

Scan data is also useful for registering and correcting existing subterranean maps. Templates are constructed from 3D point cloud data by extracting cross sections parallel to the existing map plane (see Figure 10). Templates are correlated to existing map

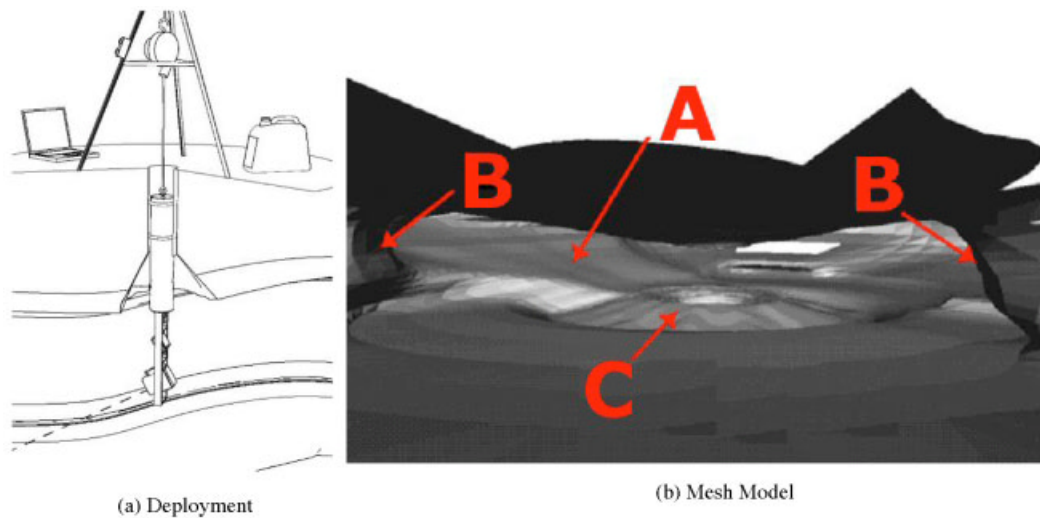


Figure 8. (a) The deployment process of BDLs. (b) A rendering generated from laser data of a partially backfilled limestone mine 45 m below the surface. Visual features in this rendering include (A) a slope of hardened backfill material, (B) two mine columns, and (C) a debris pile formed during the borehole drilling process.

regions and locked into place using template matching techniques described in Olson (2000). The original map can then be corrected to reflect any measured changes or inaccuracies. When multiple templates are available, matches can be established over different map regions and correlated with borehole GPS data. Subterranean maps aligned to surface coordinates have tremendous value when drilling additional boreholes or blocking-off hazardous locations on the surface.

The correlation of void geometry with surface coordinates has also led to the development of a technique known as “frontier drilling.” Frontier drilling gives BDLs a more active role in the drilling process. A major problem with borehole drilling is that many boreholes miss the intended void [see Figure 11(b)]. Surface-linked subterranean maps can optimize the placement of additional boreholes, saving time, money, and manpower by reducing the number of required boreholes. Once a subterranean access is established, BDLs can direct the remaining boreholes to be drilled over known areas of void location. Figure 11 gives an example of this method.


BDLs have proven invaluable for data collection and quick void analysis. The current platform, known as Ferret III [see Figure 7(b)], has been deployed into both limestone and coal environments. The onboard laser averages 70 readings per second,

allowing reasonably high-resolution 3D scans (approximately 0.5° azimuth \times 0.5° elevation) to be accrued in 80 min. Example results are shown in Figure 9.


Ferret III’s minimum angular resolution is around 0.1° in both elevation and azimuth. Range measurements are calibrated to 65 m on rock surfaces with a ± 10 mm error. Deployment depth is logged to within 30 mm from a survey tape directly linking the instrument to the surface. To date, all deployments have been less than 100 m deep and GPS-surveyed surface coordinates have established surface deployment locations.

One common assumption in all BDL deployments is that boreholes are vertically straight from the surface position to the void breach. In reality, horizontal drift causes displacement relative to surface coordinates. At long drill depths (i.e., around and beyond 300 m), drift may become significant. Drilling companies have methods for recovering drift measurements; however, past field deployments have not yet reached such depths and drift has not been a significant problem.

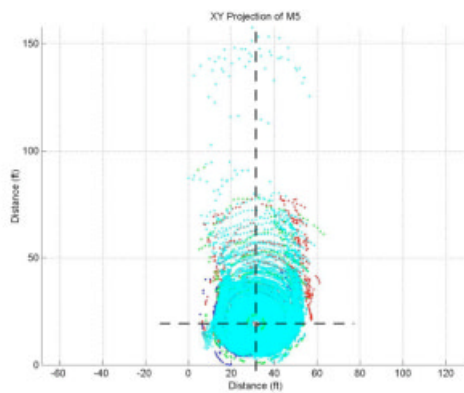
During BDL deployment, orientation recovery is a more important concern than borehole drift. Absolute device heading is necessary to properly align




(a) Point Cloud



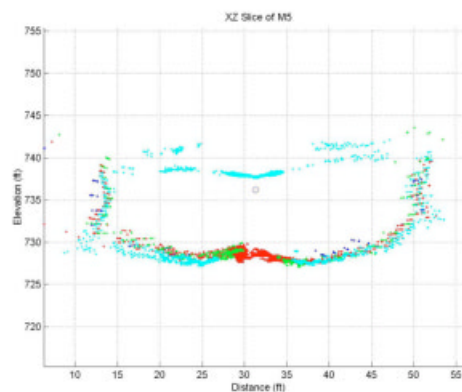
(b) Mesh Model



(c) X-Y Plot



(d) Y-Z Plot



(e) X-Z Plot

Figure 9. A series of plots showing a limestone mine void that was originally believed to have been backfilled several years prior. The void was scanned with a BDL to reveal an open space approximately 16 m wide, 15 m long, and 7.5 m tall. With a discrepancy of this degree, a further investigation was launched to verify other believed backfill locations.

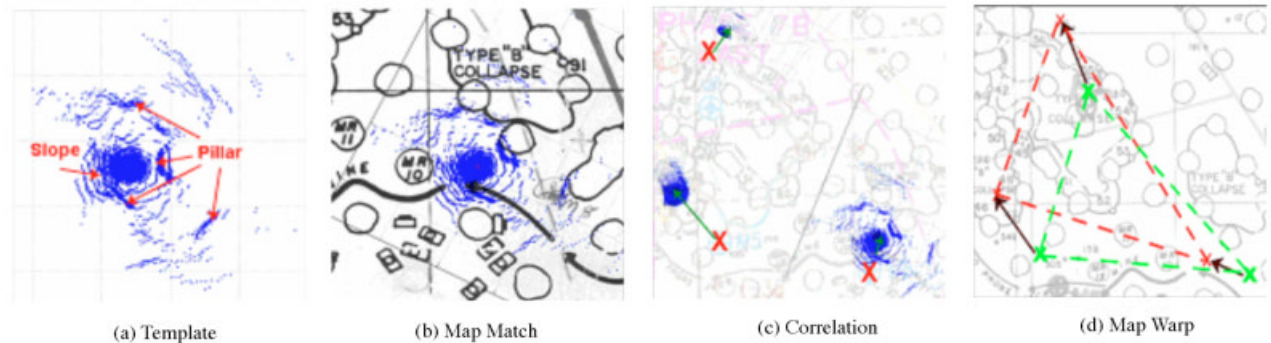


Figure 10. The process of map correction from template matching. First, a point cloud is compressed into a plane to provide a baseline template for template-to-map matching. Next, the template is matched to features in the existing map. Some interpretation may be required due to discrepancies between the map landmarks and scan data. Next, multiple scans are locked to a single map with GPS information to synchronize scan coordinates with surface coordinates. Finally, the map is warped to correlate with scan and GPS data.

scan data. Orientation recovery is accomplished through mechanical alignment, sensing, scan matching, or map matching.

- Mechanical alignment uses interconnecting tubing to transfer the orientation of the BDL

to the surface where the heading is referenced. This method is extremely useful and reliable, but limited to shallow depths.

- Sensing uses a compass and/or gyroscope to establish orientation. Sensing is useful pro-

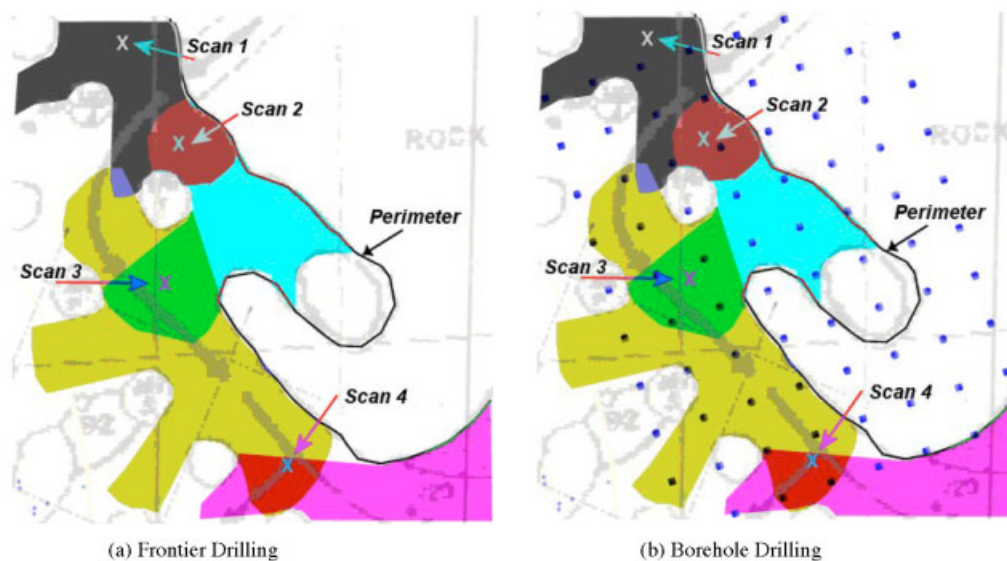


Figure 11. (a) The process of frontier drilling. To establish the perimeter of this mine, an existing borehole is used to take Scan 1. The borehole for Scan 2 is then drilled by correlating Scan 1 data with surface coordinates. The scan data from each successive scan thereafter are used to place the next borehole, until the entire perimeter of the mine has been established. (b) The number of boreholes that may be needed to establish a uniform distribution of ground sampling to cover the same amount of area as frontier drilling.




Figure 12. A borehole deployable sonar unit.

vided the device does not experience hard impacts during deployment.

- Scan matching recovers heading by correlating features from overlapping scans from different borehole vantages into a common model. This method, however, requires two or more boreholes.
- Map matching involves orienting the model to old map features, such as pillars, rails, and corridors. This method assumes the existing map is sufficiently reliable.

Field work conducted at Carnegie Mellon¹⁰ and the NIOSH research mine near Pittsburgh, PA have enabled verification of these orientation recovery techniques on prototype BDL systems.

The greatest caveat to these systems is that they are highly limited by vantage obstruction. Scans taken from a single location suffer when nearby objects obscure the sensor's view. Careful selection of borehole placement can alleviate this problem, but borehole drilling is quite literally a "hit-or-miss" process. If a system could move around local obstructions, however, occlusions would not be an issue. The use of mobility and the additional challenges created by motion are addressed in Section 3.3.

3.2. Borehole Deployable Sonar: Static-Borehole Wet

A borehole-deployable sonar (BDS) is operationally identical to a BDL with the primary difference being the technology (i.e., sonar over laser) used to acquire range measurements. Comparatively, broad sonar beams are less accurate than laser in dry conditions; however, acoustic range sensors outperform lasers in water since light refracts and reflects when underwater.

¹⁰See www.subterraneanrobotics.org.

BDSs are similar in construction, mechanisms, and inertial sensing to BDLs. The additional conditions that must be addressed for BDSs are water and pressure. BDSs are sealed to prevent leakage from shorting electronics and structurally enhanced to withstand pressure when hundreds of meters below the water's surface.

The deployment process for BDSs is the same as for BDLs. Scans are taken at incremental vantages along the elevation of the void. Scan time varies for BDSs since sonar measurement rates depend on range limitations and scan resolution. For example, the plot seen in Figure 13 consists of 68,872 data points taken over a void with an estimated volume of 4,500 cubic meters. These data are a composite of 20 scans recorded at elevation increments of 0.3 m and acquired in 3.2 h. The same orientation recovery mechanisms employed by BDL are likewise utilized by BDS.

The current BDS prototype, known as Wet Ferret (see Figure 12), has been deployed into a number of submersed environments, including mines, sewers, and drainage pipes. Results from these deployments can be seen in Figure 13.

Obstruction and limited vantage also restrict BDSs in the same ways as BDLs. Furthermore, BDSs are also affected by water flow. Strong flow pushes material onto the sensor head and reduces scan quality. Methods for dealing with flow are currently being investigated.

3.3. Subterranean Mobile Robots: Mobile-Portal Dry

Subterranean Mobile Robots (SMRs) define a class of systems that navigate and sense subterranean spaces both autonomously and under teleoperation. Navigating through cluttered and often hostile subterranean environments necessitates far greater system sophistication than seen in BDLs and BDSs. In ex-




Figure 13. This series of plots was recorded in a zinc mine 52 m below ground and demonstrate the capability of sonar in wet conditions. (a) Sonar data are logged as planar intensity images that encode the strength of an echoed signal. (b) Range points, which reflect the loudest echoes, are recorded at incremental elevations and stacked to generate a point cloud similar to the one created by laser. (c) This point cloud can also be meshed to show wall structure and give a sense of void geometry.

change for added complexity, mobile systems enable data acquisition from multiple vantage points, thus providing more diverse and content-rich datasets than are possible with static systems.

SMRs are typically slow moving with high torque capabilities. High torque allows the robot to surmount more objects while low speed is more amenable to mapping and permits simplifications in autonomy software. For example, the Groundhog platform [seen in Figure 14(a)] is based heavily on components from commercial all-terrain vehicles but has hydraulics motors that provide speeds of 0.15 m/s (1/3 mph) and 0.45 m/s (1 mph). At the lowest speed, Groundhog can easily surmount objects 16 cm in height and climb slopes of 30° under a 725 kg (1600 lbs) payload of electronics, motors, hydraulics, and a steel-reinforced chassis. Slow mo-

tions also allow SMR navigation systems to plan paths while ignoring vehicle dynamics, a simplification not possible for a fast-moving robot of Groundhog's weight.

Several other void characteristics dictate the mechanical and structural design of an SMR. Void cross section and obstacle distributions determine the width, ground clearance, and wheel base of an SMR. Explosive gas motivates the use of hydraulics over electric motors that could spark. Water and mud require waterproof enclosures and protective casings for electronics. Sharp corners and narrow passageways require drive systems that allow for tight maneuvering. Cave Crawler, for example [seen in Figure 14(b)], has four independently steered wheels that allow it to spin in place, which is a useful feature in narrow coal mine corridors.



Figure 14. Groundhog and Cave Crawler, our two subterranean mobile robots.




Figure 15. (Left) A survey map of a section of the Bruceton mine traversed by Groundhog during an autonomous run. Intersections that were encountered by the robot were recorded as nodes in its topological map (marked as circles). (Right) Groundhog's path from online pose estimation during this run. Even though the pose estimates for a particular intersection differ each time it was detected, the topological map provides a mechanism for determining the correct location.

In addition to rugged mechanics, SMRs must also address capable sensing. Sensing is usually specific to a particular subterranean environment. In the case of Groundhog designed for abandoned coal mines, two scanning lasers, one mounted on each end of the platform, provide range measurements for mapping and navigation. Two gas sensors monitor air quality and prevent the robot from entering a potentially explosive pocket of methane gas. Eight flotation switches line the bottom of the robot to detect sinking in water or mud. A gyroscope, encoders, and tilt sensors provide heading, velocity, and tilt information for odometry and rollover prevention.

SMR software must allow the robot to make productive decisions while operating in a subterranean void. One of the greatest demands in SMR software is reliable localization. Limited a priori information and the absence of communication make localization in subterranean environments exceptionally difficult. Classical simultaneous localization and mapping (SLAM) approaches do not perform well in highly cyclic spaces and are computationally expensive (Ferguson *et al.*, 2003). One approach that has proven useful in coal mine environments employs incremental scan matching for short-term position estimation and topological SLAM for global localization (Silver, Ferguson, Morris & Thayer, 2004). Topological SLAM is achieved through construction of a topological map from features extracted inside the environment (such as mine corridor intersections). Figure 15 illustrates a topological map created online

during a traverse in the Bruceton Research Mine where corridor intersections were used to localize the robot.

Many other algorithms are also responsible for the decision-making capabilities of an SMR. These algorithms include exploration, navigation, goal selection, obstacle identification, motion planning, and path following. Currently, navigation is split into a two-tiered hierarchy. First, global navigation, which uses a topological map to choose goals, selects a target area of the environment to explore based on high-level objectives, such as perimeter discovery¹¹ or complete coverage.¹² Second, local navigation utilizes 3D range data to produce local goals and motion plans based on a kinematic model of the robot (Morris, Silver, Ferguson & Thayer, 2005). Figure 16 presents a sample binary traversability map generated from a mine intersection, along with a generated path from an initial configuration to a goal configuration. This navigation scheme has proven successful in test-mine conditions; however, much more work will be needed to provide system reliability.

Although high reliability resides as future work for these systems, the first steps toward fault recovery have already been examined. Some SMRs em-

¹¹A planned series of goal points that will move the robot along the perimeter of the subterranean void.

¹²A planned series of goal points that guarantees all unexplored sections of void will be explored.




Figure 16. (Left) A 3D point cloud generated at a corridor intersection in the Bruceton Mine. (Center) A binary traversability map of the intersection. Also shown are the robot's initial configuration (1) and its goal configuration (2). (Right) The path planned from the initial configuration to the goal configuration, taking into account steering constraints.

ploy a hierarchy of degraded operation modes that are triggered by changes in system status (Baker et al., 2004). Errors in software modules or hardware components trigger preprogrammed behaviors intended to drive the robot back to a location where humans can recover the system.

Preliminary results from SMRs demonstrate tremendous potential for automated subterranean mapping. Groundhog, for example, has been used for exploring and mapping both research coal mines and abandoned mines. From May until November of 2003, Groundhog performed a series of experiments in the abandoned Mathies mine outside of Pittsburgh, PA. Originally used as a coal haulageway, Mathies's 3 km long central corridor was under investigation by the PA-DEP¹³ for use as a possible pipeline route (see Figure 17). To provide information regarding the structural integrity of the main haulageway, Groundhog was deployed to gather data and produce an updated mine map.

Table I summarizes the Mathies campaign. In total, 800 m of the possible 2100 m of mine corridor

were mapped during eight different deployments. The robot also autonomously traversed in excess of 2 km while in the process. Some of the results obtained during these eight deployments are displayed in Figure 18, showing the resultant two-dimensional (2D) maps and 3D mine scans.

In addition to exploration of isolated corridors, experiments involving more general exploration were conducted in the Bruceton research coal mine outside of Pittsburgh, PA. Each experiment consisted of an autonomous exploration of an acyclic environment. For each exploration, a topological map was constructed, along with a log of detected intersections. These experiments were designed to verify topological exploration where robot position and goals are determined by detection of corridor intersections.

Figure 15 shows the topological map built during the largest single experiment. This topological map is superimposed on a survey map of the same corresponding section of the mine. In this run, Groundhog was teleoperated from the portal to a position near the first intersection. Groundhog then

¹³Pennsylvania Department of Environmental Protection.



Figure 17. The prior survey map of Mathies Mine, with the three portals labeled as (1), (2) and (3). The corridors of interest run horizontally from left to right. Note the inaccuracy in the map on the right side (the corridors are misaligned). The image represents approximately 1200 meters from left to right.


Table I. Summary of field deployments of Groundhog into the Mathies Mine during May and October 2003 (see Figure 17 for port references).

Mission	Date	Port	Goal (m)	Comp (m)	Return caused by	Duration (min)	Egress	Comments
1	05/30	1	500	308	Roof fall	155	No	Hardware error
2	10/01	2	100	100	Mission Complete	48	Yes	Success
3	10/01	1	100	100	Mission Complete	43	Yes	Success
4	10/01	3	100	60	Submergence	30	Yes	Stuck in mud
5	10/08	2	500	140	Roof fall	81	Yes	Success
6	10/22	3	100	20	Software failure	20	Yes	Navigation error
7	10/22	3	100	10	Software failure	9	Yes	Navigation error
8	10/30	3	330	230	Fallen cable	140	Yes	Success

proceeded to autonomously explore as much of the environment as was traversable and autonomously returned to the portal. Over the course of 2 hours, Groundhog traversed more than 400 m of mine corridor.

At the time of these experiments, Groundhog was unable to recognize cycles in the environment (known as “closing the loop”) and so was limited to acyclic sections of mines. Known generally as the loop closure problem, Groundhog, unable to recognize previously visited locations, would endlessly traverse a cycle believing it was exploring unseen corridors. Recent work (Silver, Carsten & Thayer, 2005) has demonstrated the feasibility of using local 3D information registered to topological features for solving this problem.

While great strides in system development have been made in Subterranean Mobile Robotics, there is much more work to be done. The tremendous amount of uncertainty in the subterranean world require robustness and reliability that is still being worked towards, but not yet achieved.

**Figure 18.** Results from the Mathies Mine: The 2D maps are approximately scaled and aligned to match the orientation in Figure 17. The 3D scans are, from left to right: (A) The roof-fall encountered 140 m into Portal 2, (B) the fallen timber encountered 308 m into Portal 1, and (C) the fork in the corridor encountered 200 m into Portal 3.

3.4. Submersible Subterranean Robots: Mobile-Portal and Borehole Wet

Submersible subterranean robots represent a class of mobile robots capable of mapping liquid filled subterranean spaces. The operational challenges of traversing and mapping wet voids are different than those of dry voids. For example, water permits motion in a higher-dimensional space. Objects that would be obstacles to a vehicle constrained to crawling on the floor are avoidable in this domain. In addition, buoyant systems consume less power when moving compared to locomotion over rough terrain for a dry system. This enables buoyant or swimming robots to traverse longer distances than dry mobiles with the same specific energy.¹⁴ The added mobility and maneuverability of submersible systems, however, place higher demands on localization and path planning, especially in the presence of water flow.

Commercial submersible robots have already addressed many of the platform design issues necessary for a subterranean system (Whitcomb, 2000). These vehicles are currently in operation for sea-based tasks such as commercial missions, oceanographic research, and military excursions. Remotely operated vehicles are tethered platforms that come laden with sensors, manipulators, and video equipment for remote inspection and scientific research. Autonomous underwater vehicles (or AUVs) are long-range fast moving submersibles designed to map coast lines, discover sunken equipment, and patrol large offshore areas. Both platform types are constructed for work in open sea; therefore, size con-

¹⁴The ratio of the energy output of a robot's power source to its weight.

straints imposed by subterranean spaces require modification of these commercial system designs.

In the confined spaces of subterranean voids, submersible robots favor designs that allow tight cornering and slow speed navigation. So-called “fast” subs rely on propellers for propulsion and dive planes for orientation control. This coupling between propulsion and orientation control results in systems with large turning radii. Therefore, for subterranean voids thruster or impeller-based platforms are favored for their ability to turn in place and maintain position and orientation, even when fighting currents.

Deployment size constraints also restrict the construction of submersible subterranean robots. The trade-offs between borehole and portal entries are largely unchanged between wet and dry systems. Portal systems can be larger with greater on-board energy storage, improved sensing, and more powerful computing; however, navigation and void accessibility prefer small submersible robots to bulkier systems. As a result, development of submersibles has and will continue to focus on borehole deployable systems.

No longer constrained to crawl on the floor, navigation and mapping in a 3D space requires recovery of the vehicle’s pose in six dimensions: roll, pitch, yaw, and three translational components. For thruster-based submersibles, these six dimensions can be reduced to four (yaw and three translation axes) by restricting roll and pitch motions through proper buoyancy design. In addition, the translation axes corresponding to depth can be explicitly measured to within centimeters by conductivity, temperature, and depth sensors. This further reduces the robots localization domain to three dimensions.

Recovery of yaw and the two translation axes still remains a difficult problem for autonomous or semiautonomous submersibles. A standard approach to underwater robotic positioning and navigation is long baseline (LBL) acoustic navigation (Hunt *et al.*, 1974), which uses time of flight to fixed transponders. For LBL to work in the subterranean domain, transponders would have to be placed in such a manner as to ensure a continual line of sight with the robot, which would require multiple boreholes and knowledge of local geometric structure. For this reason, LBL is not a practical solution for subterranean spaces.

Doppler-based velocity estimation (Whitcomb, Yoeger & Singh, 1999) is another approach coming

into use in underwater vehicles. Unfortunately, the necessary hardware is often both large and heavy. While recent advances have resulted in systems small enough to be built into a borehole deployable robot, their weight and cost are still prohibitive at this time.

Sensor-based localization for underwater systems, which rely on sonar for range sensing, is complicated by the low angular resolution of sonar versus laser. However, with roll, pitch, and depth directly measurable, the localization problem is reduced to a planar problem.

Several algorithms have been developed recently to improve sonar-based underwater horizontal localization. One such method for robust 2D scan matching with sparse and noisy range data is presented in Silver, Bradley & Thayer (2004). This method was validated for subterranean environments by using laser data from several mines which have been subsampled and noise corrupted. The resulting data, when combined with a rotating field of view, are believed to simulate an underwater profiling sonar.

Any local scan-matching system, however, will accumulate error over time and requires a method to close loops in order to maintain robust positioning. Figure 19 shows a sonar-based feature for global localization in flooded environments known as a slide image (Bradley, Silver & Thayer, 2004). This feature was designed to minimize sensitivity to accumulated positioning errors and allow for reliable topological localization in flooded subterranean tunnels. This method has been tested on sonar data recorded from an underground tunnel system in Wakulla Springs, Florida (Stone, Ende, Wefer & Jones, 2000). An occupancy grid approach to underwater sonar SLAM, along with an analysis of the localization performance of several sonar configurations, is presented in Fairfield, Kantor & Wettergreen (2005).

Currently, development of a submersible unit for exploration of submerged mines is underway. Figure 20 shows a conceptual rendering of a borehole-deployable submersible known as “Mine Fish.” When complete, the underwater localization and mapping methods of Silver, Bradley & Thayer (2004) and Bradley *et al.* (2004) will enable Mine Fish to navigate the extents of abandoned mines and produce maps that would otherwise be unobtainable.

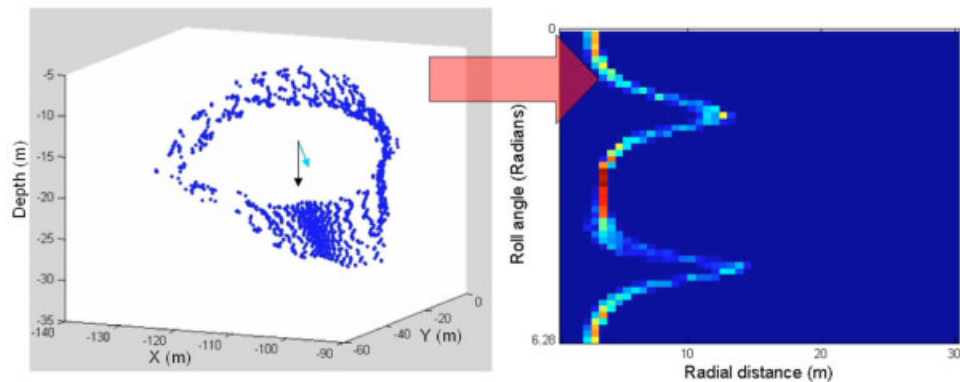


Figure 19. Slide images create a descriptor for tunnel-like environments that is invariant to accumulated errors in positioning in the horizontal plane. A coordinate system for the most recent sensor data is defined by the gravity vector (dark black arrow) and an estimated tunnel axis vector (light arrow). The points are then histogrammed by range and angle around the tunnel axis to create a compact feature vector.




Figure 20. Conceptual rendering of Mine Fish.

Table II. Mobile borehole robot concepts.

Concept	Deployment	Drive	Steer	Scan
Two-wheel gravity reaction differential drive	Inflatable wheels	Two-wheel drive	Differential steering	Serpentine path, fixed line scanner
Four wheel explicit steer inflatable	Inflatable wheels	Four-wheel drive	Hinged center link	Tilting line scanner
Linked independent differential drives	Mechanically deployed wheels	Four-wheel drive	Coordinated differential steering	Serpentine path, fixed line scanner
Two roller explicit steer	Mechanical reconfiguration	Two drive rolls	Explicit steering (each roll)	Tilting line scanner

3.5. Borehole-Deployable Mobile Robots: Mobile-Borehole Dry

Borehole-deployable Mobile Robots bridge the gap between SMRs and BDLs. Small enough to fit within a borehole or pipe, these systems eliminate the problem of limited vantage by providing mobility to see around obstructions.

Four critical abilities that differentiate Borehole-deployable Mobile Robots from BDLs and SMRs are:

1. Self-reconfiguration from borehole compatible to subterranean void compatible,
2. Linear locomotion sufficient to overcome expected void obstacles,
3. Steering consistent with navigation in narrow passages, and
4. Scanning adequate to produce globally consistent 3D datasets.

From these requirements a broad range of configuration concepts is possible, with four representative concepts outlined in Table II and Figure 21.

Figure 22 shows a transforming borehole-deployable mobile platform chosen for development from the four concepts shown above. When stowed, the robot has a cylindrical shape compatible with borehole deployment. Once deployed, it drives on two-drive cylinders, much like a steamroller, that are




Figure 21. Stowed and deployed configurations for four borehole deployable mobile robot configurations. From top left: Inflatable differential drive with laser scan path shown; inflatable four-wheel explicit steer concept, with base station; linked differential drive concept with mechanically deployed wheels; two roller helical drive/steer concept.




Figure 22. (Left) Photograph of Helix, a mobile subterranean robot that can transform for deployment through boreholes. (Right) A diagram of Helix as it unfolds from a deployment configuration to a mobile configuration.

articulated in the plane of motion to steer. A fixed laser-range finder reflects off a spinning 45° mirror, providing a 360° planar scan. The center section of the robot, which contains the laser and mirror assembly, pivots to provide 3D laser scanning capability. While the system described does not provide the simplicity of static borehole sensors, or the long range of mobile portal systems, it increases the quantity and quality of information obtainable from a single borehole.

4. SUMMARY

In Section 1, subterranean maps were introduced as crucial documents for prevention of disasters and for planning of surface and mining operations. Existing maps, however, were found to be of unacceptable quality for delivering reliable information needs. This motivated the need for subterranean mapping methods. Section 2 went on to characterize subterranean spaces as highly uncertain and treacherous environments. The tremendous challenges humans and machines face when attempting to inhabit a subterranean space for exploration, mapping, or profiling were presented. Section 3 then presented a spectrum

of systems developed to achieve the goals of Section 1 when faced with the environmental constraints of Section 2. This section examines the question, “How successful are these solutions to addressing the problem of subterranean mapping?”

Table III is an overall comparison of the systems described in Section 3. Included in this comparison are human surveys, geophysical, and borehole methods. Please note that the geophysical category encompasses a variety of methods that are generalized to fit within the context of this comparison. This table distinguishes the particular strengths and weaknesses of each methodology.

The attributes of Table III are as follows.

- **Entry.** Means of void access: P (portal), B (borehole), and S (surface).
- **Water.** Works in water: Y (yes) and N (no).
- **Drilling.** Relative number of boreholes needed: Arranged by magnitude: Most, Fewer, and Fewest.
- **Depth.** Usable void depths: Shallow (ground level), Deep (hundreds of feet), and Very Deep (thousands of feet).
- **MTBF.** Mean time between failures: Short

Table III. Comparison of subterranean mapping methodologies.

Method	Entry	Water	Drilling	Depth	MTBF	Density	Accuracy
Human Survey	P	N	—	Shallow	Very Long	Sparse	High
Geophysical	S	Y	—	Deep	Very Long	Dense	Med
Borehole	S	Y	Most	Very Deep	Very Long	Sparse	Low
BDL	B	N	Fewer	Very Deep	Long	Dense	High
BDS	B	Y	Fewer	Very Deep	Long	Dense	Med-high
Mobile	P	N	—	Shallow	Short	Very Dense	High
Submersible	B	Y	Fewest	Very Deep	Short	Very Dense	Med-high
Reconfigurable	B	N	Fewest	Very Deep	Short	Very Dense	High

(frequently needs repair), Long (performs multiple operations without repair), and Very Long (rarely needs repair over entire lifetime).

- **Density.** Quantity of data recovered: Sparse (very little quantitative information), Dense (rich amount of data), and Very Dense (vast amounts of data).
- **Accuracy.** Reliability of data: Low (very ambiguous or not informative), Med (have some ambiguous properties), Med-high (has ambiguous properties that can be statistically resolved), and High (has very little ambiguity).

This table provides a general interpretation of subterranean mapping methods. First, human surveying methods are among the most resource intensive with fewest results. A human requires both portal entry and boreholes for surface-to-void correlation and is restricted to shallow deployments. Where humans are superior is their ability to cope with unexpected incidents. Geophysical methods prove to be powerful mapping tools; however, they are not applicable to all domains (e.g., sewers) and are adversely effected by increasing depth. In some cases, acquired data are subject to so much ambiguity that alternative methods are needed for verification.

Among the robotic systems, BDLs and BDSs are the most applicable for immediately addressing the subterranean mapping problem. Together, they can venture into a myriad of subterranean conditions with a high likelihood of survival. In addition, they offer much richer data than surveying, geophysical, and borehole methods with very little ambiguity. SMRs are systems currently under development, but have proven to excel at mapping when portals are available. The most promising of these systems, however, are submersibles and borehole-deployable mo-

bile robots. These systems have the deployment versatility of BDLs and BDSs with the capacity to move. Movement from a borehole deployment will increase the quantity of data acquired per borehole and allow these methods to become the most economic for subterranean mapping.

The most likely solution to the subterranean mapping problem will not be a single approach, but a combination of techniques. For example, in a mine mapping operation, geophysical methods could be deployed to establish a rough idea of void location. Boreholes can then be drilled, which would allow a BDL or BDS to verify and map discovered void. In later scenarios, the autonomy methods developed on SMRs will be adapted for submersibles and borehole-deployable mobile robots. These systems could then be able to augment or replace the BDLs and BDSs altogether.

5. CONCLUSIONS

Subterranean robots are transforming the exploration and mapping of underground voids and revolutionizing operations in sewers, mines, and caves. These robots are becoming competent agents for entering, viewing, and exploring voids as they evolve for future challenges, such as fire fighting, rescue, reclamation, and active mining operations.

This research chronicles the basics of void entry, mobility, sensing, navigation, modeling, and subterranean operations. In this work, subterranean robots have demonstrated huge advantages for moving through underground voids that preclude human access. Superb void models have been created from voluminous range data. These models are valuable in

their own right as void maps, but have also shown even greater value for machine guidance and operational autonomy.

This research is ongoing. Development is active in mechanisms, sensing, and software. Future machines will become smaller, more reliable, and incorporate new sensing configurations, such as thermal imaging, tactile, and pressure sensing. New classification algorithms, SLAM, and navigation methods will emerge to yield more capable autonomous behavior. Together, these components will coalesce into a bold future for subterranean robots.

ACKNOWLEDGMENTS

Special thanks to previous members of the subterranean robotics group at CMU, particularly Dirk Hähnel, Carlos Reverte, Mike Montemerlo, and Sebastian Thrun, and students of the Mobile Robot Development class at CMU who helped build Groundhog. Thanks are also given to the Bruceton Research Mine (Paul Stefko), the National Institute of Occupational Health and Safety, Mine Safety, and Health Administration, Pennsylvania Department of Environmental Protection, and the various people in the mining industry who provided guidance and supported this work. Many thanks go to John Hatch, owner of the Mathies site of Mon View Mining, for granting us access to his mine as well as Workhorse Technologies, LLC, for technical and operational support in all aspects of the project.

REFERENCES

- Baker, C., Morris, A., Ferguson, D., Thayer, S., Whittaker, C., Omohundro, Z., Reverte, C., Whittaker, W., Hähnel, D., & Thrun, S. (2004). A campaign in autonomous mine mapping. Paper presented at the IEEE International Conference on Robotics and Automation (ICRA), New Orleans, LA, pp. 2004–2009.
- Belwood, J., & Waugh, R. (1991). Bats and mines: Abandoned does not always mean empty. *Bats*, 9(3), 13–16.
- Bradley, D., Silver, D., & Thayer, S. (2004). A regional point descriptor for localization in subterranean environments. Paper presented at the IEEE Conference on Robotics Automation and Mechatronics (RAM 2005), pp. 440–445.
- Champany-Bares, L., Coppersmith, L., & Dowling, K. (1991). The terregator mobile robot (Tech. Rep. CMU-RI-TR-93-03). Pittsburgh, PA: Carnegie Mellon University, Robotics Institute.
- Duff, E., & Roberts, J. (2003). Wall following with constrained active contours. Paper presented at the Fourth International Conference on Field and Service Robotics.
- Duff, E., Roberts, J., & Corke, P. (2003). Automation of an underground mining vehicle using reactive navigation and opportunistic localization. Paper presented at the IEEE/RSJ International Conference on Intelligent Robots and Systems, pp. 3775–3780.
- Ellis, J., & Dodd, H. (2000). Applications and lessons learned with airborne multispectral imaging. Paper presented at the 14th International Conference on Applied Geologic Remote Sensing.
- Fairfield, N., Kantor, G., & Wettergreen, D. (2005). Three-dimensional evidence grids for slam in complex underwater environments. Paper presented at the 14th International Symposium of Unmanned Untethered Submersible Technology (UUST), Lee, NH: AUSI.
- Ferguson, D., Morris, A., Hähnel, D., Baker, C., Omohundro, Z., Reverte, C., Thayer, S., Whittaker, C., & Whittaker, W. (2003). An autonomous robotic system for mapping abandoned mines. Paper presented at the Conference on Neural Information Processing Systems (NIPS).
- Hunt, M.M., Marquet, W.M., Moller, D.A., Peal, K.R., Smith, W.K., & Spindell, R.C. (1974). An acoustic navigation system (Tech Rep. WHOI-74-6). Woods Hole, MA: Woods Hole Oceanographic Institution.
- Kearey, P. (2002). An introduction to geophysical exploration, 3rd ed. Oxford, UK: Blackwell Science Ltd.
- Kuipers, B., & Byan, Y. (1991). A robot exploration and mapping strategy based on a semantic hierarchy of spatial representations. *J Robot Auton Syst*, 8, 47–63.
- Morris, A., Silver, D., Ferguson, D., & Thayer, S. (2005). Toward topological exploration of abandoned mines. Paper presented at the IEEE International Conference on Robotics and Automation.
- Morris, A., Kurth, D., Huber, D., Whittaker, C., & Thayer, S. (2003). Case studies of a borehole deployable robot for limestone profiling and mapping. Paper presented at the International Conference on Field and Service Robotics (FSR).
- Olson, C.F. (2000). Maximum-likelihood template matching. Paper presented at the 2000 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR), Vol. 2.
- Reynolds, J. (1997). An introduction to applied and environmental physics. New York: Wiley.
- Roberts, J., Duff, E., & Corke, P. (2002). Reactive navigation and opportunistic localization for autonomous underground mining vehicles. *J Inf Sci*, 145(1–2), 127–146.
- Roberts, J., Duff, E., Corke, P., Sikka, P., Winstanley, G., & Cunningham, J. (2000). Autonomous control of underground mining vehicles using reactive navigation. Paper presented at the IEEE International Conference on Robotics and Automation.
- Scheding, S., Dissanayake, G., Nebot, E., & Durrant-Whyte, H. (1997). Slip modeling and aided inertial navigation of an LHD. Paper presented at the IEEE

- International Conference on Robotics and Automation.
- Silver, D., Carsten, J., & Thayer, S. (2005). Topological global localization for subterranean voids. Paper presented at the Fifth International Conference on Field and Service Robotics, pp. 556–567.
- Silver, B., Bradley, D., & Thayer, S. (2004). Scan matching for flooded subterranean voids. Paper presented at the IEEE Conference on Robotics and Mechatronics (RAM), pp. 422–427.
- Silver, D., Ferguson, D., Morris, A., & Thayer, S. (2004). Feature extraction for topological maps. Paper presented at the IEEE/RSJ Conference on Intelligent Robots and Systems, pp. 773–779.
- Stentz, A., Ollis, M., Scheduling, S., Herman, H., Fromme, C., Pedersen, J., Hegadorn, T., McCall, R., Bares, J., & Moore, R. (1999). Position measurement for automated mining machinery. Paper presented at the International Conference for Field and Service Robotics, pp. 299–304.
- Stone, E., Ende, B., Wefer, F., & Jones, N. (2000). Automated 3D mapping of submarine tunnels. Paper presented at the Robotics 2000: ASCE Conference on Robotics for Challenging Environments.
- Wells, B. (1999). The latest development in laser profiling, borehole deviation, and laser enhanced videometry. Aberdeen, UK: Measurement Devices, Ltd.
- Whitcomb, L.L. (2000). Underwater robotics: Out of the research laboratory and into the field. Paper presented at the IEEE International Conference on Robotics and Automation, pp. 706–716.
- Whitcomb, L.L., Yoerger, D., & Singh, H. (1999). Advances in Doppler-based navigation of underwater vehicles. Paper presented at the IEEE International Conference on Robotics and Automation, pp. 399–406.