Robot Configuration for Subterranean Modeling

Zachary Meyer Omohundro

CMU-RI-TR-07-02

Submitted in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy in Robotics.

The Robotics Institute
Carnegie Mellon University
Pittsburgh, Pennsylvania 15213

February 16, 2007

Thesis Committee:
William "Red" Whittaker, Chair
Dimitrios Apostolopoulos
John Bares
Scott Thayer
Joseph Sottile, University of Kentucky

Contents

List of Figures

List of Tables

Abstract

Void models are invaluable for understanding subterranean conditions and guiding underground operations, but many maps are inaccurate, missing or difficult to acquire. Underground voids are inaccessible and often hostile to man and machine. Mapping and modeling of subterranean voids is a superb motivation for development of robots. Successful robot configurations for modeling subterranean worlds such as sewers, tunnels, mines and voids have yet to be created and characterized.

This research identifies subterranean modeling tasks and exploration constraints. Modeling tasks range from limited two-dimensional cross sections to comprehensive, globally-referenced, three-dimensional coverage of vast voids. Entry constraints such as small diameter holes limit the size of robot configurations, but once inside, muck and obstacles require significant terrainability that only comes with larger size. Based on these conditions and constraints, robot configurations are conceived, developed, tested and classified.

This research distinguishes three classes of robotic configuration for void modeling based on entry constraint and robot mobility. **Portal Explorers** are mobile robots featuring substantial locomotion, sophisticated sensing, and plentiful energy storage that can roam afar after being deployed via large diameter portals. **Borehole Sensors** provide exceptional modeling capability from fixed vantage points, while their miniaturization enables deployment down smallest diameter boreholes. **Borehole Explorers** combine some mobility with some borehole access miniaturization to acquire otherwise unobtainable model data. Mobility provides greater modeling coverage than possible with borehole sensors, while reconfigurability allows deployment via small diameter boreholes, enabling void access not possible with portal explorers.

This research classifies subterranean modeling tasks, develops robots to address these tasks, and tests these robots by entering and modeling vast voids. Unprecedented data sets thus created evaluate the merits and limits of these important robot configurations. Additionally, field work provides insight into issues and operational approaches that are not otherwise apparent from lab work alone. Finally, future robot configurations and paradigms for subterranean robotics are envisioned.

Acknowledgements

The majority of my time at the Robotics Institute has been spent doing what I love, building real robots. This would not have been possible without the support, encouragement, and example of those around me. From the first crazy days building Groundhog to recent refinements to CaveCrawler, I've learned so much from the people around me and want to take this opportunity to thank them.

My thesis committee has been an inspiration to me in the realm of real-world robots. John and Dimi both continue to produce robots which redefine the limits of robot mobility and robustness. Scott's commitment to transitioning from research to commercialization of technologies he believes in is something I hope to mirror in my own career. Joe Sottile keeps me honest, making sure my work is relevant in mining as well as robotics.

Red's intensity is a constant source of amazement and motivation. He has helped me see myself with a newfound clarity and has pushed me to discover how much I can actually accomplish as a builder, an engineer and a leader.

Chuck and Aaron have been steadfast in their commitment to subterranean robotics, more so than I. The most impressive models and results in this thesis come from Aaron's efforts, while the continued existence of the Subterranean Robotics Group is due to Chuck's refusal to allow any problem large or small to stop the team. Charlie Reverte and I are a fearsome robot development team. The design of CaveCrawler and Helix and the ruggedness of Groundhog are all results of our collaboration. My fellow builders and RI Shop denizens, particularly Mike Vande Weghe, Chris Atwood, Ben Brown, and Jonathan Hurst have shown me how to be a better builder. Groundhog and many of the core concepts of subterranean robotics could not have been developed without the aid of members of various Mobile Robot Design classes and other individuals, particularly Chris Baker, Dave Ferguson, Dave Silver, and Uland Wong.

Finally, I'd like to thank Mom, Dad and Zane for being the strongest, most consistent, and most persistent of advocates for the development, and, more significantly, the completion, of this document.

Chapter 1

Introduction

Robots are venturing into subterranean domains. Unlike well established robotics technologies for land, sea, air and space, research and experience in the subterranean world is just emerging. The applications, configurations, and challenges of subterranean robotics are yet to be distinguished. As with any new field, early obstacles to effective research are inexperience, lack of technologies and lack of information. Compounding this problem is the inherently hidden nature of subterranean environments. The insides of decades old abandoned mines, hundred year old storm sewers and breached bunkers are all examples of subterranean voids that are unobservable by humans. This lack of information both compels and hinders subterranean research. Existing robot configurations, tailored to unconstrained outdoor environments, possess some sensing and mobility required to operate within and model subterranean voids, but are too large to enter and negotiate within them. Miniaturized configurations designed for clean indoor environments can enter subterranean voids, but lack the mobility, robustness and sensing necessary to effectively model them.

Some prior research has addressed the challenges of subterranean robotics. However, this prior work has generally approached the problem from the context of improving mining efficiencies through automation and tele-operation. Mining has a rich heritage of machines and techniques for operating effectively both above and below ground. However, mining relies on skilled human operators, often operating in dangerous conditions performing repetitive tasks. Robotic techniques have automated underground equipment [Altafini, 1999], improved mining efficiency [Roberts et al., 1999] and demonstrated some underground modeling capabilities[Corke et al., 1998]. Some aspects of mine automation have matured beyond research into products [Elphinstone, 2006]. Some early experiments have also demonstrated crude robotic mapping of mines [Champeny-Bares et al., 1991] [Shaffer, 1995]. Finally, development of robot configurations suitable for operating in the most hostile of subterranean environments has been pursued in the context of mine distaster rescue response [Hainsworth et al., 1990] [McAteer et al., 2006]. However, these examples fail to

develop robotic void modeling configurations.

While this early research in the subterranean domain has demonstrated the utility of robots underground, it has fallen short in three key respects. First, prior work deals almost exclusively with subterranean environments which are human accessible. Second, prior work automates existing procedures, rather than creating new approaches, such as comprehensive 3D models, for understanding subterranean voids. Finally, prior research utilizes existing, large device configurations rather than developing new subterranean robot configurations tailored both to the environment and its access constraints. Research into modeling subterranean voids inaccessible to humans with new robot configurations is a distinguishing feature of this thesis.

This thesis explores the configuration space of subterranean robots. Three distinct classes of devices are hypothesized, constructed and tested. Data from real voids is acquired and incorporated into subterranean models of unprecedented size. The design space is then segmented on the two axes of entry constraint and device mobility. Configurations are evaluated based on their effectiveness at acquiring full coverage range data of underground voids.

1.1 Motivation

Numerous robots have been configured for working in tightly constrained environments. However, these environments are simplified by the nature of the constraints. For example, operating in a pipe of fixed diameter enables assumptions to be made about the environment, simplifying the robot designers task. Underground voids such as abandoned mines, cave complexes and collapsed bunkers present a different problem, wherein *access* is tightly constrained, but the environment inside the void is often expansive, unmaintained and filled with challenging obstacles. Robot configuration and operational concepts that address this pair of conflicting conditions have yet to be demonstrated.

Recent advances in the miniaturization and commercialization of components such as laser range finders, computers and inertial measurement units enable constrained access compatible robot configurations. Developments in geo-referenced model generation using laser range data allow for the creation of detailed, accurate underground void models. However, the fusion of robot configurations with model generating requirements first requires robot presence and data gathering underground, and has yet to be demonstrated. This thesis demonstrates new operational approaches and robot designs for acquiring subterranean void models.

1.2 Applications

Quantitative knowledge of subterranean voids is desirable in a number of disciplines. Having robotic systems capable of acquiring such information therefore has high utility. Motivating application areas include miner safety, environmental impact mitigation, and improved economic efficiency underground, among many others.

1.2.1 Safety

Thousands of mine deaths occur annually worldwide [CNN, , Bailey, 2003], and many of these could be prevented by emerging robot capabilities. Structural failure of subterranean voids due to inaccurate understanding of void geometry can result in dangerous flooding. Often, underground voids are close to large volumes of water, like old, flooded mines or adjacent lakes. Three distinct types of flooding scenarios are possible. First, workers in active mines can accidentally breach into adjacent inaccurately mapped flooded mines. Secondly, barriers holding back water can be subject to erosion and degrade over time to the point of failure. Finally, above-ground bodies of water can break into underground voids, with flooding comparable to dam failures. Each of these scenarios results in uncontrolled and dangerous flooding, putting mine workers or those downstream of underground voids at risk.

Breaches between active and abandoned mines can be caused by inaccurate or non-existent abandoned mine maps. Breaches involve the sudden failure of mine walls between active and abandoned, flooded mines, usually where miners are working. Breaches can cause harm to miners in multiple ways. The explosion-like energy of the initial failure at a breach can harm miners. Should the miners survive the chaos of the initial failure, they face the potential of being trapped by rising water levels. Modern mining practices dictate that portals be upstream of the active face. Therefore, water will continue to fill the active mine, rather than safely draining out the portal. Based on mine geometry, this can result in miners being trapped underground until water levels can be brought down via pumping. One specific example of this type of accident is the Quecreek/Saxman mine accident, described in detail in section 3.3.2.

Insufficient monitoring and modeling of mine walls subject to erosion can also result in flooding. Erosion can be caused by rapidly moving water or through the slow dissolution of barrier material. Many minerals and geologic structures have salt like properties, meaning they dissolves into water until the water becomes saturated. Therefore, if a barrier wall composed of such a substance is in proximity to a continuous source of fresh water, gradual erosion of the barrier will occur. Different water levels on different sides of such a barrier setup conditions for flooding should failure occur. Monitoring of barrier wall thickness through subterranean modeling is useful in determining if such conditions are developing.

A Wyoming Trona mine (Trona is a salt-like mineral) provides a specific example of this type of scenario where robotic modeling could prevent a dangerous failure, and is described in section 3.3.2.

Flooding of lakes into an underground void occurs when the material between the lake and void can no longer support the hydrostatic pressure of the lake. In effect, such failures are comparable to undercutting a dam, with water escaping without compromising the perimeter barriers containing the body of water. Underground voids further worsen the situation as they often direct water to multiple void portals, which can be widely separated and are unexpected sources of flooding. Section 4.2 details one such incident where contaminated water flooded into two separate watersheds via a subterranean void.

Flooding is not the only threat to the safety of underground workers that could be alleviated with robotic void modeling. Miners also face the danger of breaching dry abandoned mines. While the resulting breach is not accompanied by flooding, the atmosphere in abandoned mines is usually low in oxygen, and high in explosive gases such as methane, resulting in potentially hazardous conditions for the miners. Either the lack of oxygen, or the mixture of the oxygen in the active mine with the methane in the abandoned mine and the resulting explosion hazard can result in the injury or death of miners.

Robotic void modeling can provide the information needed to detect, quantify and remedy potentially dangerous underground conditions. Accurately measuring wall thickness between adjacent underground voids prevents miners from unintentionally breaching into flooded or gas filled voids. Modeling robots occupy voids, providing direct measurements of void location and geometry unmatched in accuracy, resolution and density by current techniques. Hence, robotic modeling of subterranean voids has the potential to improve miner safety.

1.2.2 Environmental

Robotic modeling can provide the information necessary to detect, monitor and correct environmental problems arising from subterranean voids. Underground voids, particularly mines, can negatively impact the ground, air and water. Environmental damage from mines falls into three primary categories: subsidence, mine fires, and acid mine drainage, as shown in figure 1.1. Remediation of these problems is beyond the capacity of current robots, but providing accurate maps and information about the environmental conditions within voids can improve the effectiveness of conventional reclamation methods.

Robotic void modeling can prevent subsidence from spreading or reoccurring by improving subsidence remediation techniques and verification. Subsidence results when a mine roof collapse propagates to the surface, and manifests as a depression, sinkhole or fracture. Houses, other structures, and roads are routinely damaged by subsidence [Ruegsegger and

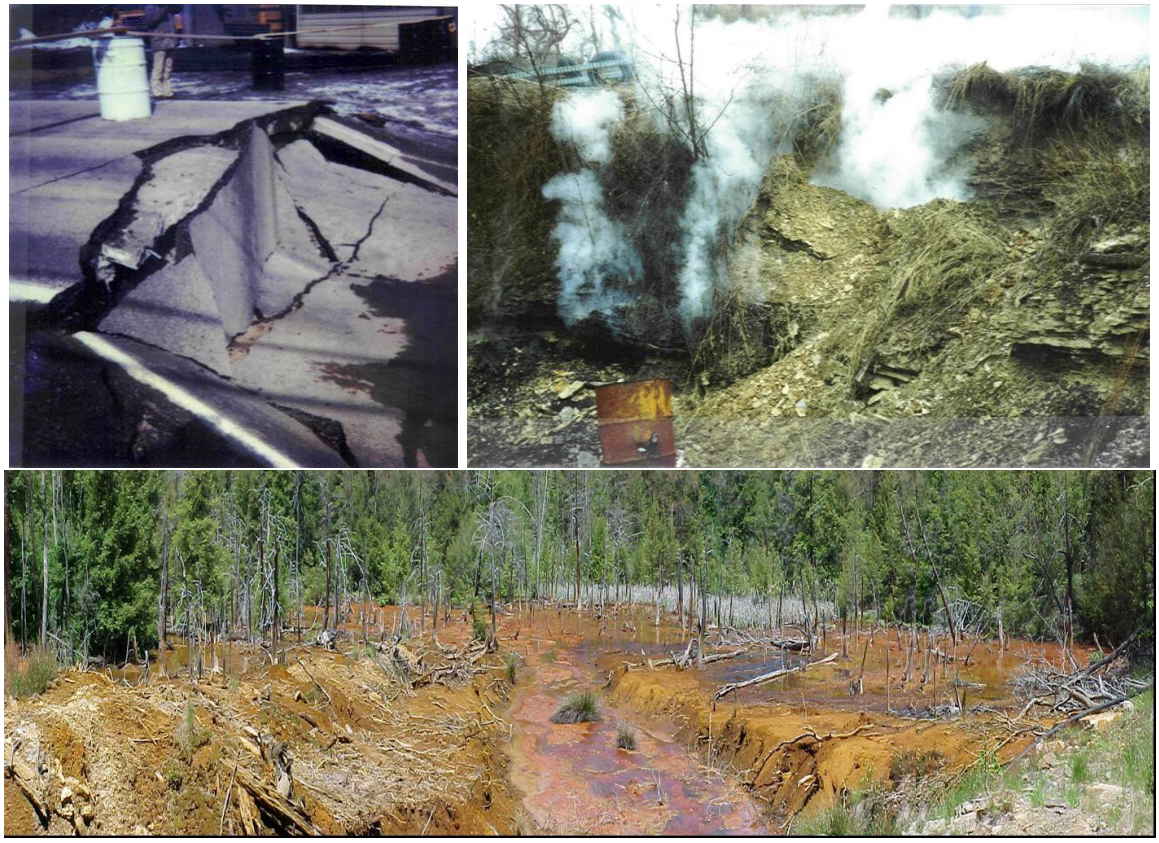

Figure 1.1: Problems caused by subterranean voids: Subsidence (top left), mine fires (top right), and acid mine drainage (bottom center).

Lefchik, 1999]. In order to prevent further subsidence from a given abandoned mine, boreholes are drilled into the mine, and grout, fly-ash, or concrete is pumped into the mine in an attempt to fill the mine with structural material, thus preventing any further problems. However, the fill material is usually pumped in from a limited number of boreholes, and coverage of the entire mine is difficult to determine, especially if the volume of the void is not initially known. Robots can provide volumetric and topological information regarding problem mines. Knowledge of where and how much backfill to inject can improve subsidence remediation quality while reducing costs. Additionally, post-backfill modeling can verify successful remediation and identify any void volumes with the potential for additional subsidence.

Robotic mine mapping can assist in planning for acid mine drainage treatment. Acid mine drainage results from water filling and flowing through mines. The sulfur and other chemicals within the coal leach into the water, turning it into a highly acidic solution, which then drains out of the mine and into the surrounding ecosystem. Acid mine drainage

is treated by adding basic compounds, such as lye, to the mine water in treatment pools. Routing and piping of contaminated mine water can use existing mine voids if accurate maps of these voids exist. Using robots to acquire maps for acid mine drainage pipe routing has been demonstrated as part of this research [Baker et al., 2004].

Robotic subterranean modeling can assist in combating mine fires. Mine fires, fueled by numerous vents and an unlimited supply of combustible coal, can quickly grow so large that it is no longer feasible or safe for people to enter the mine or fight the fire using conventional techniques. Instead, the fire is isolated and starved of fuel through the emplacement of concrete plugs. As with subsidence problems, void modeling robots could be used to direct and verify the placement and completeness of these plugs. Mine fires have been known to persist for decades [DeKok, 1986], turning the surface above the fire into a smoking wasteland, so any improvement in fire-fighting techniques that can be provided by robots is an environmental victory.

1.2.3 Economic

The Quecreek accident and its wide public exposure resulted in regulator agencies imposing additional and stricter requirements on permitting and operating of mines near abandoned mines, resulting in increased costs to mine operators. Specifically, the safety perimeter around abandoned mines was expanded from 200 feet to 500 feet to account for uncertainties in maps of surrounding voids [PADEP, 2002]. Maintaining this buffer would leave a fortune in unmined coal around the perimeter of every new mine. Mine operators are only allowed to encroach on this perimeter if they can show hard evidence of the locations and extent of adjacent abandoned mines. Therefore, the capability to accurately and cost effectively provide abandoned mine maps serves to increase the allowable size of a mine, and hence its output. A mobile borehole-deployable robot could provide many hundreds of feet of mine map from a single borehole, thus significantly increasing range and reducing the required cost of mapping.

While exhaustive borehole drilling and directional drilling can both be used to crudely validate the location of nearby abandoned mines, only subterranean robots can provide comprehensive 3D map data. While 3D data is not required for determine mine extent, it is useful in determining the amount of coal removed and remaining in an abandoned mine. Many strip mining operations intersect older deep mine operations. This results in a loss of revenue for the strip mine as the ratio of coal to overburden decreases. At some point it becomes uneconomical to continue strip mining, but this determination is difficult or impossible to make without a map of the original underground mine. Therefore, by providing an accurate estimate of the volume of coal removed from an abandoned underground mine, the strip mine can be run more effectively.

Volumetric surveying of abandoned mines can also reduce uncertainties in mineral rights ownership, taxation and production royalties. Proof of the absence of minerals can result in taxes being dropped. Therefore, the cost of robotic mapping of abandoned mines somewhat offsets or displaces recurring fees, especially when a single mapping cost is amortized over many years of taxation.

1.3 Subterranean Modeling

Subterranean voids possess numerous features that uniquely distinguish them from other domains. These features translate into design constraints placed on subterranean robot mechanisms, sensing and software. For example, lack of global positioning information necessitates high quality inertial sensing or improved software localization. Similarly, lack of communications requires higher levels of autonomy than with surface robots. This research was tightly coupled with complementary research for autonomy, pose estimation and model fusion. However, this thesis doesn't overemphasize issues such as robot intelligence, communications, or localization except as they impact robot mechanical and sensing configurations. This thesis distinguishes features of subterranean voids and identifies how these features simultaneously compel robot usage while constraining mechanical and operational configuration design.

Inhospitable to humans: Underground spaces, such as mines, rescue sites or enemy bunkers, are often hazardous to humans, and in many cases, such as sewers and abandoned mines completely inaccessible to human entry. This motivates the use of robots and precludes human rescue of defective robots, thus requiring robust and reliable robot configurations.

Limited access: Entering underground voids is often difficult. Fully sealed abandoned mines (i.e. no entry points remain) are common. Limited access compels robot systems which can provided extensive void modeling information from a single entry, while also compelling configuration compatible with small diameter access points. Two primary types of access constraints are:

Unconstrained Accesses or Portals are gateways that have cross-sections of comparable size to void interiors. Portals are usually located on the perimeter of voids and are limited in number and fixed in location.

Constrained Accesses or Boreholes are small diameter holes drilled into voids. Boreholes allow greater flexibility in entry placement, but severely limit entry diameter.

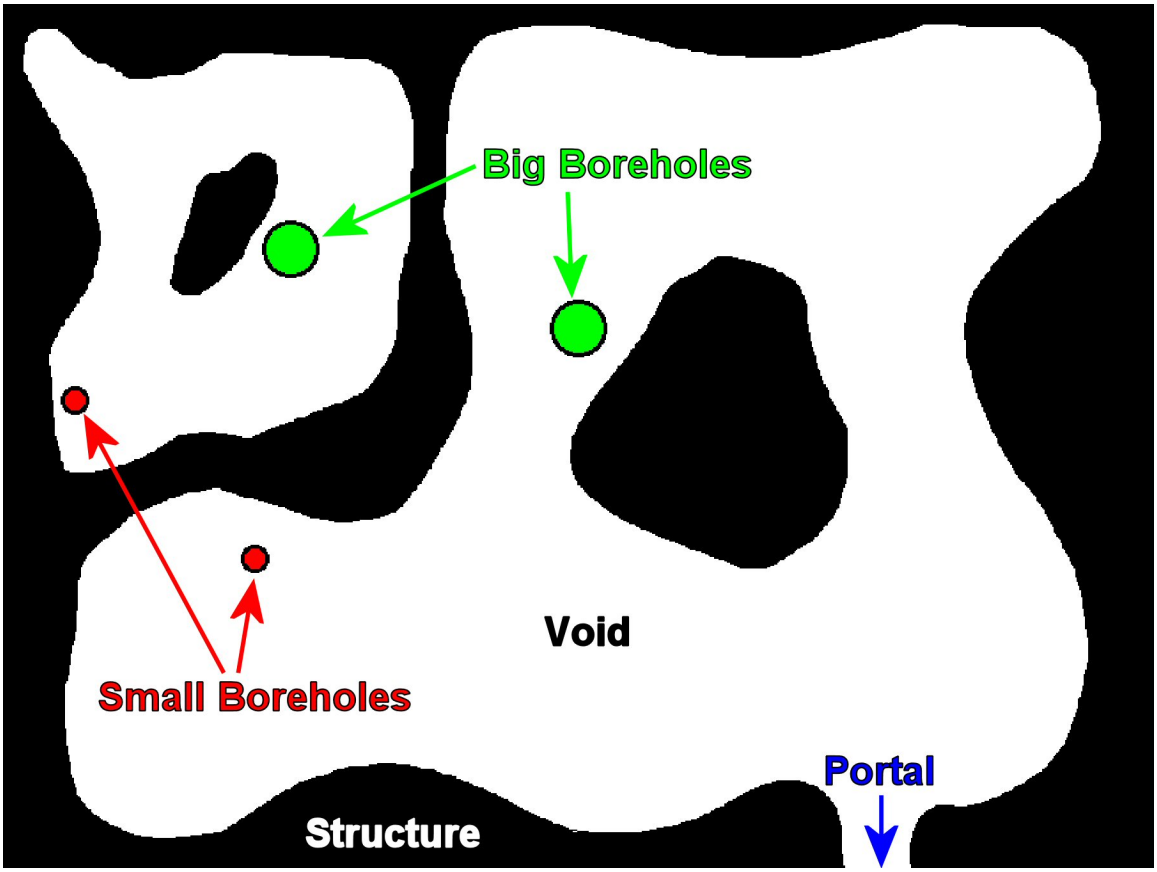

Figure 1.2: Representative geometric abstraction of subterranean environments

Unknown void topology: The topology, geometry and extent of many underground spaces is completely or partially unknown. Beyond motivating the work, lack of a priori information about void topology necessitates exploration and eliminates many sensor coverage planning techniques which rely on accurate prior knowledge. Further, unknown topology means that flexible operational approaches to subterranean void modeling are required in order to take full advantage of limited void access.

Non-convex geometry: A limited set of underground voids, such as buried fuel tanks, have geometries which can be fully observed from a single vantage point. However, the more prevelant class of voids, of which figure 1.2 is an example, has no single vantage point from which complete sensor coverage is possible, regardless of sensor range and resolution. This motivates robotic systems capable of easily achieving multiple vantage points.

Subterranean modeling and exploration can be described at the highest level by a simple procedure. First, a robot enters a void. Second, it maps and models visible portions of the

void using its sensors. Third, it changes vantage point in order to acquire additional data not obtainable for its prior locations. Steps two and three are repeated until the void is modeled or the limitations of robot modeling abilities are reached. Finally, the robot exits the void. More detailed descriptions of these elements are below:

Entering/Exiting: Underground spaces can be accessed either through portals or boreholes. Small diameters entries usually require robot reconfiguration for deployment. Exiting requires robots to both successfully return to an enter/exit point and reconfigure to be compatible with void egress.

Sensing: Range data can come from laser, radar, or sonar depending on void environmental conditions. Point sensors provide a single range data reading. Mechanically scanning point sensors along one or two axes provide 2D or 3D information respectively. Additional sensing such as still or video imagery, environmental monitoring, and pose estimation is also often acquired and utilized in the construction of integrated models.

Moving: Robot motion changes viewpoints to accumulate sensor coverage of underground environments. Overlapping data sets from multiple viewpoints can be fused into comprehensive full coverage models. For example, a mobile robot with accurate pose estimation and 2D range sensing can create 3D models by moving through an environment. A static 3D range sensor capable of multiple deployments into different entries could obtain multiple overlapping data sets. Finally, a mobile robot with 3D sensing could move in an underground space, acquiring high resolution 3D scans when stopped at known locations. In all of these examples, motion of the robot allows for more complete sensor coverage of subterranean voids.

1.4 Thesis Statement

Robots enter and exit voids through holes and portals that range from small to large. Smaller holes introduce smaller robots with lesser capabilities, while larger portals permit large robots with greater modeling abilities. Mobility within voids varies inversely with robot size. The span of subterranean modeling tasks requires robot configurations with various combinations of mobility and miniaturization. An understanding of possible robot configurations is required in order to meaningfully approach subterranean modeling.

This thesis asserts that subterranean robots can be distinguished by robot access constraint and robot mobility. Furthermore, the span of subterranean modeling objectives can not be effectively achieved with a single subterranean robot configuration.

For this thesis:

Access constraint is the minimum diameter through which a robot must pass during deployment and recovery.

Robot mobility is the capability to reach geometrically accessible points within a void.

1.5 Robot Configurations

Mobility increases with increasing size of robots that can enter voids. Large void entries admit large robots that have mobility to roam afar and map large regions. Mobility decreases with decreasing size of robot that can enter a void. The smallest void entries, like boreholes, admit only sensors that lack mobility. These sensors cannot roam and cannot map regions beyond what they can sense from their fixed vantage. This thesis conceives, develops and demonstrates a new form of robot that can enter voids through small holes and then roam and sense to an extent to create extensive void maps. From these core concepts three classifications for robot mobility emerge.

Zero Mobility classifies static sensors which can only move through external forces, such as the lowering of a deployment tether.

Limited Mobility configurations provide mobility sufficient to move beyond the view of a static sensor, but not everywhere within the void.

Unlimited Mobility system are able to reach any geometrically accessible point within a void.

Three robot mobility classifications, when crossed with the two access constraint classifications outlined in 1.3 result in six potential design space regions of interest in the domain of subterranean robot configurations.

Constrained entry, zero mobility: consist of miniaturized scanning mechanisms coupled with borehole compatible range sensors.

Constrained entry, limited mobility borehole-deployed mobile robots are constrained in size, thus limiting energy storage capacity, computing, and available sensing. However, they poses mobility and can therefore "see around the corner," something not possible with static sensors.

Constrained entry, unlimited mobility Robots which could be deployed down borehole but without any limitation on energy storage, sensing or computing would be the ideal solution to subterranean modeling. However, such configuration are unachievable.

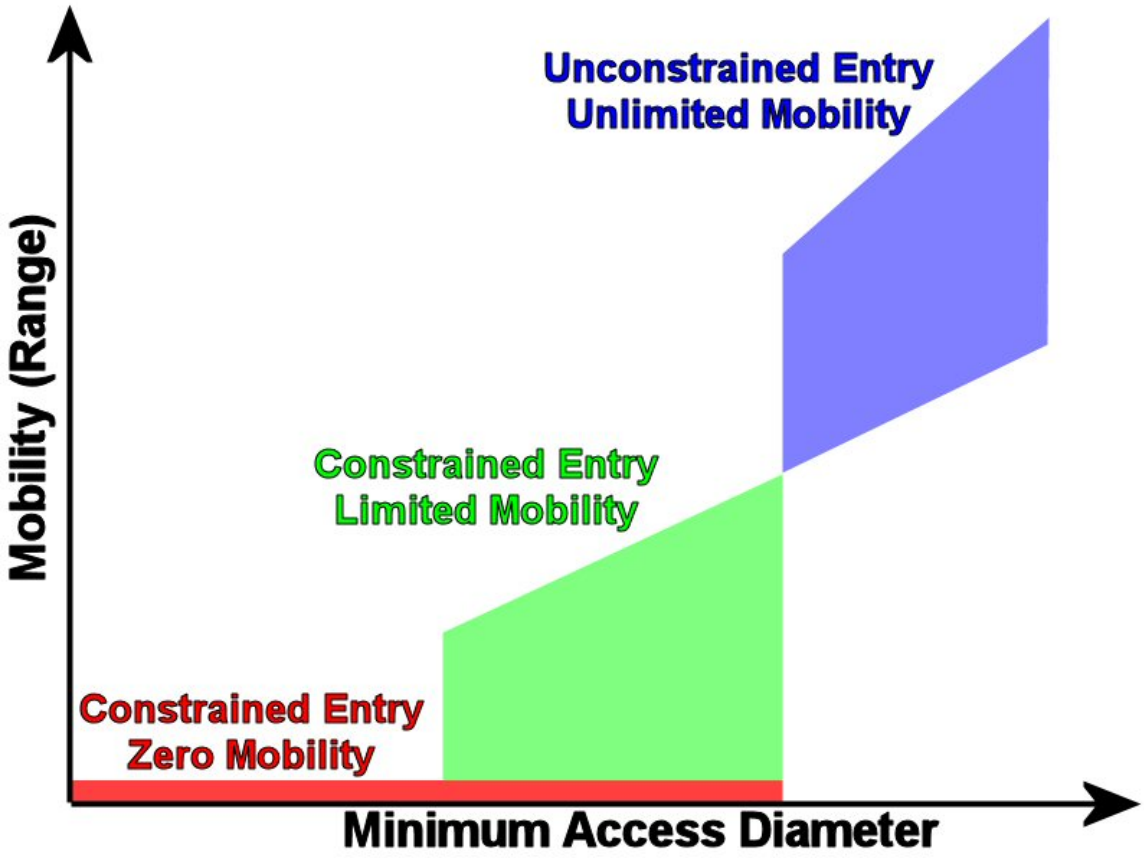

Figure 1.3: Regions in the design space of subterranean robots

Unconstrained entry, zero mobility Without the requirement for miniaturization due to constrained entry static sensors for 3D modeling are easily found commercially. Specifically, laser range finders for architectural and as-built modeling are commonly available.

Unconstrained entry, limited mobility Non-specialized mobile robots designed for purposes other than subterranean modeling can still gather data and create models within underground spaces. However, these configuration lack the combination of long range, terrainability, and non-GPS dependent pose estimation necessary to successfully fully traverse subterranean voids.

Unconstrained entry, unlimited mobility Mobile robots specialized for underground operations can move to any point within a void, but require portal access.

The research creates, develops, and deploys robot configurations falling into the three classes highlighted in figure 1.3. **Borehole Sensors** are constrained entry, zero mobility

robots useful in the most challenging voids. **Borehole Explorers** enter voids through constrained entry, then explore with limited mobility, providing more data from one borehole than possible with borehole sensors. **Portal Explorers** require unconstrained entries, but once inside a void exhibit unlimited mobility and produce extensive models. Development, testing and analysis of these three core robot configurations demonstrates the capabilities, limitations, and relevance for subterranean void modeling of these robot configurations.

1.6 Outline

This thesis explores the design space of subterranean modeling robots. Chapter 2 details prior work relevant to the task, while chapter 3 outlines the steps and specifications involved in subterranean modeling. Chapter 4 describes the over-arching philosophy of this thesis. Chapters 5 identifies existing configurations that have been retasked to subterranean modeling objectives. Portal explorers are examined in Chapter 6, while borehole sensors are reviewed in Chapter 7. Chapter 8 analyzes the middle ground of the design space, borehole explorers, and Chapter 9 summarizes the contributions of this thesis and imagines possible future endeavors in the field of subterranean modeling.

Chapter 2

Background

This research innovates by expanding on prior work. Because this research concerns configuration or system level design, a broad foundation of prior work is required. Modeling quality of traditional mapping methods like geophysical sensing, prior maps, and borehole drilling must be understood to be evolved. Prior automation of underground machinery demonstrates navigation algorithms, sensing modes and methods for hardening automated devices. Algorithms used in indoor, GPS-denied modeling can be applied underground. Robot configurations which reconfigure to negotiate tight entry constraints provide mechanisms and inspiration for designing borehole deployed robots. This thesis builds upon prior work in these varied domains to achieve new robot configurations and produce never before seen subterranean models.

Traditional manual methods currently used in modeling subterranean voids provide minimum levels of accuracy, detail, and extent in void modeling that must be meet or exceeded to make robotic modeling relevant. These manual techniques also provide insights into effective operational procedures for modeling and serve as the basis for some operational approaches to robotic modeling.

Underground work machinery has been successfully automated in the past. While this automation has only improved production efficiency in known subterranean environments rather than modeling unknown ones, the mechanisms, algorithms, and approaches to surviving hostile underground environments are of use. However, subterranean robots must do more than simply survive the environments, they must model it.

Indoor environments provide a strong analog to subterranean ones. Modeling of indoor environments is a long standing field of research, and algorithms to solve canonical problems such as the Art Gallery and Next Best View coverage problems are of use, particularly in the case of borehole-deployed robots where properly locating boreholes greatly impacts model coverage, total number of required boreholes, and overall modeling system efficiency. Approaches to indoor modeling can be applied underground once successful subterranean

robot configurations are developed and demonstrated.

Examining prior robot configurations designed for deployment or operations within constrained environments provides system and subsystem concepts that can be built upon or utilized in developing borehole-deployed mobile robots. However, the combination of borehole deployment constraints and expansive, hostile subterranean voids requires innovation beyond that demonstrated by existing reconfigurable robots.

While many other disciplines and algorithms are touched upon by this research, these four core background areas are the most relevant and therefore expounded upon in detail below.

2.1 Traditional Subterranean Modeling

It is essential to understand traditional modeling techniques for three reasons. First, traditional methods represent accepted standards in the mining industry. Robotic modeling must produce results of higher quality at lower cost for robotic modeling to be relevant. Second, not all subterranean modeling tasks can be accomplished using traditional techniques. Therefore, demonstration of robot configurations and modeling approaches that can attempt these previously impossible modeling tasks validate both subterranean robotics in general and this thesis specifically. Finally, robotic modeling approaches can make use of traditional techniques by adapting, integrating, or supplementing them, but only if traditional techniques are understood. In many cases traditional void characterization techniques combined with robotic systems results in models not otherwise obtainable.

Traditional methods of mapping subterranean voids provide some data as to location and extent of underground voids, but have fundamental limitations in applicability, resolution, and accuracy. Historical maps provide detailed information about void geometry and topology, but are often inaccurate, out of date, or non-existent. Indirect geotechnical methods infer void location and extent through surface measurements combined with soil models, but lack resolution and are often ambiguous. Direct methods such as borehole drilling provide unambiguous data, but at such limited density it is rarely useful for modeling. Three primary traditional methods of underground modeling, historical maps, indirect geotechnical methods, and borehole drilling are described in further detail in the sections below. Composite methods that utilize some combination of the above are also described.

2.1.1 Prior Maps

Prior maps of underground voids, including centuries-old historical documents, are the first, easiest, and cheapest reference for determining the extent and topology of voids. However, errors and omissions in these maps often create false information about the extent and

location of underground voids that can lead to disaster, as occurred at Quecreek. Historical maps can be inaccurate in four different ways. First, maps of a void might not exist, or might have existed at some point but have since been lost or destroyed. Secondly, void maps can be inaccurate or idealized due to the methods used in their original creation. Thirdly, accurate maps can go out of date as an underground void is expanded by human or natural actions. Finally, intentional inaccuracies in maps can be introduced by dishonest operators or map makers. Despite all of these factors, historical maps will likely continue to be the first source examined when determining underground void geometry, but they must not be the only source.

If surveyed maps of an abandoned mine adjacent to a proposed new mine exist these are often used with little or no additional measurements of the abandoned mine undertaken. However, this approach has some critical vulnerabilities as demonstrated by the Quecreek mine accident. The primary failing of existing maps is their reliability. Older maps are hand drawn to various quality levels, and while most states with mining industry now have standards in place for the accuracy and quality of mine maps, these standards vary from state to state, and have evolved through time. Even given a highly detailed and accurate map, there is no guarantee that any map accurately reflects a mine void. Additionally, many mine maps have been lost or accidentally destroyed, such as the loss of 30,000 mine maps in a fire at the Kentucky State Department of Mines and Minerals in 1948 [Council, 2002]. Finally, while regulatory agencies benefit from accurate mine maps, mine operators, who are taxed based on the amount of coal removed, benefit from maps which under represent the size of a mine. While such a practice is illegal and presumably uncommon, the potential for intentional inaccuracies in abandoned mine maps does exist. The current method for dealing with faulty maps is to prohibit active mines from being within a given distance of the presumed location of an abandoned mine, but such a precaution were insufficient in the case of Quecreek, so an improved method of determining the extent of abandoned mines is required.

2.1.2 Indirect Geophysical Techniques

A number of geotechnical methods can partially detect and characterize mine voids, but these methods require knowledge of the surrounding strata, void condition, and the resulting data deliverables are of limited resolution and accuracy. A summary of the available techniques, taken from [Council, 2002], is shown in Figure 2.1.

One common requirement of all indirect void detection methods is the need to drill boreholes in order to confirm the hypothesized void location. Geophysical techniques can guide, but not supplant, exploratory borehole programs. Therefore, even when conditions are favorable for the application of geophysical techniques, the use of more advanced void

Technique	Material	Depth	Cost	Comment
Airborne imaging	Air	Surface	Low	Fast but Low Resolution
Electrical Resistivity	Water	656 feet	Medium	Good if water-filled but poor in air-filed
Electromagnetics (active)	Water	656 feet	Low	Fast, easy data collection but poor in air-filled
Ground-penetrating radar	Air	33 feet	Low	Fast, easy data collection but clay wipes out signal
Induced polarization	Water	656 feet	Medium	Unknown, success undemonstrated
Magnetics	Scrap metal	164 feet	Low	Easy data collection but need metal in mine
Microgravity	Air	66 feet	Medium	Simple theory but impacted by geologic noise
Seismic reflection	Air or water	656 feet	High	Produces image but intensive data processing
Seismic refraction	Air or water	656 feet	Medium	Easy data processing but indirect detection
Seismic surface waves	Air or water	98 feet	Medium	Easy interpretation but limited depth of penetration

Table 2.1: Indirect geotechnical void sensing techniques

mapping technologies, deployed via borehole, can significantly improve the quantity and quality of void data obtained. Additionally, the limitations on depth, resolution, and accuracy of geophysical techniques increasingly require borehole based methods as mine depth grows.

2.1.3 Borehole Drilling

Exploratory borehole drilling is the only method of conclusively determining the existence of mine voids from the surface. However, each borehole only provides a single bit of data indicating the presence or absence of coal beneath the drill rigs position, and while lack of coal conclusively indicates a void, the inverse is not true. Therefore, borehole drilling must explore the perimeter of abandoned mines exhaustively in order to confidently determine void extent. The number of boreholes required by exhaustive search is both cost prohibitive and environmentally damaging, but decreased pattern density results in lowered void extent confidence levels.

Directed horizontal drilling is a technique that drills from active mine working faces toward theorized abandoned mine locations. In contrast to the binary information provided when drilling from the surface, each horizontal hole provides one analog value, the distance

from the drill operator to the abandoned mine void. Additionally, while the path of the horizontal hole does not necessarily represent the minimum distance between mines, it does provide conclusive evidence of the absence of voids along said path. Therefore, the perimeter of the adjacent abandoned mine can be more accurately and efficiently mapped than with surface drilling. Unfortunately, horizontal drilling provides a direct connection between the abandoned and active mines, thus allowing for water and gases to flow from abandoned to active mine. While horizontal drilling companies have developed and deployed techniques for ameliorating these potential hazards, the risk to miners is increased over drilling boreholes from the surface. Also, since horizontal drilling must be conducted from within an active mine, the drill rig required must be compatible with the regulations and geometric restrictions associated with underground mines. These factors drive up the cost of horizontal drilling as compared to vertical drilling, but the greater data utility has resulted in this being the preferred method of determining abandoned mine void location.

2.1.4 Composite Techniques

Borehole cameras are lowered down holes to provide qualitative visual information about subterranean voids, but fail to provide useful modeling data. Cameras provide more information than drilling, expanding a single bit of data into estimates of void geometry with the aid of skilled human observers. However, cameras fail to provide quantitative data about void topology or geometry. Without quantitative model data, camera surveys from multiple boreholes or locations can not be integrated into comprehensive models. Cameras only provide glimpses, not models, of subterranean voids. Further, cameras rely on scene illumination, which decreases with the square of distance, limiting effective range.

Borehole cameras are established technologies, and make use of commercial miniaturization and low-light illumination techniques. Therefore, they can fit down very small diameter boreholes, and can make use of existing mine certified housings. Subterranean modeling robots will likely included camera systems in order to improve tele-operation, provide human-accessible imagery, and verify range data based models. However, useful subterranean modeling requires data accuracy, density, and extent only possible with range-finder equipped robots.

2.2 Subterranean Robots

Robots have been configured to operate underground, but have not yet been designed to model underground. Prior subterranean robots provide insight into environmental constraints on configurations. Previous developments of subterranean robots clarifies lessons learned, challenges yet to be overcome, and minimum levels of capability required for new

configurations to achieve relevance. Background work in mining equipment automation, sewer inspection, and mine disaster response robots all have relevance toward developing new configurations for subterranean modeling.

2.2.1 Mining Equipment Automation

As an example, robotic systems such as Autonomous Guidance Vehicles (AGVs) [Automated Mining Systems, 2006] and CSIRO's Load-Haul-Dump truck (LDH) [Scheding et al., 1997, Roberts et al., 2000, Roberts et al., 2002, Duff and Roberts, 2003, Duff et al., 2003], 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 AGV navigates a mine 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 map[1] and information gathered from prior runs to guide the vehicle from corridor to corridor.

Mineral extraction platforms, such as continuous miners, 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 on-board sensors that read off-board survey lasers [Stentz et al., 1999].

2.2.2 Sewer and Pipe Inspection Robots

In non-maintained or natural environments, SANDIA's cave robot, RATLER [The RATLER Project, 2006], and RedZone's sewer robots [RedZone Robotics, 2006] navigate within cluttered caves and sewage-filled pipes, respectively, for the purpose of inspection. These systems are tele-operated via communication links such as tethers or radio signals. Due to the difficulties encountered in non-maintained environments, these robots do not utilize autonomy as part of operations.

The first robot to capture map data from a mine environment was "Terregator," a six wheeled multi-purpose research vehicle [Champeny-Bares et al., 1991]. Terregator achieved semi-autonomous navigation using a combination of sonar and laser scanners for position estimation and obstacle avoidance. These sensors also enabled Terregator to acquire high resolution scans of coal mine interiors.

[1] A topological map is a graph representation where the nodes and edges of said graph correspond to distinct locations in the environment [Kuipers and Byan, 1991].

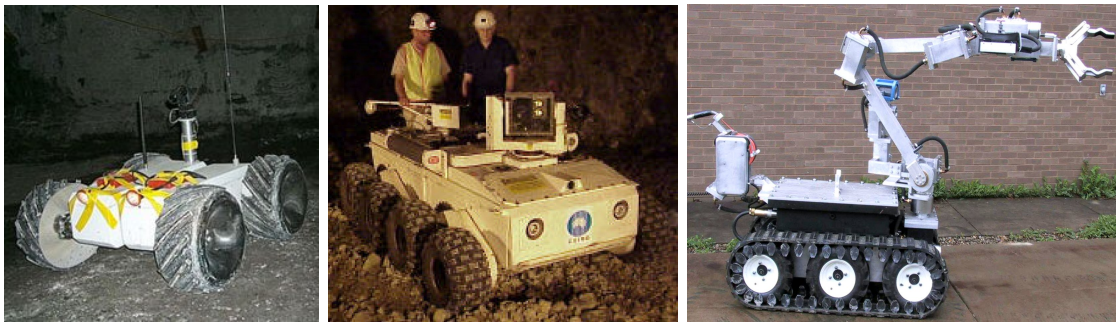

Figure 2.1: Robots designed or deployed in response to mine disasters: RATLER mobile robot (left) Numbat mine disaster response robot (center) MSHA customized ANDROS response robot (right).

2.2.3 Mine Disaster Response Robots

Mine disaster response robots have been developed to operate underground and aid human understanding of dangerous subterranean voids. However, these prior systems are not equipped for void modeling and require unconstrained entries in order to enter voids. Three specific systems (seen in figure 2.1) embody different approaches to applying robots to mine disasters.

RATLER is a research robot platform that was retasked to demonstrate robotic capabilities underground. RATLER is a differentially driven two-part robot with passive, body averaging suspension. Deployed underground in response to a mining accident in 1999, RATLER demonstrated basic mobility in an underground mine. RATLER acquired camera imagery, but did not attempt to create void models. RATLER was tele-operated throughout its deployment and served primarily as a remote camera platform.

The Mine Safety and Health Administration (MSHA) customized an existing disaster response robot platform, the Andros, to be compatible with underground mine conditions and regulation. MSHA's Andros features multiple visible and thermal camera systems as well as gas sensors on a mast atop a tracked base. Tele-operated via fiber optic tether, MSHA's Andros also functions primarily as a remote camera and environmental sensor platform. It is not equipped or designed for modeling, but it does feature a manipulator for moving small obstacles out of its path. MSHA's Andros was deployed as part of the response to the Sago mine disaster (see section 4.2.3), however it became bogged down in deep water only a short distance from its control station and ultimately failed to speed or improve the disaster response.

The robot Numbat, unlike the prior two examples, was designed from the ground up with mine disaster response as its only mission. Numbat, an 8-wheel differentially driven robot with rocker-bogey-like passive suspension provides remote video feeds to a tele-operator via

a fiber optic tether. Numbat was developed in response to mining accidents in Australia and was deployed and tested underground. Numbat was designed with mine rescues in mind and features an interface which allows communications between trapped or injured miners and remote teleoperators accross regions of a mine which are too dangerous for human rescuers to traverse.

2.3 Indoor Modeling

Hallways, intersections and obstacles commonly found inside buildings have topology, geometry and GPS-denial similar to that found inside many subterranean voids. Therefore, exploration approaches, sensor configurations, and localization techniques suited for indoor environments have relevance for subterranean modeling. However, indoor systems ignore configuration requirements such as terrainability, range, and volumetric constraints that pertain underground. Optimal sensor placement algorithms for full coverage have also been developed for indoor environments. These algorithms can be directly applied to borehole placement for improving subterranean modeling efficiency. Algorithms are not the principal thrust of this research but a few are utilized here to illuminate issues and approach developed in the context of indoor robots that are applicable to subterranean modeling.

2.3.1 Next Best View (NBV)

Selecting subsequent borehole locations based on prior data is canonically referred to as the Next Best View (NBV) problem. NBV is an offshoot from the Art Gallery problem, in which the goal is to provide complete sensor coverage of known polygonal spaces using the minimum number of sensors. Optimally (from a number of boreholes, but not necessarily cost, perspective) mapping completely known mines devolves into a modified Art Gallery problem with the added constraints imposed by sensor range and incidence angle limitations [H.H.Gonzalez-Banos and J.C.Latombe, 2001]. Uncertain mine topology requires a mapping approach that only relies on data directly acquired from the environment, which leads to determining which next sensor location, or view, will yield the most new data, and therefore be the "best" view. By combining the non-uniform borehole cost map with the NBV algorithm presented in [Gonzalez-Bantos and J.C.Latombe, 2002] and based on [Banta et al., 1995], [C.I.Conolly, 1985], and [J.Maver and R.Bajcsy, 1993], an optimal next insertion location can be determined.

2.4 Constrained Access Robot Configurations

Robots designed for reconfigurability or compatibility with small diameter entries incorporate specialized mechanisms and stowage designs in order to achieve high levels of functionality in small volumes. Borehole-deployed subterranean robots must satisfy the combined constraints of limited diameter boreholes and harsh, extensive subterranean terrain. Such configuration have yet to be created and are the aspirations of this research. However, subsystems and mechanisms employed by miniaturized or transforming robots in other closely related applications can be adapted and used in borehole-deployed robot configurations. Relevant systems include tank inspection devices, pipe exploration systems, snake robots, and other configurations with deploy in some fashion. The following sections extract and highlight mechanisms which could be useful in borehole-deployed configurations from existing systems.

2.4.1 Pipe Inspection

Pipe inspection systems have form factors similar to the stowed configuration required of borehole-deployable devices. Pipe systems assume operation within pipes of a given range of diameters, and are not suited for operation in unconstrained terrain, however, the knowledge about packing and configuring such systems for small diameter constraints can be applied to the individual segments of borehole-deployable devices. Additionally, many pipe inspection systems are designed to navigate around sharp bends and feature highly miniaturized componentry. Therefore, pipe inspection devices can act as a source for borehole-deployable robot components and packing schemes rather than entire subsystems.

2.4.2 Serpentine Robots

Hyper-redundant snake robots might seem to present ready made solutions to the problem of borehole-deployable mobile robots. Snakes have inherent cylindrical configurations combined with rough terrain mobility. However, beyond the lack of demonstrated performance outside of laboratory environments, snake robots suffer from bad energetics and being very low to the ground, which while acceptable in dry environments results in cameras and range finders being blinded almost immediately in wet or muddy environments. However, the highly-refined joint designs [Shammas et al., 2003] found in many snake robots can likely serve as deployment and mobility mechanisms in borehole-deployed systems.

2.4.3 Tank Inspection Robots

Tank inspection systems are perhaps the closest analog to borehole-deployable systems. In both cases ingress and egress are limited by small diameter pipes and fittings, and once

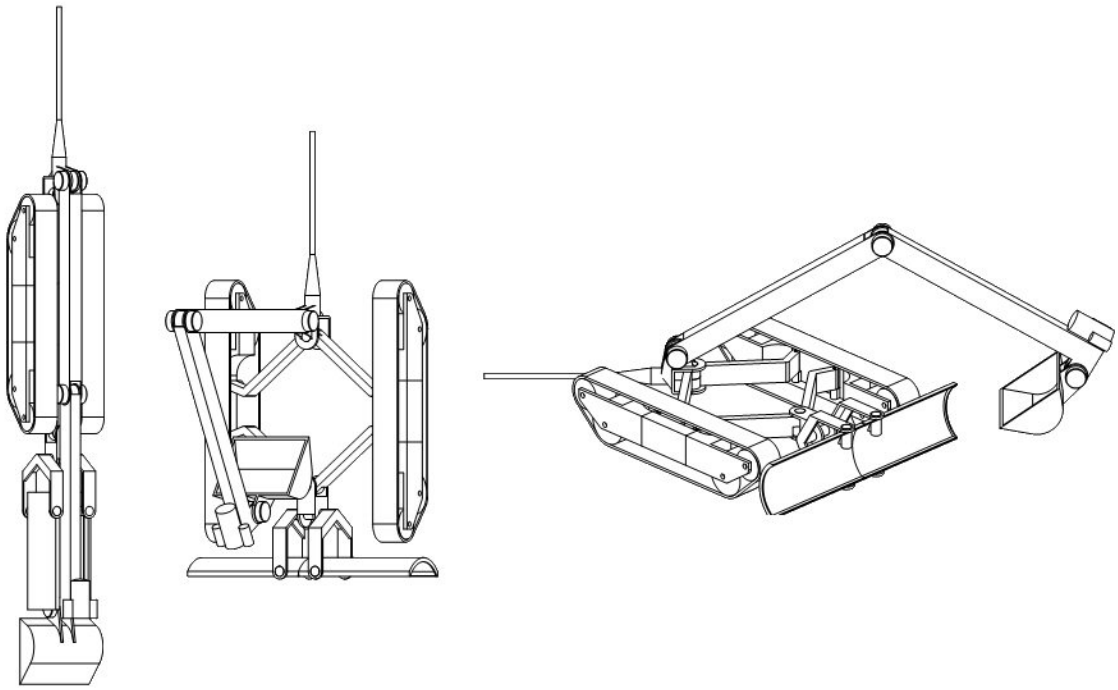

Figure 2.2: Self-reconfigurable Houdini tank inspection and cleanup robot.

ingress is achieved, the volume of operation is orders of magnitude larger than the deployment volume, and likely hostile to robotic mobility. However, most tank situations are highly uniform, with the size, variety, and frequency of obstacles likely to be encounters being less than what is expected in abandoned mines. Additionally, most tanks are fully visible from the point of ingress, thus allowing for tethered operations and improving the effectiveness of teleoperation. Finally, many tank inspection systems, such as the Neptune II fuel storage tank inspection system [Schempf et al., 1995], take advantage of tank wall magnetic properties for locomotion over the full tank surface, and are thus ill suited for operation in mine mud. One strong example of a mobile tank worker is Houdini [Schempf, 1995], a folding frame, hydraulically powered bulldozer/backhoe designed to clean radioactive material from the bottom of underground storage tanks. Houdini can fit through a 22.5" diameter pipe and then expand into a roughly three foot cube as shown in figure 2.2. However Houdini is not well suited for mine mapping as 24" diameter and larger drill rigs are rare, and the system relies on externally generated hydraulic pressure delivered via tether, which is not feasible for deep mines or for exploration beyond the first corner.

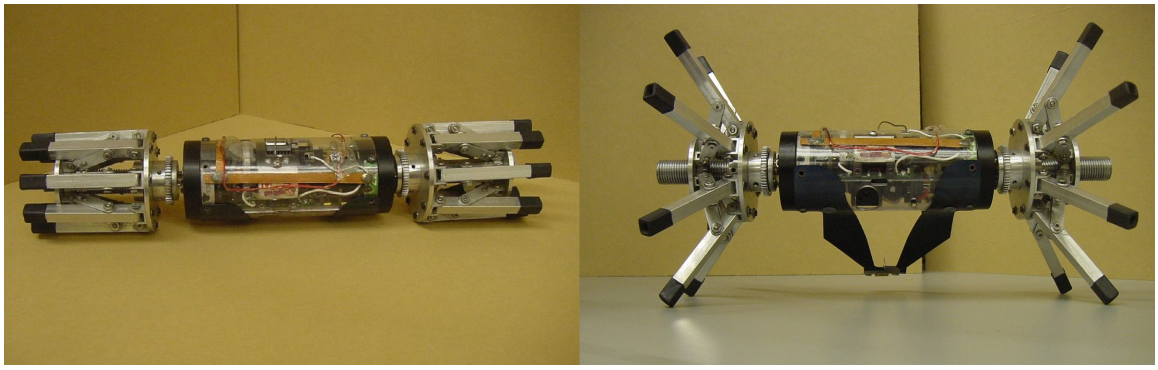

Figure 2.3: Self-reconfigurable Scout robots improve ground clearance through mechanically expanding wheels.

2.4.4 Transforming Robots

Reconfigurable systems which transform geometrically to provide locomotion or to meet required packing constraints span a broad array of applications, from planetary rovers to serpentine robots. For example, Scout robots are designed to be launched as small diameter cylinders from larger Rover robots, but then transform, via mechanically expanding wheels [Drenner et al., 2002], into differential drive systems capable of driving through rubble.

In order to fit within strict fairing dimensions, most planetary rovers travel to their destinations folded, deflated, or otherwise compacted from their operational configuration. Explosive bolts, deployment mechanisms, and inflation serve to transform the geometry of such rovers, and while explosive bolts are not appropriate for mine applications, both inflatable technologies [Jones, 2001] and deployment mechanisms [Eisen et al., 1997] taken from planetary rovers could be applied to borehole-deployable mobile systems.

Chapter 3

Subterranean Modeling

When modeling is the ultimate aim of deploying and operating robots in subterranean voids, it follows that requirements for modeling become requirements on robot configurations. Mapping two miles of underground corridor requires a robot with a two mile range. Modeling one inch diameter roof bolt heads on a mine ceiling sixteen feet high requires angular scanner resolutions and overlapping 3D scans which result in multiple range returns from each bolt head. Accurately placing a sewer pipes location hundreds of feet from the nearest man holes requires localization at least as accurate as the margin of error around the sewer pipe. Modeling a mine that is only accessible by a single small diameter hole that is over one thousand feet long is only possible if robotic configurations exist that are compatible with such size constraints and deployment depths. These and other modeling objective impact subterranean robot configuration and must be understood and accounted for in developing effective subterranean modeling robots.

Subterranean modeling tasks span a spectrum from full, dense coverage of expansive voids to limited coverage of a few critical underground points. No single robot configuration will be optimally effective for all void modeling tasks. Therefore, it is necessary to enumerate and understand different subterranean void modeling tasks and the implications these tasks have on robot configuration.

Subterranean modeling tasks can be specified by three parameters: data density, data representation, and model extent. Data density is the distance between adjacent data points, a continuously variable value with a minimum value specified. Data dimensionality is a cross of binary choices consisting of either two dimensional versus three dimensional and individual points versus connected geometries. Model extent specifies which void areas are to be modeled through bounding boxes or bounding volumes.

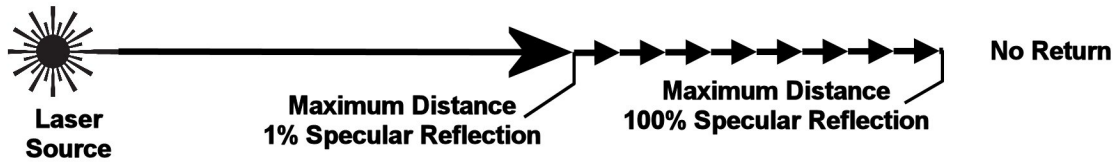

Figure 3.1: Impact of range sensor characteristics on data density in one dimension

3.1 Data Density

Subterranean robot sensing configurations are impacted by the desire to detect small features in voids. Data density is the number of range data points that strike a square area of void surface (i.e. one data point per square inch of void structure). Specifying minimum data densities ensures small features, such as bolt heads, are detectable in void models. Achieving a desired minimum data density requires an understanding of range scanner function and resolution as well as robot mobility. This section examines the relationship between data density and sensor geometry using laser range finders as a baseline. This analysis can then be extended to alternate sensing technologies if needed.

Range sensing devices such as lidar and sonar are one dimensional devices, providing range information along a line. Data density along the line of sensing is either unlimited when the device is in range of a target, or zero when out of range, as shown in figure 3.1. This range limitation must be included when calculating data density of scanned (more than one dimensional) range finders.

Two dimensional scanning range finders, sometimes referred to as line scanners, consist of a single point range finder that rotates like a lighthouse, sweeping out some sector of a circular disc. Data points are usually acquired at angular increments, as shown in figure 3.2. The angular increment between data points (α) is determined by the rotary position resolution of the scanning mechanism. The number of data points in a single scan is determined by the angular resolution and the scanner field of view, θ_{fov}. The distance between adjacent data points d is a function of scanner angular increment α, target incidence angle θ_i, and target range r. Scanners with variable angular increments are possible, but uncommon and therefore not further discussed. Applying the law of sines to the geometry shown in figure 3.2 yields the following:

$$\frac{d}{\sin\alpha} = \frac{r}{\sin\left(180-(90-\theta_i)-\alpha\right)} \tag{3.1}$$

$$d = \frac{r\sin\alpha}{\sin\left(90+\theta_i-\alpha\right)} \tag{3.2}$$

$$d = \frac{r\sin\alpha}{\cos\left(\theta_i-\alpha\right)} \tag{3.3}$$

In addition to the above equation it is also meaningful to examine the impact of a specified data density on robot operations. The update rate of the range sensor f_{rs} determines how much time acquiring each data point requires. Scanner field of view divided by angular resolution determines the number of data points acquired in one complete 2D scan, and hence the minimum time per scan t_{smin}, as shown below:

$$t_{smin} = \frac{\theta_{fov}}{\alpha f_{rs}} \tag{3.4}$$

Two common schemes for scanning are unidirectional continuous rotation and oscillatory. Both of these scanning methods require additional time to return the scan head to its home position, resulting in actual scan times larger than the specified minimum. For a continuously rotated scanning head, θ_{fov} should be set equal to 360^o to account for the time required to move the scanner head from the leading edge of the field of view to the trailing edge. For an oscillatory scanning head, time is required to decelerate and accelerate the mass of the scan head at either edge of the field of view. Equation 3.4 is therefore accurate for electrically scanned sensors with negligible transition times between sensor headings and for continuously scanned sensors with a full 360^o field of view. While scan time does not effect data density or model quality, it does effect the time required to obtain a model and hence impacts mobile robot range and static sensor cost effectiveness.

The representation for data density analysis in three dimensional scanners is a composite of two 2D representations. The three dimensional scanning pattern can be divided into a fast axis and a slow axis. The scanner sweeps through the full field of view of the fast axis while transitioning through one angular increment in the slow axis. The slow axis can be either a slow continuous motion, or a rapid incremental action.

3.2 Data Representation

Both maps and models have utility in understanding subterranean voids. Models contain three dimensional information and are best represented as point clouds, wire meshes, or polygons. 3D models are useful for examining blockages in mine corridors, analyzing sewer

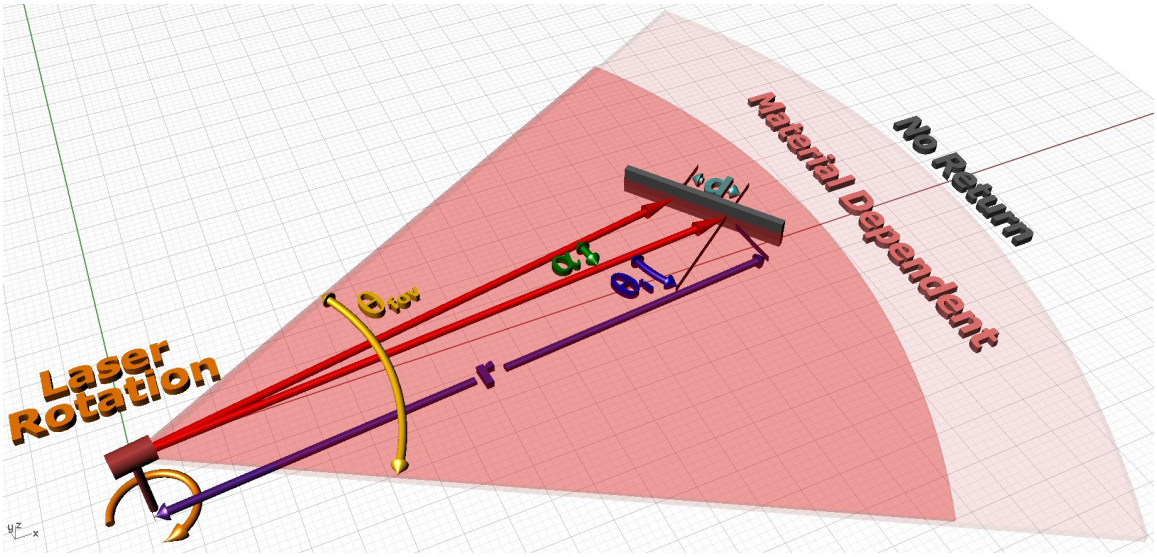

Figure 3.2: 2D range scanner representation.

pipe structural degradation, and determining damage to underground bunkers. Maps contain a single plane of two dimensional information, like a drawing, and are best represented with points or continuous line segments. Maps can be thought of as a subset of models, with a map representing a single cross sectional slice through its corresponding model. Maps are commonly used in applications such as underground mining, city sewer diagrams, and underground facility layouts.

The selection of which type of data to obtain and use is dependent on both sensing requirements imposed by robot configurations and subterranean exploration objectives. Robots designed for subterranean modeling usually utilize the same primary sensor for modeling and for obstacle avoidance, due to the superiority of modeling sensors over alternate possible navigation sensors. Therefore, even in cases where only 2D data is required by the void modeling task, 3D data might be necessary for robot navigation. Alternately, while data acquired for robot navigation might be extensive, it is often necessary to acquire additional data to maintain specified model data densities.

Representation of model data depends on modeling objective. Model data from range finders inherently consists of points rather than lines or images. Therefore, 2D and 3D point clouds are a direct, complete, and uninflated data representation. Generation of other data representations, such as 3D polygonal meshes or 2D closed polygons, involve inference, interpolation, or assumptions in connecting nearby data points together. However, polygonal representations produce better imagery that is easier for humans to interpret. Data representation does not directly impact robot configuration. However, experience has shown that some form of human accessible visualization is required to fully utilize model data.

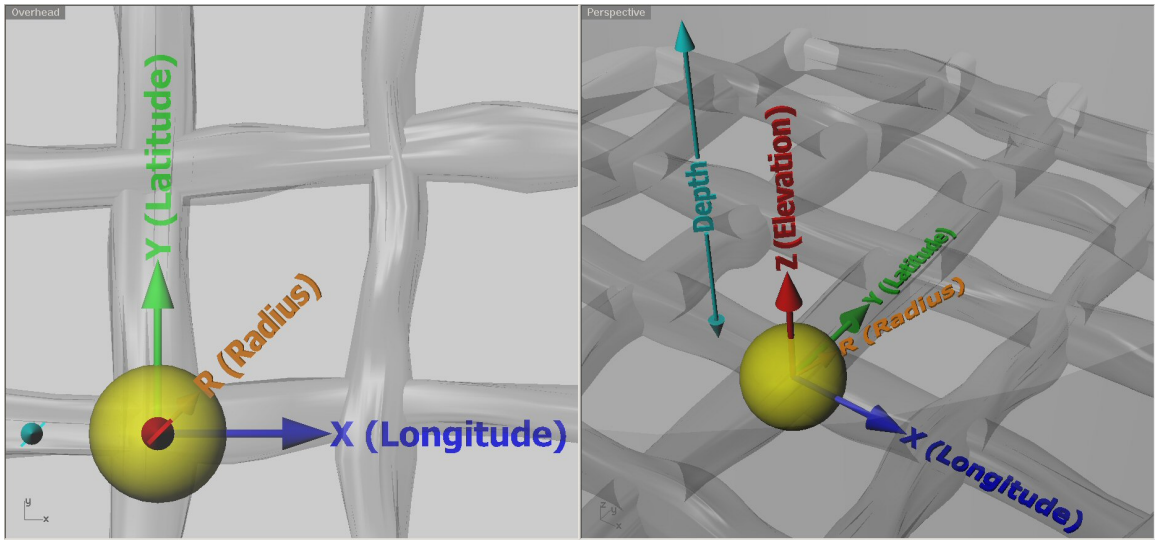

Figure 3.3: Graphical representation of local coverage.

Therefore cameras should be incorporated into robot configurations if meshing software for generating polygonal representations is unavailable.

3.3 Model Extent

While data density and data representation primarily impact the nature and characteristic of the primary modeling sensor, model extent is more tightly coupled to robot mobility and deployment configuration choices. The extent of subterranean models varies from complete coverage of square miles of underground mine to modeling a few critical square feet of a blocked sewer pipe. Four general categories of model extent are presented with associated modeling tasks in the sections that follow.

3.3.1 Local Coverage

Many subterranean modeling tasks require detailed information about a critical underground object, structure or area that resides in a larger subterranean structure. The resultant model is completely contained within a sphere centered at location X, Y, Z of radius R, as shown in figure 3.3. R is generally of the same order of magnitude as sensor range R_{sensor} and by definition less than R_{total} the radius required for the sphere to encompasses the entire subterranean void. Typical local coverage tasks include backfill verification, sewer diameter measurements, and mine structural failure analysis.

Backfilling is the process of injecting fill material, such as concrete or flyash, into an underground void, usually via boreholes. Backfilling is used to shore up failing structural

supports, isolate void sections, or plug corridor. Unfortunately the fluid nature of fill material means that precision placement of material is not possible. Additionally, many backfilling operations are conducted blindly, without any visual or quantitative feedback on the ultimate location of backfill material within the void. Usually the only feedback to surface operators is whether or not the borehole into which the backfill material is being pumped is still unobstructed. When the borehole stops accepting new material, the void is assumed to be filled. Therefore, backfill verification is often required in order to demonstrate that fill material is properly emplaced.

An abandoned underground limestone mine in Kansas City serves to demonstrate one type of backfill verification task. Developers wished to build above this mine, but the mine structure was unstable. The weak structure led to concerns about subsidence, the process by which collapse in an underground mine propagates to the surface and creates sinkage depressions. In order to obtain permits for surface construction the structure of the mine was reinforced with backfill material. Backfill verification was required in order to verify that the backfill reached fully to the the ceiling of the mine, thus providing the needed structural reinforcement. Alternate methods of backfill verification, such as direct human presence or borehole deployed cameras were deemed too expansive and would not yield any quantitative data. Therefore, robotic subterranean modeling was employed.

The required data consists of range surveys looking for gaps between the backfill and mine roof. Therefore a series of local coverage tasks were used in order to obtained the required data. Ultimately, the data obtained demonstrated insufficient backfill material and resulted in additional fill being pumped into the mine. A graphical representation of the prior map, targeted local areas and obtained data is shown in figure 3.4.

Another example of local coverage modeling is measuring sewer diameters for flow rate calibration. Over time sewer pipes fill with sediment and suffer from structural degradation. Cross-section geometry affects volumetric flow rates, which is of interest for monitoring and maintaining sewers. Volumetric flow is the product of flow velocity and flow cross section. If pipe cross sectional geometry is known, the volumetric flow is fully determined by measuring liquid level and flow velocity. Local coverage subterranean modeling can be used to determine pipe cross sections, and is thus an example of high data density, 2D linearly connected, local coverage modeling.

A final example of local coverage modeling is the generation of models for underground structural failure analysis. Fallen roof beams (Figure 3.5), breaches into mines (Figure 3.6), and blown out seals are all examples of structural failures that benefit from local modeling.

Local coverage has only minimal impact on robot configuration, as most local coverage tasks can be accomplished from a single fixed vantage point. However, most local modeling tasks require very accurate and detailed data of the small volume of void actually modeled. Therefore, local modeling requires higher data densities, such as in the sewer cross section

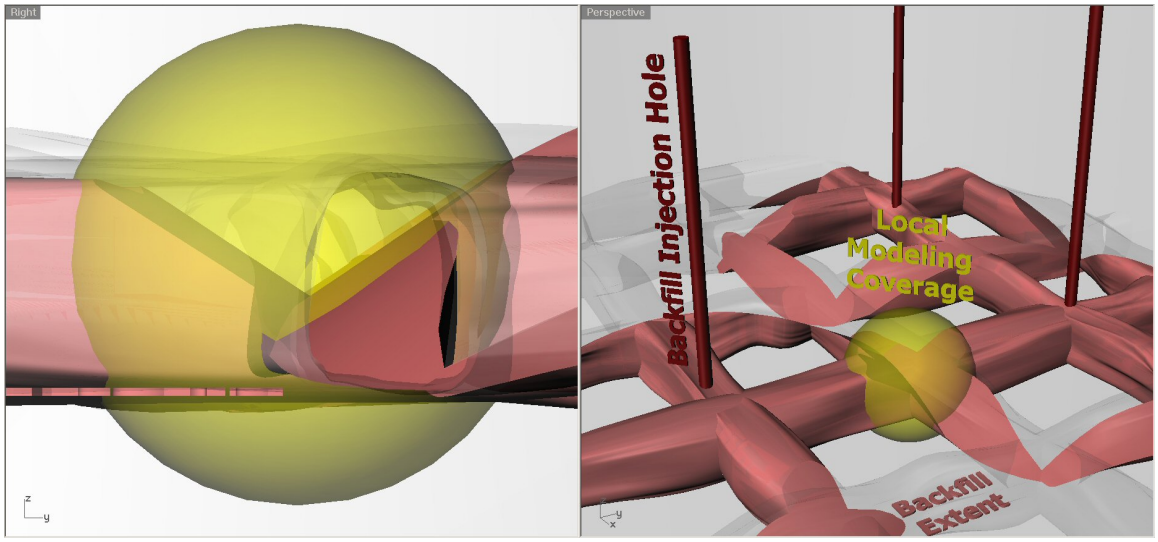

Figure 3.4: Modeling for backfill verification, an example subterranean modeling task requiring local coverage.

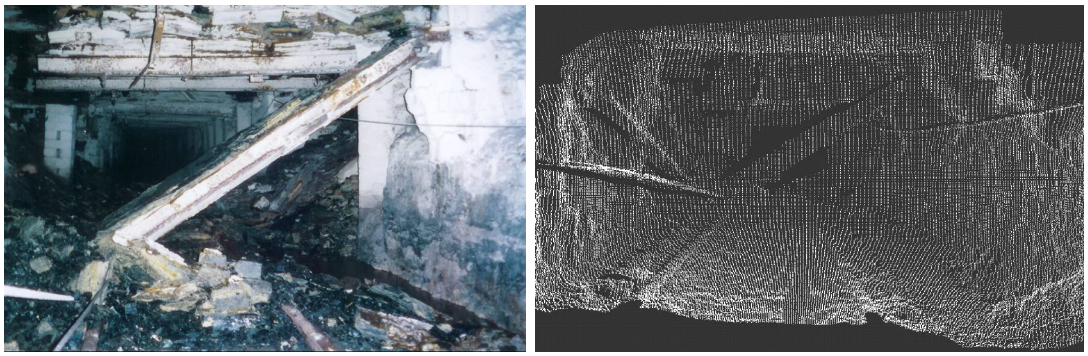

Figure 3.5: Image (left) and model (right) of a fallen roof beam in an abandoned coal mine.

Figure 3.6: Breach between Quecreek and Saxman mines through which flooding occurred.

example, and hence impacts sensor scanning mechanism design.

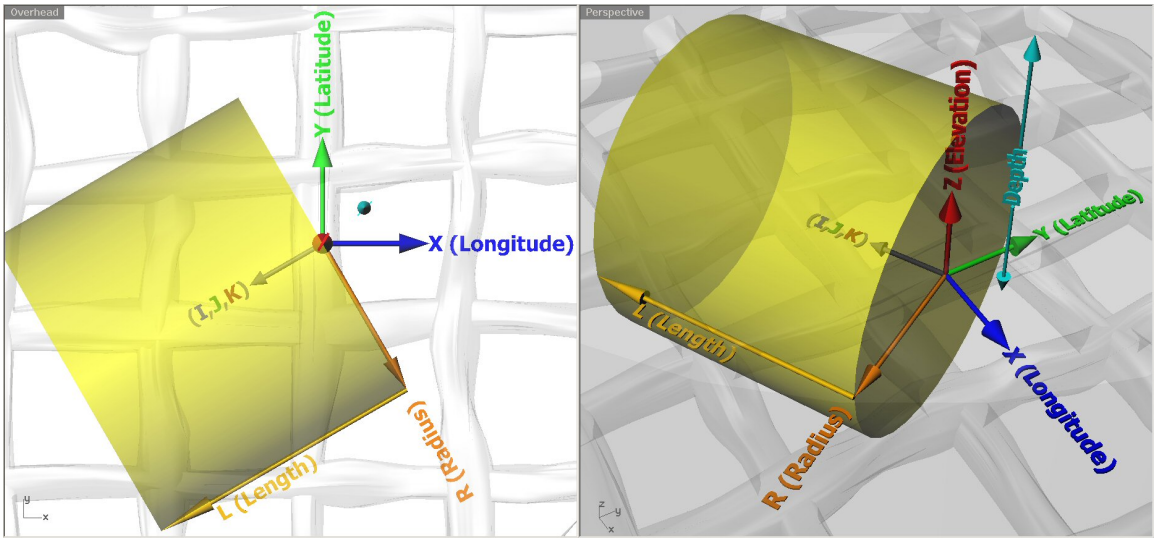

Figure 3.7: Graphical representation of edge coverage modeling.

3.3.2 Edge Coverage

Many subterranean modeling tasks are targeted toward discovering the perimeter, extents, or conditions of voids in a given direction. These modeling tasks can be described by a cylinder of radius R and length L. A unit vector I, J, K at a given location X, Y, Z can describe the direction of and start point from which to model respectively, as shown in figure 3.7. Edge coverage is utilized in modeling tasks such as measuring barrier wall thickness and determining abandoned mine proximity.

Barrier inspection tasks determine the structural and geometric condition of walls between adjacent but disconnected subterranean voids. A flooded trona mine provides a concrete example of this type of task, as shown in figure 3.8. Trona is a salt-like mineral which will dissolve into water until the water becomes saturated. Stagnant water in contact with a trona wall will absorb some amount of trona, eating into the wall, until an equilibrium is reached. Fresh or flowing water will continuously eat into the trona structure, thus setting up a situation in which walls might be eroded to the point of failure.

A trona mine in Wyoming presented such a situation, wherein the mines of two competing companies are separated by a narrow trona barrier with a width varying between 170' and 200'. Flooded abandoned portions of the two mines exist on both sides of the barrier. However, the water level on one side of the barrier is 65' above the mine floor on the other side. Therefore, a failure of the barrier wall would result in an equalization of water levels and a flooding of the lower mine by approximately half a billion gallons of water contained in the upper mine. The data desired by the mine operators and regulatory agencies is the thickness of the barrier wall. An edge based inspection along the length of the barrier, in

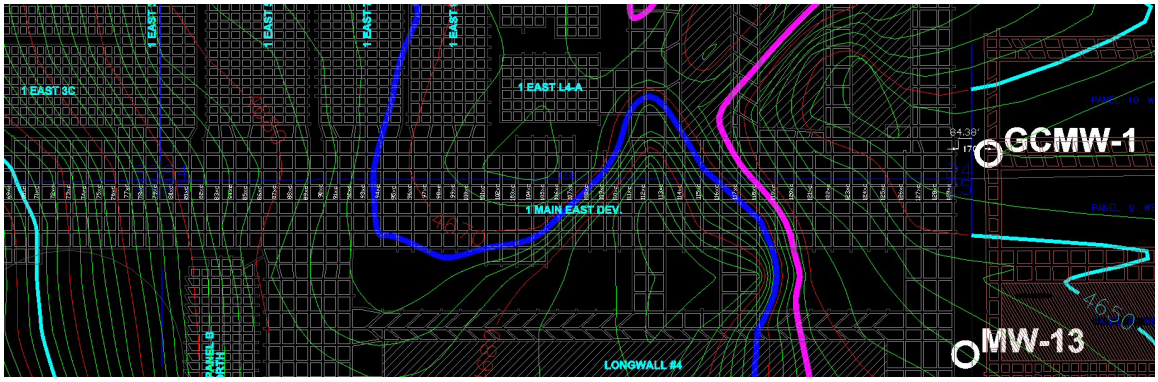

Figure 3.8: Barrier inspection task via borehole GCMW-1, an edge based modeling task.

the direction of the abandoned mines is therefore required.

The only access into the flooded portion of the 1500 foot deep lower mine is via a single four inch diameter directional drilled borehole. Due to the depth, material, and placement accuracy required, this single borehole cost upwards of $100,000. While the hole depth and cost do not directly effect the modeling task, they serve to motivate the usage of mobile modeling platforms and provide calibration as to the value of accurate modeling data.

Determining the proximity of abandoned mines to active or planned mines is important in order to ensure miner safety and maximize material removal in the active mine. Figure 3.9 shows a portion of the Quecreek and Saxman mines. The Saxman mine map shown was later demonstrated to be inaccurate. An edge based modeling task, as shown, could have determined the true extents of the abandoned Saxman mine, potentially preventing the breach that trapped nine miners for three days. Proximity detection requires exploration of voids with unknown topology in addition to the unknown geometry encountered in barrier inspection tasks. In a sense, barrier inspection is a simplified subclass of proximity detection.

The Quecreek accident [Directions Magazine , 2003] is a typical motivation for accurately mapping abandoned mines. Nine miners were trapped when they accidentally breached into an adjacent abandoned mine. Existing mine maps showed the perimeter of the abandoned mine to be more than 200 feet from the active mine (as required by law), when in fact they were only a few feet from millions of gallons of water contained in the adjacent Saxman abandoned mine. The water in the abandoned mine flooded into the active mine, trapping the miners for three days before they were rescued. The ensuing investigations determined that the existing mine map used during the permitting process for the Quecreek mine were not the last version of map generated prior to the closing of the Saxman mine. A different version of the map obtained from the family of the surveyor in the process of the investigation showed the correct extent of the abandoned mine, including the point at which the breach occurred. In order to prevent similar events from occurring in the future, the

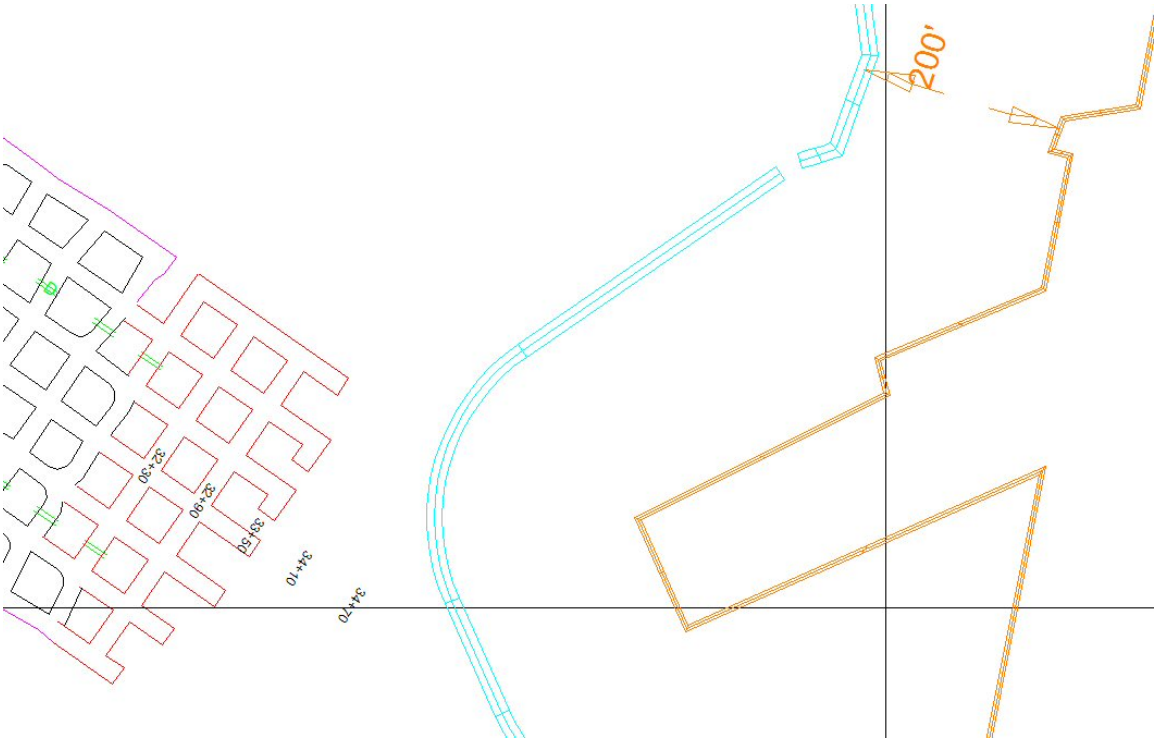

Figure 3.9: Quecreek active mine (grid, left), Saxman abandoned mine perimeter (right) and 200 foot safety barrier (center).

accuracy and reliability of abandoned mine maps must be improved, and since conditions within abandoned mines preclude human access, robots are a superior means of providing such maps.

Unlike local modeling tasks, edge coverage tasks usually require multiple vantage points. Therefore either multiple entries into the void or mobile modeling platforms are necessary. Additionally, in edge coverage tasks the extent of required coverage often changes during modeling. If unexpected void regions are discovered in the direction of exploration then additional modeling into the newly discovered void volume is required. A progression of vantage points as part of an edge based exploration task could have determined the true extent of the Saxman mine, shown in figure 3.10. The locations and density of vantage points is determined by the required data density and the topology of the void being explored. Moving from one vantage point to the next can be accomplished either by a mobile robot, by creating new axis points into the void by drilling additional boreholes, or some combination of the two. The progression of scanner locations is dependent on scanner characteristics and desired model quality, but not on the method used to move the scanner from point to point.

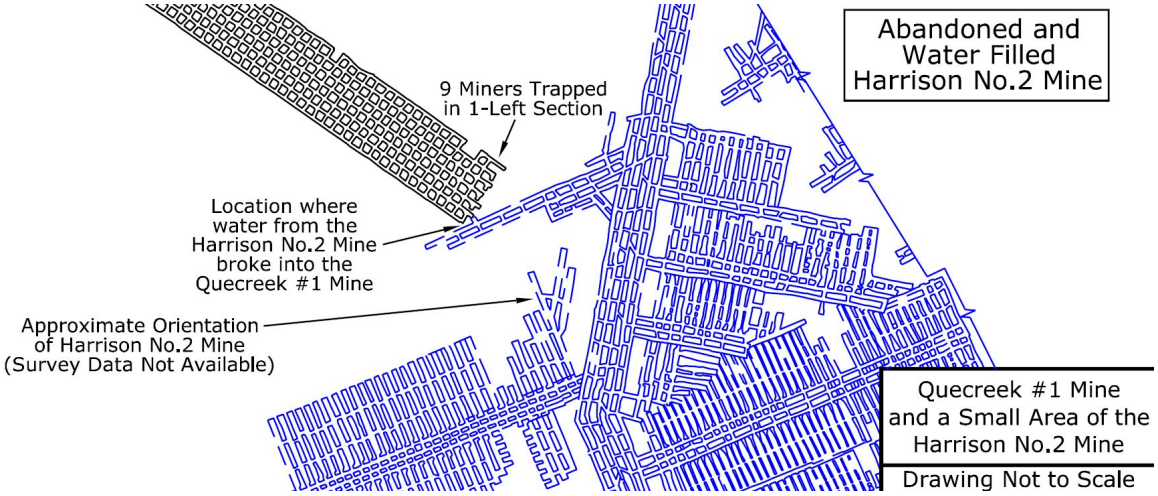

Figure 3.10: Intersection of Quecreek and Saxman mines.

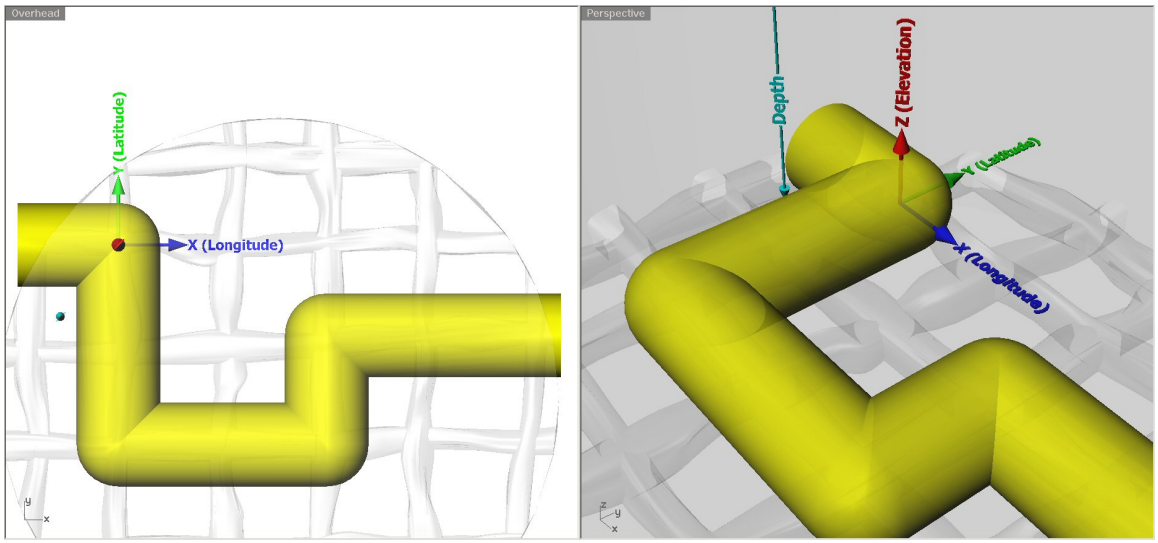

Figure 3.11: Graphical representation of a path-based modeling task.

3.3.3 Path Coverage

Finding a clear path between two points in an underground void is another common subterranean modeling task. The task can be defined using a piecewise linear path and a circle of radius R swept along that path, as shown in figure 3.11. In most path coverage tasks the desire is to provide complete information about the structural conditions along the path. Common path-based tasks include sewer pipe inspections, modeling of abandoned mines for use as haulageways, and determining safe paths for rescue workers after underground accidents.

Path modeling inside sewers serves two purposes. 3D models of the sewer can be used to determine the structural conditions along the length of the sewer pipe. This enables targeted repairs to the most damaged or degraded portions of the sewer as well as more informed preventative maintenance decisions. 2D maps of the sewer pipe provide improved underground operations adjacent to sewer pipes. Placing new pilings for building foundations or connecting a new lateral into the mapped and modeled sewer main are both improved with accurate mapping of sewer pipes.

Abandoned mines are sometime partially reopened in order to improve operations in adjacent active mines. For example, the primary passageway of an abandoned mine might be reopened to serve as a haulage way for transporting coal between an active mine and a processing plant. Alternately, piping might be installed in the abandoned mine in order to facilitate water treatment of an adjacent active mine.

Both of these situations were present at the Matties coal mine outside of Pittsburgh, Pennsylvania. A route through an old mine original used as a haulage way was later desired to be used for transporting acidic mine water for treatment. The conditions of the primary passageway were unknown and regulatory agencies wished to know the conditions prior to installation of water treatment pipes. While the abandoned mine was extensive (covering square miles), only information about the short (a few thousand feet) central passageway was desired, thus a path based modeling task was performed. The passable portions of this corridor were modeled and mapped using the mobile robot Groundhog, as discussed in section 6.1.5.

Accidents and disasters within underground mines often result in miners being trapped deep underground and miles away from the nearest portal. Mine rescue teams are specially trained to advance into mines after accidents, finding or establishing a safe path from portal to trapped miners. However, this process is very slow and hindered by lack of knowledge about the structural and environmental conditions along the planned rescue path. Therefore subterranean modeling of rescue paths could speed rescue of trapped miners and reduce the risks faced by mine rescue teams. A detailed, specific example of this type of modeling task is outlined in section 4.2.3.

3.3.4 Complete Coverage

As its name suggest, complete coverage provides comprehensive modeling of the entirety of subterranean voids. As such it is the most time and resource intensive of modeling tasks and usually only used when volumetric information about the underground void is required. Some modeling tasks fall into both the complete coverage category and one of the preceding categories, as an underground void sometime consists of only one path or a single small area of interest. However, while the preceding coverage patterns can be specified by

vectors and radii, complete coverage needs no further specification. Complete coverage is useful for modeling damage to underground bunkers, determining mineral volumes removed in mining, and in analyzing underground tank contents.

Modeling of underground bunkers before and after they have been breached enables an understanding of the forces and damage involved during collapse. Such an analysis requires accurate accounting of void volumes before and after collapse. Detailed modeling identifying both features left unchanged by the collapse and features altered during collapse is also useful in identifying exactly what happens when a subterranean bunker is compromised. Such modeling requires numerous vantage points, as the random and cluttered nature of bunkers post-breach results in significant occlusions for most scanner locations. This type of modeling has been conducted as part of this research with both static sensors and mobile robots.

Accurate estimation of mineral volume extraction is of interest to numerous parties, including mine operators, government regulators, and private property owners. Mine operators want to track material extraction rates from active mines and estimate the volume of material extracted from an abandoned mine before reopening or strip mining over said abandoned mine. Government regulators are interested in accurately assessing taxes against mineral removal and ensuring that mine operators accurately report their extraction numbers. Private property owners are often taxed based on the minerals below their surface properties. Demonstrating that all mineable minerals have been removed through full coverage modeling could be used to eliminate such taxes. All of these parties are interested in determining the volume of mineral removed and remaining underground. Accurate volumetric analysis requires full coverage 3D models of the mine in question, but can be accomplished with relatively low data densities.

Accurate, comprehensive, detailed subterranean models can reduce the risk entailed in reopening abandoned mines. Analyzing an abandoned mine in order to assess the cost-effectiveness of reopening requires full coverage and high data density since the structural conditions of the mine are of interest. As mineral prices fluctuate over years and decades, some mines will be abandoned before they are fully mined when prices are low, only to be reopened when prices go up. The point at which it becomes economical to reopen an abandoned mine is a combination of current mineral price and the expense involved in reopening the mine.

Underground storage tanks are commonly used for storing fuel, chemicals, and nuclear waste. Geometrically and topologically tanks are one of the simplest types of underground void. For tanks containing liquid, simple level sensors can determine the material level within a tank if the sediment layer on the bottom of the tank is accurately modeled. For solid materials, accurate modeling of the top surface of the material provides sufficient information for estimating volume.

3.4 Composite Model Integration

Creating useful subterranean models requires two key procedures; data acquisition, and composite model integration. Data acquisition tasks can be specified by data density, data dimensionality and model extents as outlined in the prior sections. Model integration is specified by local error bounds and global error bounds. Model integration requires localization, scan matching algorithms, and global referencing. Definition, impact on model error and impact on robot configuration of each of these three model integration tasks is examine in the following sections.

3.4.1 Localization

Stitching disparate laser scans together and globally orienting models both require localization. Localization is the process whereby a robot or sensor determines its position and orientation in the world. Techniques for surface robot pose estimation are well known, but the sensors, data and operational approaches for subterranean robots differ somewhat from surface systems. This thesis leverages these existing localization techniques and algorithms with only minor improvements or innovations where appropriate. A brief introduction to localization with emphasis on its impacts to robot configuration is presented below in order to make clear the role of localization in subterranean modeling.

On the surface, each range data point a robot acquires can be tagged with GPS position information, thus establishing the global position of that data point. With all data points referenced to the same global frame, no additional localization is necessary to either stitch together scans or translate and orient models into global coordinates. Underground, GPS information is unavailable thus requiring alternate localization techniques. Additionally, approaches to stitching scans together and globally aligning models become potentially separate problems, requiring different methods of localization. Localization between scans is necessary in order to create self-consistent, error free models. Registration of the self-consistent model to known points in the world is required in order to accurately model subterranean voids in world-relevant coordinates.

Localization information can be provided from three potential sources. Inertial measurements, integrated over time, can provide estimates to robot motion between scans. Odometry information provides similar information to inertial measurements. Finally, measurements from a fixed sensor to each scanner vantage point can provide accurate localization of multiple scans. Scan matching, which uses statistical methods to determine the alignment and offset between two scans, can be used to provide localization information, but due to its integral nature in the model building process it is treated separately in a section below.

Inertial measurements are taken using accelerometers, which provide linear acceleration measurements, and gyros, which provide rotation velocity measurements. Double integra-

tion of three orthogonal accelerometers fused with single integration of three orthogonal gyros results in positional information which can be used to track the complete pose of a robot as it moves. Inertial Measurement Units (IMUs) have well known problems with drift, which occurs due to small biases in inertial sensor readings being compounded into larger errors through integration. Drift compounds over time, and IMU cost, size and power requirements increase dramatical in order to reduce drift rates. However, IMU data is useful in determining the relative orientation between two scans acquired within a short period of time.

Inertial data has utility beyond the positional information provide by an IMU. Accelerometer data taken from a static platform provides a gravity vector which can be used to fix two axes (roll and pitch) of scan orientation. Unlike position estimation from accelerometers, gravity vector determination does not involve any integration and therefore has a constant error bound. Similarly, north seeking gyros mounted to a static platform can use repeated measurements of a gyroscope in multiple orientations to determine the direction of rotation of the Earth. Knowing the direction of the Earth's rotation, coupled with the knowledge of a gravity vector fully defines the orientation of a scan. Magnetometers can be used to determine magnetic north, thus providing data similar to that provided by a north seeking gyro. While magnetometers are effected by local magnetic fields, either on the robot or in the surrounding subterranean structure, their significantly lower cost and complexity compared to north seeking gyros makes them the preferred choice for providing supplemental yaw information.

Odometry information can provide a direct measurement of robot attempted motion. By placing position sensors on steering and locomotion motions it is possible to record all motions that the wheels or tracks of a robot go through. However, odometry information fails to account for slipping, sliding and other disturbances introduced by the environment. Odometry errors compound over distance in a similar fashion to how inertial errors compound over time. Therefore, odometry can be useful in determining relative orientation between two scans taken a short distance apart.

Tethers used for providing power and/or communications to a robot can also be used to provide odometry information. For static sensors deployed by being lowered into a borehole, tether payout provides an accurate measurement of scanner depth, thus locking down the z location of a scan. For mobile systems tether payout can either occur at the robot or at a base station at the void entry. In an environment that doesn't impact tether location it is theoretically possible to determine robot position by measure tether payout length and direction of tether pull. However, for such measurements to be accurate the tether must be under tension. This requires that the robot continuously expend tractive force in order to maintain tether tension. Additionally, tether elongation is possible, and if the tether is suspended at any point tether tension impacts tether length. In complex environments

where the tether is no longer in a straight line between base station and robot the situation is even worse, as different segments of the tether might be under different tensions and it is no longer possible to determine the angle between the base station and robot. With all this being said, tether payout measurements can still contribute to an estimate of robot location. When a robot is moving in a straight line, relative tether payout measurements can be used to set the distance between two adjacent scans. Tether based odometry has been most successfully employed underground for borehole deployed devices and in sewers where the linear, unobstructed nature of the environment minimizes impact to tether measurement accuracies.

The types of localization information discussed so far are all common within the robotics trade, but are largely unheard of in conventional underground surveying and mapping. Manually generated maps of subterranean voids, which can be thought of as very low data density subterranean models, are usually acquired using well established surveying techniques and equipment. Theodolite survey instrument provide precise azimuth and elevation, which determine direction toward a survey target. Total stations automate the reading and recording of theodolite angles and add laser range finders in order to precisely determine distance as well as direction. Robotic total stations automate the motions of a total station, enabling the tracking of a moving survey target while maintaining the high precision measurements of a basic theodolite.

Once the position of a robotic total stations is established at a portal, it can be used to track the position of a prism mounted to a mobile robot as it progresses through an underground void. The robotic total station has bounded and fixed angular and range errors, thus range errors remain constant while lateral position errors increase linearly with range. The data from the robotic total station determines the x, y, z position of the robot mounted prism and hence the scan, but does not provide any information about scan orientation. Having two target prisms, or actuating a single prism in on dimension provides sufficient information to establish scan location completely and two axes of scan orientation. Three target prisms, or actuating a single prism in two dimensions (i.e. rotation in a circle) provides sufficient information to fully determine the pose of a scan.

Using surveying equipment to localize scans is only effective as long as line of sight between the robot and survey instrument is maintained. After the first corner or dip in the void, survey-based localization is no longer possible. Human surveyors combat this problem by setting up a survey traverse, wherein the survey instrument and target positions are switched. This process is repeated with new target locations until all relevant areas have been surveyed. Two robots, both equipped with survey instruments and prisms, can perform a similar traverse, leapfrogging each others position to achieve full void modeling. Robot A starts as a static surveying instrument while robot B explores and models a subterranean void until robot B reaches an intersection or the limits of line of sight range. Then the

robots switch roles, with A following a similar route to that taken by B, eventually reaching the location of the now static B. A then leapfrogs B and continues to explore deeper into the void. The complexity, cost and additional software requirements of such a system are large, but in return the localization information acquired is comparable to that achieved and accepted by human surveying practices. Additionally, surveying techniques for propagating and correcting errors in closed loop traverses can be applied to further improve the accuracy of localization.

3.4.2 Scan Matching

Scan matching can provide localization information, decrease the error in models, or do both simultaneously. Scan matching is the process by which two data models are registered to each other until the error between the overlapping portions of the two models is minimized. Scan matching stitches 2D scans into maps and aligns and orients 3D point clouds in full six dimensional space into models.

Scan matching provides localization information by providing an estimate of the translation and rotation between two scans. If scan matching is only applied to scans being integrated into the final subterranean model, then scan matching and model integration are one in the same. However, if additional intermediate scans exist, for example scans used for navigation but not modeling, scan matching between these non-model scans can provide additional localization information. This additional localization information obtained from non-model scans can be used to better establish the pose between two scans integrated into the final model.

Scan matching, whether for localization or model integration, can be performed with one of two primary paradigms. Scan matching can either be performed pairwise between two scans, or between a scan and a composite model. Consider three scans, A,B and C with scan A and B and scans B and C overlapping significantly, but scans A and C only overlapping slightly (see figure 3.12). Additionally, scan A's location is accurately known. Scan matching can be performed between any two scans, but will only be reliable and effective if there is significant overlap between scans. Pairwise scan matching between scans would match A to B, then match B to C, with C's location relative to A being the composite of the two transforms thus determined. Composite scan matching would match A to B, then merge the two scans into a new scan or model, AB. C would then be matched to the composite model, AB, rather than simply to B. The resulting transform between A and C is expected to be better when the composite model is used because the small area of overlap between A and C is taken into account in addition to the $B - C$ overlap.

Scan matching is the final step applied to individual scans in order to create a composite model. Scan matching is an optimization problem that minimizes the error between two

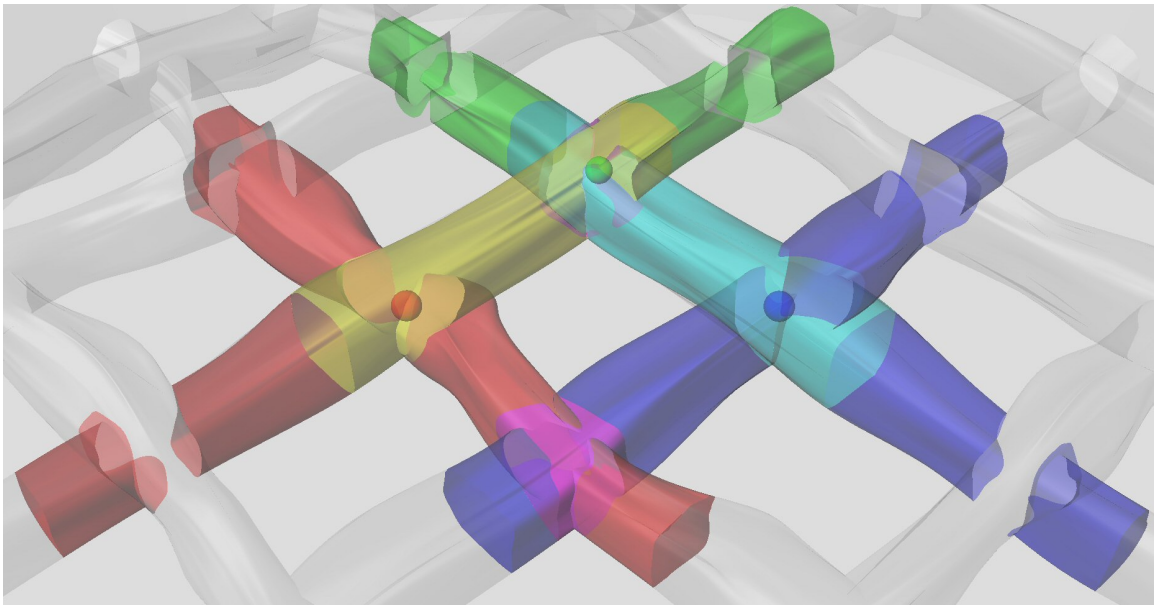

Figure 3.12: Three overlapping modeling scans.

scans. However, scan matching algorithms, such as iterative closest point, often converge to local rather than global minima. In order to improve the probability of converging to the global minimum error, scan matching requires an accurate initial guess as to the transformation between two scan or a scan and a model. The more accurate the initial guess, the better the resulting scan matching will be. Therefore, the first step to integrating multiple scans is to transform the location and orientation of each scan relative to an initial scan based on available localization information.

Localization information can be uniquely-defined, under-defined, or over-defined. When localization is uniquely defined, the initial pose of each scan is determined directly from the localization information. Under-defined localization information means that one or more of the scan degrees of freedom is unknown. For example, referring back to the previous section and the example of a robot with a single prism being tracked by a survey instrument, only the three degrees of freedom associated with scan location are determined, with orientation unknown. When localization is under-defined, scan overlap is essential so that adjacent scans can be properly oriented and placed relative to each other. In both the uniquely-defined and under-defined cases the localization information is applied directly to the scan without further processing.

When the pose of a scan is over-defined, additional processing of the localization data is required in order to reduce the available data to a single homogeneous transformation for each scan. Combining multiple measurements of different accuracy into one measurement of higher accuracy is accomplished using a Kalman filter. Once the localization information

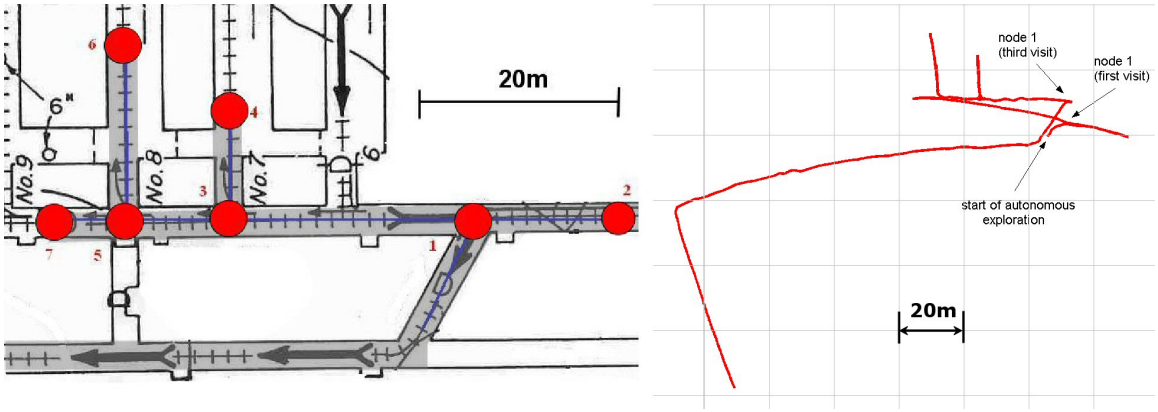

Figure 3.13: How model errors compound and can be corrected during scan matching.

is merged into a single pose estimate, scan matching can proceed.

Beyond providing an initial estimate of scan poses, localization information also impacts how convergence conditions for the scan matching process are handled. For example, if the localization information is very reliable, then bounds should be placed on what possible positions the scan matching algorithm can specify. These bounds can take the form of maximum offset or rotation from the scan pose specified by the localization information. Additionally, the bounds might vary based on scan, with some scan poses known more accurately than others based on external known reference positions.

Integrated into the scan matching step is the concept of error correction. Some scan positions, such as a scan that occurs at the bottom of a borehole, provide precise information that locks down the scan location. Consider a sequential series of scans, A,B,...,M,N with the poses of scan A and N accurately known while the intermediate scan poses are not. Scan matching starts at with the integration of scan B into the composite model that initially consists of only scan A. The modeled position of scan B relative to A has a small amount of error relative to the real world. As each new scan is integrated into the composite model, incremental errors compound. When integrating scan N into the composite model, a large difference between the scan matched location of N and the actual, known location of N is observed (as shown in figure 3.13). On possible way to correct this error is by reorienting the entire composite model to minimize the error between scans M and N.

Traditional surveying techniques provide a better means of applying error corrections. A common type of survey is a closed loop traverse. In such traverses the last survey point "closes" on the first survey point to constitute a "loop". Error accumulated in moving the survey instrument from station to station can be corrected using the known coincidence of the first and last points, with error corrections being propagated through the chain of survey points based on proportional allocation of the error correction. A similar procedure

can be used to correct the errors in a series of scans which describes a loop, with the first and last scan being of the same location.

The end result of scan matching is the generation of a composite underground void model. The error within the model can be minimized through effective use of all available localization information in addition to data contained within individual scans. While a self consistent model describes the geometry and topology of a subterranean void, its utility is improved if the model is referenced to some global coordinate system.

3.4.3 Global Referencing

Global referencing ties a subterranean model to points and coordinates on the surface. The most common and universal example of global coordinates are those provided by GPS systems. However, most countries and states also have their own vestigial coordinate systems with various origins, corrections, and projections to allow 2D maps to account for the spherical nature of the Earth. An integrated model can be transformed into a useful coordinate system through accurately locating and aligning void entry points to their measured positions on the surface, orienting to gravitation and magnetic fields produced by the earth, or referencing to artificially generated external signals.

Both portals and boreholes can be used to convey global referencing downhole to an integrated model. Scanning and modeling the area around and in a portal provides model data of points that can be reached by human without requiring any underground operations. The location of distinct points within the portal area model can be survey in, thus providing very accurate global locations for these points. Locking down the full pose of portal area scans means that all subsequent scans which are matched to the initial portal scan are locked to global coordinates. The only flaw with such referencing is that minor errors in the locations of the distinct points can result in orientation errors of the original referenced scan. With only the portal areas serving to globally reference the entire scans, minor angular offsets at the portal are magnified into significant positional inaccuracies at the points in the model furthest from the portal, as shown in figure 3.14.

Borehole locations can be accurately surveyed on the surface. When combined with accurate depth measurements, a borehole can provide an effective global position reference for an integrated model. However, the idealized view of a borehole as a perfectly vertical cylinder is rarely the case. Changing materials, soil properties, and drilling conditions result in lateral offsets between the top and bottom of a borehole. These offsets tend to grow linearly with depth. Advanced drilling techniques, such as directional drilling, which actively guide the path of a drill bit can virtually eliminate such offsets, but at significantly increased cost compared to conventional drilling. In addition to the unknown position error, a single borehole fails to provide complete orientation information, as rotations about the

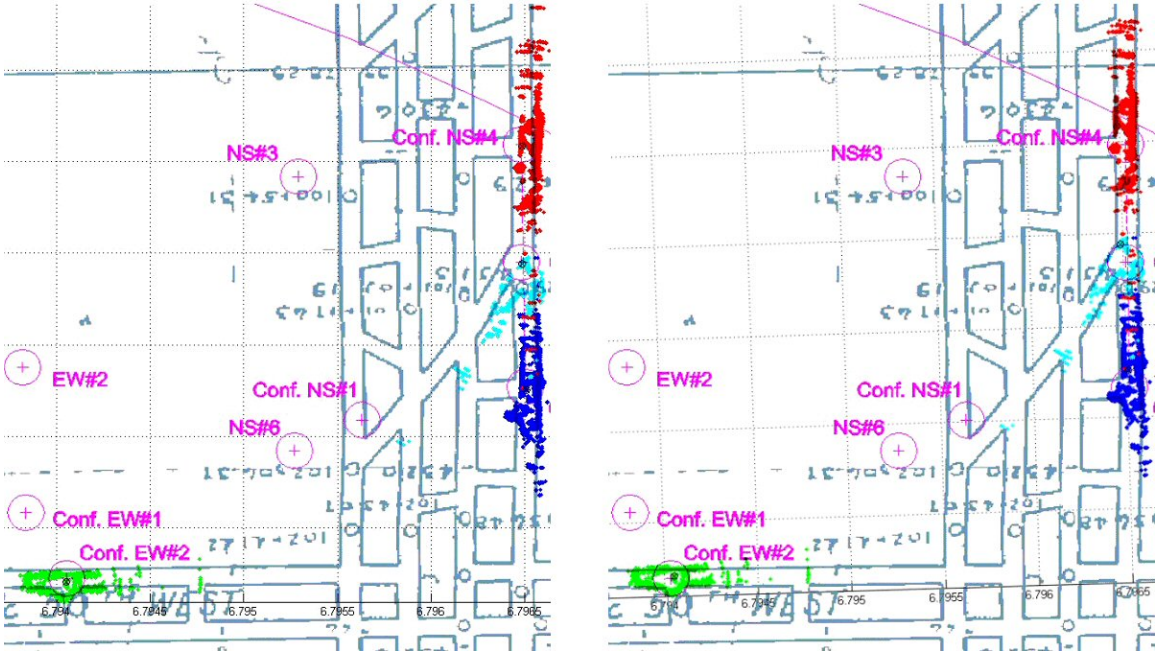

Figure 3.14: Referencing models using boreholes: Unreferenced model (left), aligned and referenced model (right)

borehole axis is unconstrained. Multiple boreholes connected to a single integrated model provide complete global referencing for the model, as the two borehole locations uniquely determine model orientation. Boreholes can also be combined with portals in order to correct the types of angular offsets described in the previous paragraph.

The process of gathering data, integrating a model, correcting model errors, and transforming the model into a global reference frame completes the subterranean modeling cycle. However, sometimes the environment changes, thus requiring additional modeling processing in order to understanding an evolving subterranean void.

3.4.4 Dynamic Modeling

While ideal subterranean environments are completely static during modeling, real-world environments are often dynamic. Subterranean environments can be altered by human actions, evolve over time, or changed by the presence of modeling robots. It is frequently of interest to know how an underground void has changed or if it has remained unchanged. Additionally, effective modeling requires dynamic changes in voids be detected, so these changes can be accounted for in models.

Subterranean modeling robots often corrupt models through self-imaging. For example, a mobile robot might model a survey station tracking its progress, or a borehole explorer might model its own tether. Creating accesses into voids can also alter subterranean en-

vironments. For example, the pile of cuttings generated during borehole drilling might obscure the actual void floor, changing data acquired by a borehole sensor. Range data acquired from objects associated with a modeling system rather than the void must be eliminated from void models. Sometimes portions of a void occluded by the modeling system can be captured in other scans. If the model contamination is limited in extent, as with an unintentionally modeled survey instrument, simply removing all data near the instrument is feasible. In the case of borehole cutting obscuring the void floor, recovery is impossible after borehole drilling and the cuttings simply have to be treated as part of the model. However, coverage of void floor from an alternate borehole prior to drilling might be possible, and both the original floor and cutting pile have validity as part of the model. The more complex case of modeling of a robot tether requires a more advanced means of correcting the model, as the tether moves with the robot, resulting in different scans placing the tether in different locations.

Multiple overlapping scans combined with nearest neighbor matching can be used to determine which parts of a model are truly dynamic in nature. Nearest neighbor comparisons find the minimum distance between a point P_A in scan A and all the points in scan B. If this minimum distance is above some threshold, P_A is said to exist only in scan A and not in scan B. If A and B are overlapping in the region around P_A then P_A represents some dynamic portion of the environment. In the example of a moving tether being accidentally modeled, both the tether and void surfaces obscured by the tether will be marked as dynamic in both scans A and B, resulting in four regions of scan data marked as dynamic. However, only the two regions modeling the tether are actually dynamic. This error is corrected with the introduction of a third overlapping scan C. Areas of the void obscured by the tether in one scan will be modeled in the other two scans. Thus area obscured by tether in A will be modeled in both B and C. Therefore, the filtering algorithm keeps points if they have nearest neighbors in at least one other scan. As additional overlapping scans are added, this effect is magnified, resulting in the tether being removed from the scan, while shadows from the tether are filled in.

The above nearest neighbor comparison procedure can also be used to identify changes over time between composite models. For example, a composite model of points that don't have nearest neighbors between two models acquired at different times in an active mine would show only the material removed between the two scans, and thus provide an accurate means of calculating the volume of mineral removed.

3.5 Summary

Effective subterranean modeling requires sensors and algorithms acting in concert. Modeling sensors must provide accurate, dense, extensive coverage. Localization algorithms

must make use of available data from any source while minimizing drift and detecting loop closures. The data desired must be both acquired and accurately transformed into relevant coordinate systems. All of these elements and more must exist and be integrated into a composite robot configuration and operational approach in order to successfully model subterranean voids.

The robot configurations and approaches to configuring robots outlined in the next sections were developed with the modeling process in mind. The forthcoming design decisions are based on knowledge of the modeling deliverables that are the ultimate purpose of these robots.

Chapter 4

Configuration Classes

While impossible to determine in practice, each void modeling challenge has an associated robot configuration that is unique and optimal. It is impossible to determine the challenge a priori, and it is impractical to specialize a unique robot for every void modeling task. Prior to exploration the true extent and difficulty of a modeling task is unknown. Therefore, determination of an optimal configuration is only possible for trivial modeling tasks where complete information is known a priori. Pragmatically, even if ideal configurations could be specified, it is not possible to specialize every robot for each modeling objective. If the desire were to solve a *single* subterranean modeling problem, then a single optimal robot configuration could conceivably be pursued. However, this thesis embarks to address *many* reasonable subterranean modeling tasks with few robot configuration classes. The challenge is to determine and develop a minimum spanning set of robot configurations which can accomplish the majority of required objectives. This spanning set of robots could theoretically consist of a single ultimate configuration or hundreds of configurations, each only effective for a narrow set of modeling tasks.

A spanning set of three configuration classes is warranted by three reinforcing considerations; spatial, operational and efficiency. *Spatial* arguments consider access constraints, sensor range and robot mobility (as seen in figure 4.1) to distinguish three different robot configuration classes. *Operational* considerations independently support the same three configuration classes: Borehole explorers are best for model beneath impoundments. Borehole sensors excel when modeling voids with traversable floors. Portal explorers are uniquely suited to modeling for mine disaster response. *Efficiency* concerns demonstrate that each configuration class provides superior cost effectiveness in a given subset of the span of modeling tasks.

Modeling capabilities relate to void access and mobility. No robot configuration can provide complete coverage of a void extending over square miles from a single one inch diameter borehole. Unfeasible, outlier modeling tasks such as this are not considered in this

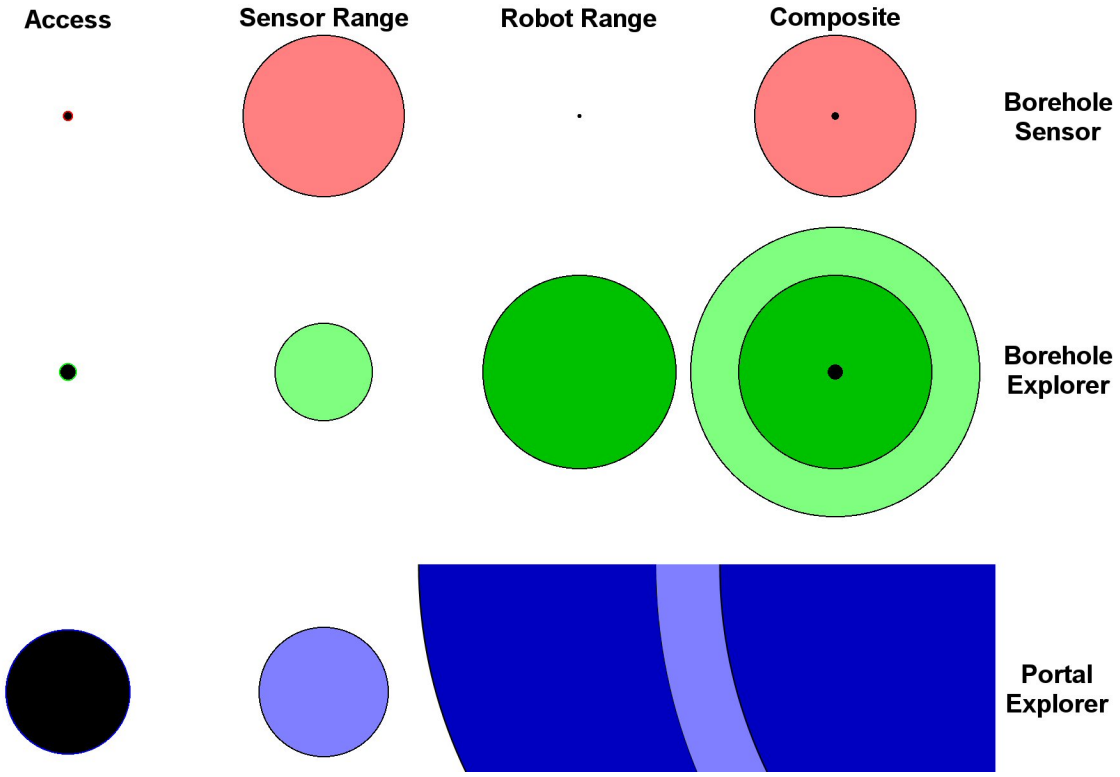

Figure 4.1: Representative access constraints, sensor ranges, and robot mobility.

research as they are beyond the capability of current component technologies, regardless of configuration ambitions.

This thesis distinguishes three robot configuration classes which address the majority of subterranean modeling challenges. Demonstration of subterranean modeling capability requires both identification of essential generalized configuration classes and development and testing of specific robot configurations within each class. This chapter distinguishes spanning configuration classes by examining contrived but comprehensive modeling scenarios, specific real-world subterranean modeling tasks, and efficiency arguments. Subsequent chapters detail specific robot configurations and the lessons learned and models acquired from their deployment into subterranean voids. Combined, these chapters investigate the hypothesis that an array of robot configurations, distinguished by entry constraints and robot mobility, is required to achieve subterranean modeling objectives.

4.1 Spatial Arguments

Consider a simplified subterranean void representation, such as the one introduced in chapter 1, and an associated overlay of borehole drilling locations based on surface features, as

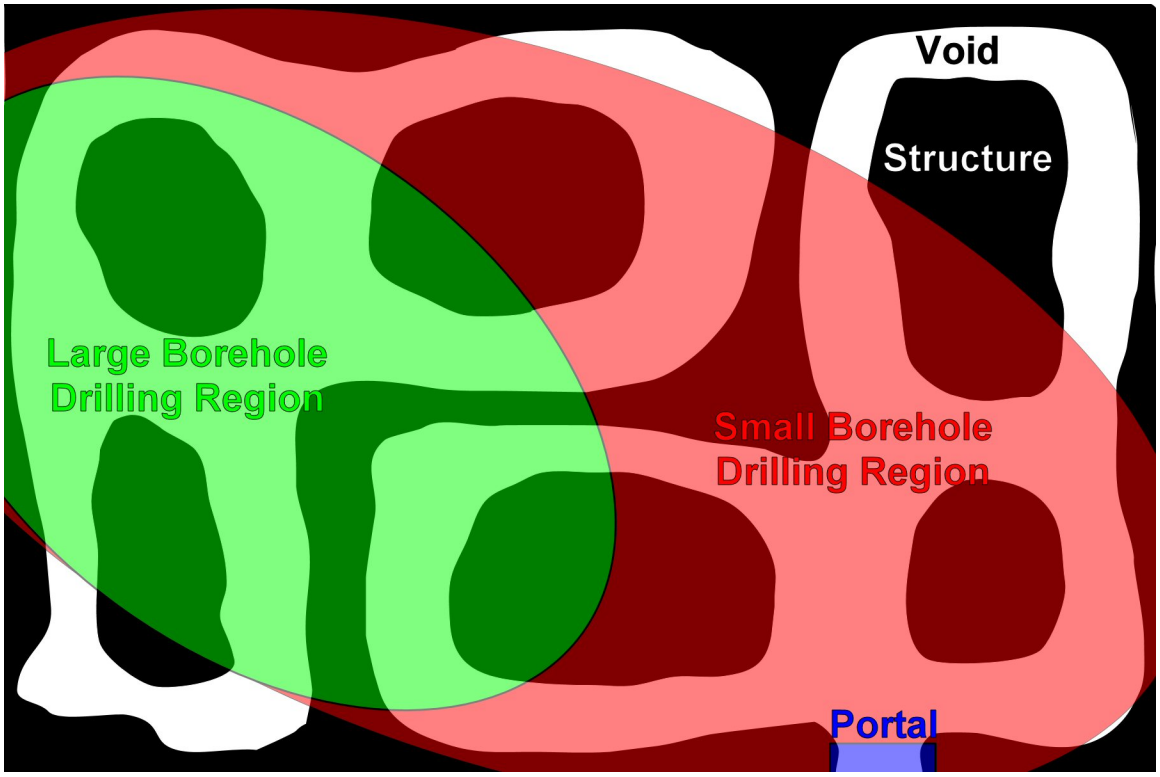

Figure 4.2: Overlay of acceptable borehole drilling locations on a representative subterranean space.

shown in figures 4.2. The topology of the space includes two unconnected regions, with portal access into only one of the regions. With a modeling goal of complete coverage, this implies that a device which requires portals to enter voids will not be capable of fully explore the space, and is limited to the coverage shown in figure 4.3. Therefore, modeling devices which are borehole deployable are required.

The simplest borehole devices are static sensors that provide range scans after entering voids, but do not have any mobility capabilities. A possible series of boreholes and models acquired from this type of device as applied to the situation outlined above results in the coverage pattern seen in figure 4.4. While this method provides some coverage of the region without a portal, even in conjunction with a long range mobile robot deployed via the portal, full modeling coverage of the subterranean volume is not possible. Therefore, some configuration which can provide both borehole deployment and some mobility is required.

A possible coverage pattern, utilizing the same borehole locations as in the prior example, but with a limited-range, mobile, borehole deployed robot is shown in figure 4.5. While some previously unobservable portions of the void are now within the robots field of view, the limited range of this robot configuration results in less than full coverage. There-

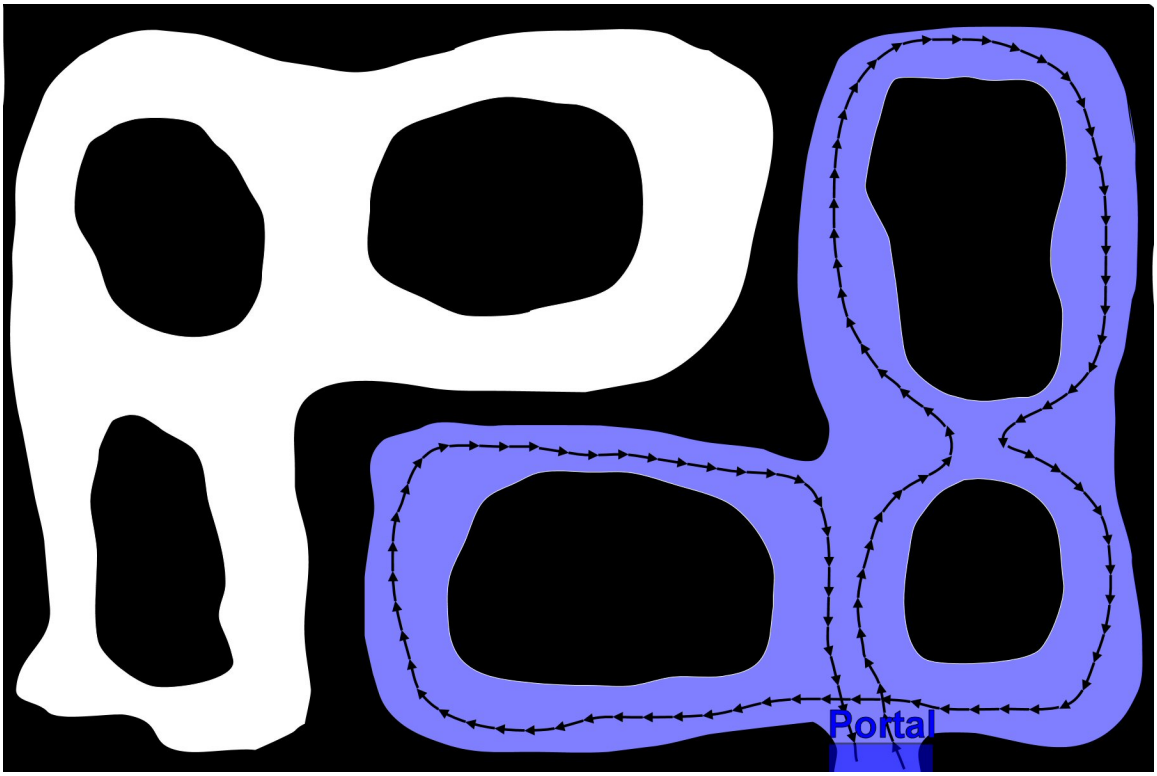

Figure 4.3: Coverage pattern achievable with an unconstrained entry, unlimited mobility robot configuration.

fore a borehole deployed mobile robot is not an "ultimate" configuration that can solve all problems. This *spatial* example therefore motivates the use of both portal and borehole deployed systems. It remains to be shown why static sensors should be employed if mobile borehole systems are available.

Small zero mobility sensors have two distinct advantages over mobile borehole deployed systems. First, they can be engineered to deploy down smaller diameter boreholes. Smaller diameter holes can be drilled by smaller rigs in locations in rough terrain not reachable by larger drilling rigs. Secondly, they take much less vertical void height to deploy into. One of the unanticipated, but significant configuration considerations revealed through this research is the importance of vertical clearance when entering and operation in reduced height voids. Therefore borehole sensors can deploy into low seam mines or underground bunkers that preclude the use of mobile systems due to requirements for deployment volumes and traversable void floor conditions. Since borehole-deployed static sensors never actually touch the void floor, conditions within the void do not effect modeling or operations. For this theoretical example, small diameter boreholes are feasible over a broad area and enable vantage points for borehole sensors that can not be achieved even when taking full advantage

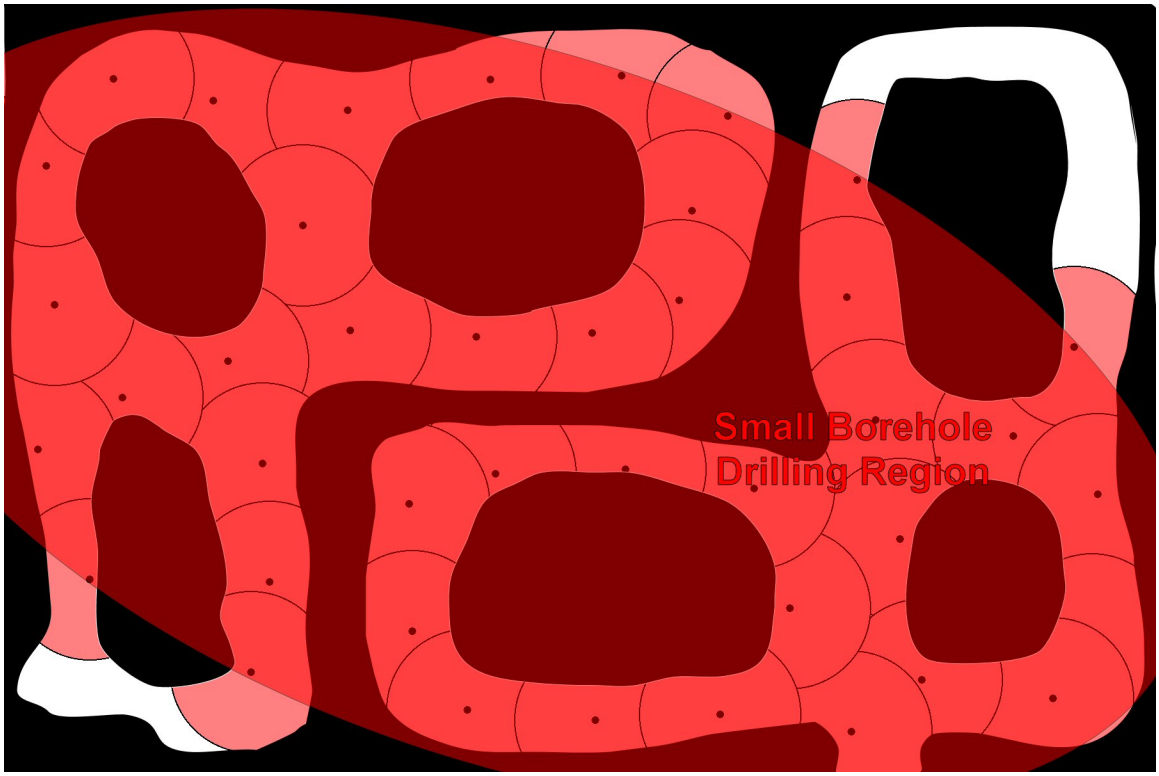

Figure 4.4: Coverage pattern utilizing borehole deployed sensors.

of borehole explorer mobility. Therefore, this subterranean modeling scenario provides *spatial* arguments which necessitate all three of the described robot configuration classes.

Do other robot mapping configurations exist? Means of void entry and robot mobility are considered to inquire about this. Portals and boreholes are the only options for entering voids. While large diameter boreholes or constricted diameter portals perhaps blur the separating line between the two entry types, no other options exist for entering voids. Similarly, a device can only be either static or mobile. No third option exists. While other meaningful divisions of the design space of subterranean modeling robots might be possible, discretization based on mobility and entry constraint results in three clearly distinct configurations. All three types of configurations are necessary for the span of subterranean modeling tasks, such as those real-world applications outlined in the next section.

4.2 Operational Arguments

Practical arguments stem from application scenarios which could benefit from robotic subterranean modeling. While not as simplified or all encompassing as the theoretical scenario described above, application scenarios distinguish and motivate the different types of robot

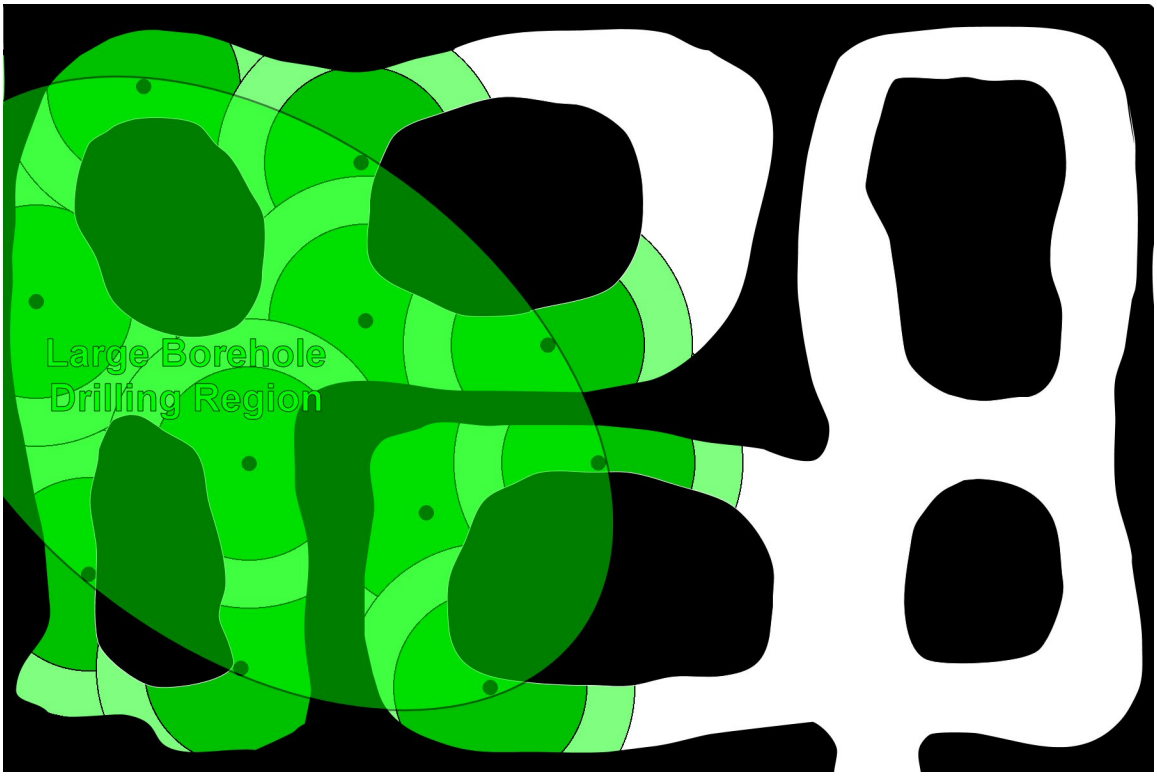

Figure 4.5: Coverage pattern with a constrained entry, limited range robot. Different traverses are shown in different colors/patterns.

configurations discussed in this thesis.

4.2.1 Impoundments Require Borehole Explorers

Modeling sealed subterranean voids underneath bodies of water motivates mobile, borehole deployed robot configurations. Sealed voids preclude using portal deployed systems, while borehole locations are forced to waters edge, away from where modeling is required. Modeling to prevent flooding due to mine impoundments placed over abandoned mines is a specific example of such a scenario. Modeling abandoned mines under impoundments has utility both in impoundment construction or expansion and in detecting structural degradation under existing impoundments. Impoundment permiting and construction requires a minimum barrier thickness between impountment and underground mine. Determining and ensuring this thickness requires models of surrounding subterranean voids. Water seepage can lead to impoundment failures, as shown in figure 4.7. Detecting erosion and monitoring barrier thickness also requires abandoned mine models.

Figure 4.8 comes from a report on the failure of a mine impoundment which breached into an abandoned underground mine. Underestimation of wall thickness at the point of

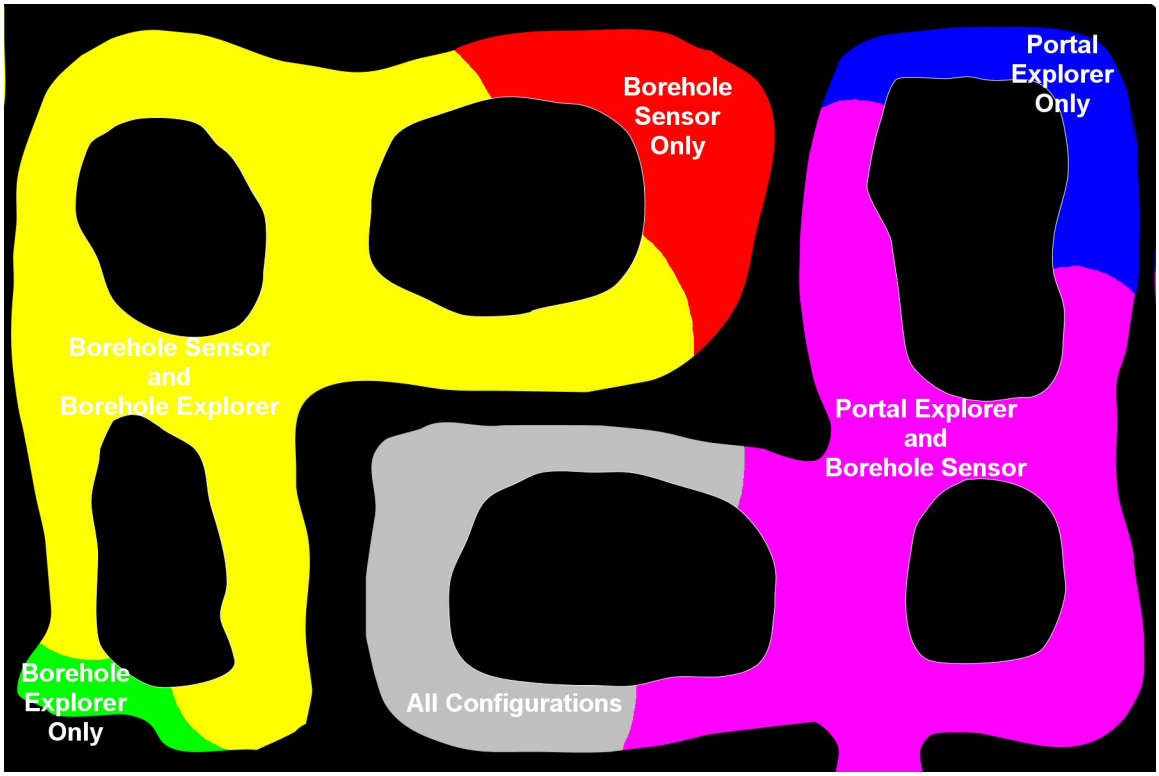

Figure 4.6: Composite coverage using three configuration classes.

breach between the abandoned mine sections and impoundment due to insufficient modeling was partially responsible for the breach. Figure 4.9 shows locations of boreholes drilled to better understand why the breach occurred following the failure.

The Inez, Kentucky mine impoundment failure provides a real-world motivation and justification for the existence and development of borehole explorers. Only robots in the borehole explore configuration class could have detected erosion of the barrier. A borehole deployed mobile robot could have been inserted through hole DH1-11 (figure 4.9) prior to the breach, driven to the end of Entry 1 and identified structural degradation or determined wall thickness. The breach location was in an abandoned section of an active mine, with at least two sets of seals between it and the nearest portal (figure 4.8). Therefore portal explorers would not have been suitable for inspecting the barrier wall. The point of failure was also close to one hundred feet away from the shore line of the impoundment, thus sensors deployed down safe boreholes would be unable to resolve details of the eroding barrier.

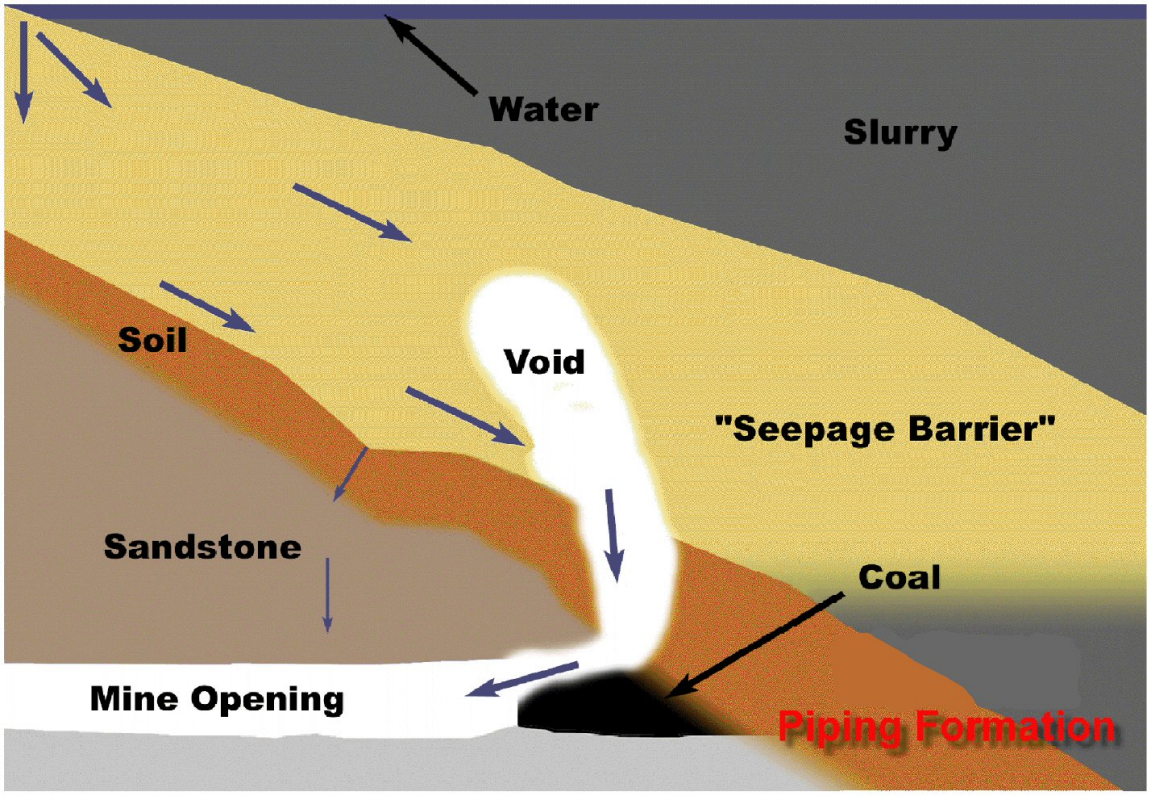

Figure 4.7: Erosion prior to mine impoundment failure.

4.2.2 Untraversable Floors Require Borehole Sensors

Modeling voids with untraversable floors motivates using borehole sensors. Voids such as collapsed underground bunkers can be so filled with rubble as to be untraversable. Random distributions of rubble result in pockets, pinch points, and other geometries which limit where mobile robots can reach. Vertical caverns or well-like voids often possess features of interest far above what is in range of mobile robots constrained to operations on void floors. In both these cases, mobile robots are constrained by floor conditions, thus limiting the span of modeling objectives achievable by mobile systems. Therefore, borehole sensors are necessary in order to accomplish these modeling tasks.

Understanding how explosions impact underground voids is of interest in the defense community and requires modeling of tunnels before and after blasts in order to calculate forces, velocities, and energies involved. While pre-blast modeling can often be easily performed using a variety of approaches, post-blast modeling requires robot configurations compatible with chaotic rubble fields, reduced void heights, and dangerously unstable conditions. Rubble fields present a dual challenge to mobile robots. First, mobility in rubble ranges from difficult to impossible. If a mobile robot can't traverse a rubble-filed void, it

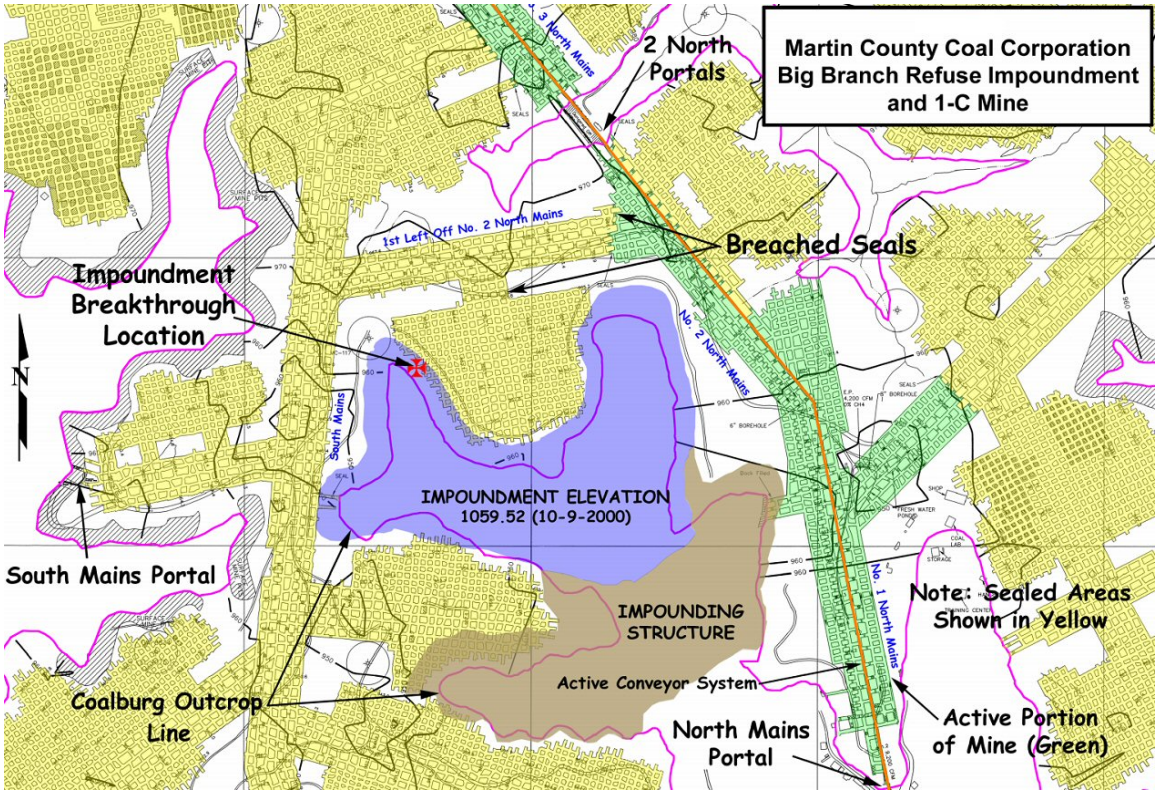

Figure 4.8: Annotated map of the Inez, Kentucky impoundment breach.

can not model it. Second, driving over unstable rubble changes void geometry, resulting in corrupted or incomplete models. Modeling conducted to determine forces and energies involved in an explosion losses value when the modeling system introduces new energies into the environment.

Only borehole sensors are capable of modeling untraversable rubble in thin voids without disturbing the void environment. Because borehole sensors remain suspended from tethers during the entire modeling process no environmental disruption occurs. Additionally, borehole sensors gain new vantage points by deploying down different holes rather than traversing void interiors. Therefore, traversability no longer impacts modeling. Finally, borehole sensors are simpler and smaller than mobile systems, allowing deployment into confined spaces not accessible to mobile systems. The combination of modeling from void ceilings, moving via external forces and deploying into tiny void volumes makes borehole sensors the only viable option for modeling voids with floors rendered untraverable due to rubble.

Borehole sensors are the only option for modeling vertical voids. Consider a 3,000 foot deep, 20 foot diameter silo-like vertical mine shaft[1]. Modeling the mine shaft walls to

[1]North Ore Chute, Kennecott Utah Copper Corporation, Salt Lake City, Utah

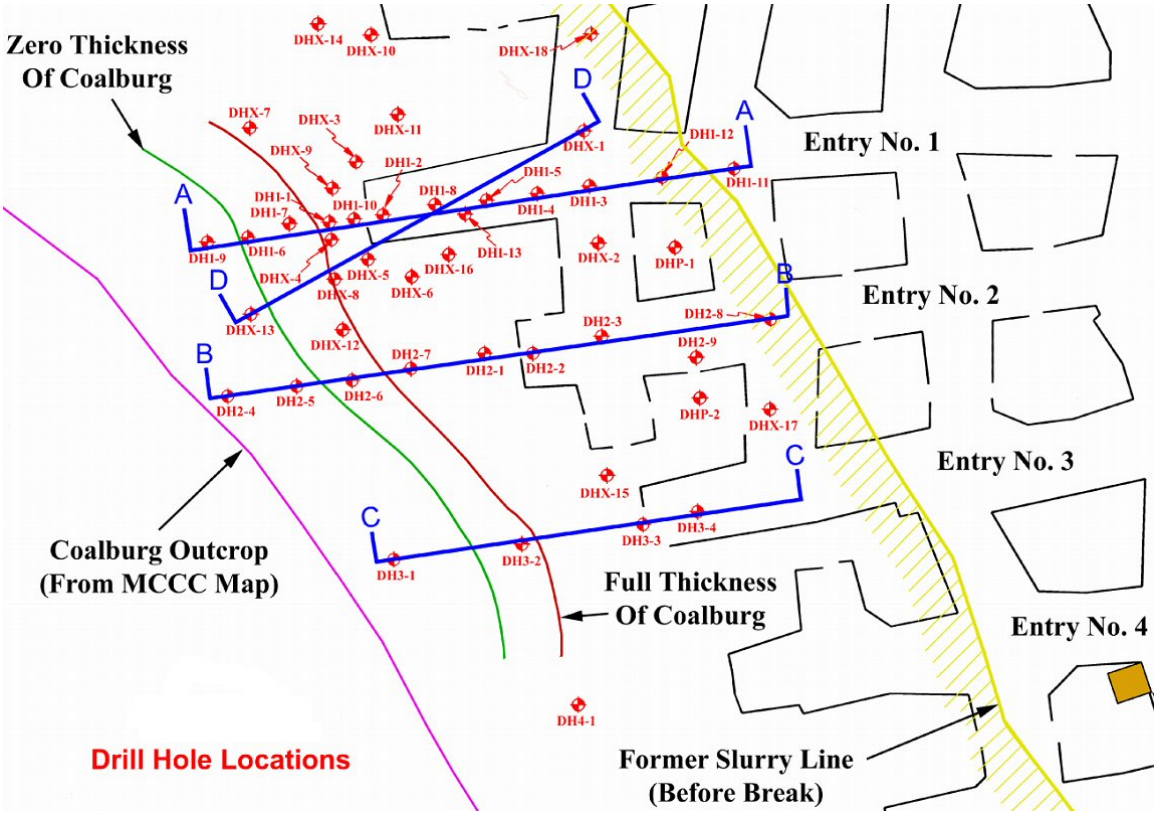

Figure 4.9: Post-breach exploration drilling. Note only borehole DH1-11 would have been possible prior to the breach.

determine structural conditions and detect concrete spalling is required in order to plan for rehabilitation and reopening of the mine shaft. Mobile or static devices could be deployed through 18 inch diameter holes, as shown in figure 4.10. Robots that drive on the floor provide no advantage in modeling such a vertical structure, thus vertical voids motivate subterranean modeling robot configurations of the borehole sensor class.

4.2.3 Mine Rescues Require Portal Explorers

Subterranean modeling tasks which require rapid exploration of paths between two points in portal accessible voids motivate portal explorers. Mine disasters require rescue workers to rapidly progress from mine portals to a working face (often the point in the mine furthest from the portal) following established escapeways. Long, linear, path based modeling tasks naturally lend themselves to being modeled with mobile robots. In the absence of time constraints, both portal and borehole explorers could model the entire escapeway. However, the overhead involved in drilling boreholes and deploying mobile borehole configurations, combined with the decreased mobility of borehole explorers, eliminates this class

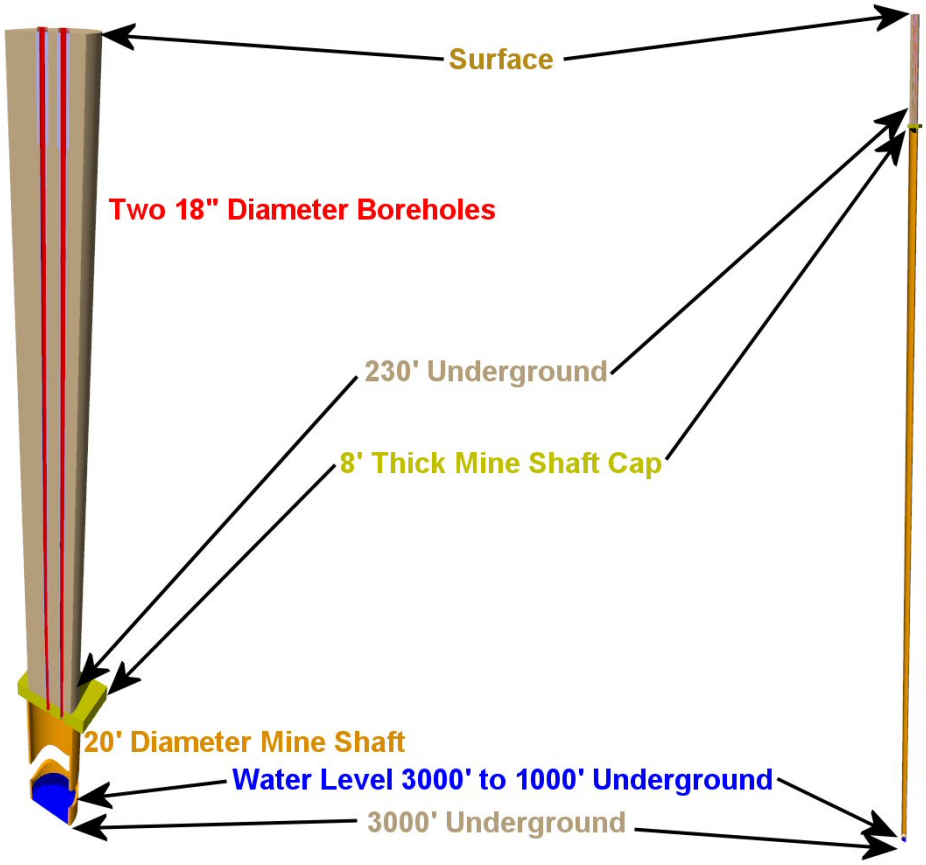

Figure 4.10: Abandoned mine shaft cross section detail (left) and at full scale (right).

of configuration from consideration in mine rescue responses. Only portal explorers possess the combination of mobility, range, operational speed and deployment speed necessary to successfully obtain models in mine rescue scenarios.

One specific example of this general mine rescue scenario occurred at the Sago mine. An explosion underground trapped 12 miners at the working face of the coal mine, two miles away from the portal. The miners were trained both to exit the mine following a specific route in the case of emergency and to go to a specific location, the working face of the mine, if they found themselves trapped by fire or debris. Therefore, the area of interest to be modeled (the primary escapeway) and the ultimate destination (the working face) were both clearly defined modeling objectives. Unfortunately no subterranean modeling robot was deployed as part of the rescue attempt at the Sago mine. In the future such systems will play an integral role in determining the environmental and structural conditions of escapeways and cutting a 30 hour rescue using traditional human rescue team techniques down to a 3 hour rescue employing robot subterranean modelers as advanced scouts.

Mine impoundments, vertical voids and mine rescue scenarios provide specific, oper-

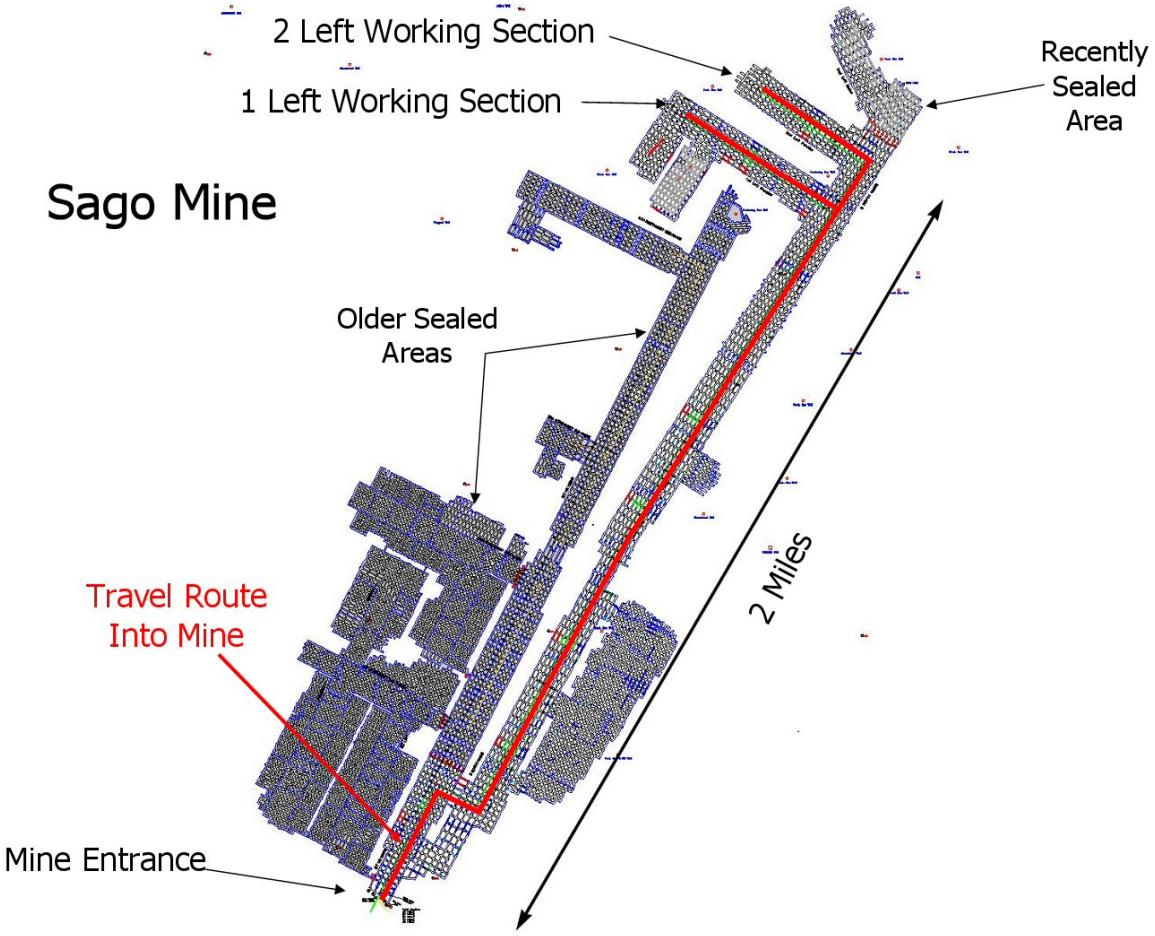

Figure 4.11: Pathway followed by rescue team following the Sago mine accident.

ational reasons for the development of borehole explorers, borehole sensors, and portal explorers respectively. These three scenarios represent a span of real-world subterranean modeling tasks that can not be effectively solved with robots from only two configuration classes. Therefore, development and demonstration of specific robots in three key robot configurations classes is warranted.

4.3 Efficiency Argument

For a given subterranean modeling objective and set of robot configurations, one configuration can accomplish the task with less operational cost than other configurations. Cost to enter subterranean voids increases with hole diameter and depth. Long range portal explorers require significant energy storage, leading to large, heavy platform which requires big, expensive holes. Small borehole sensors only require small, cheap holes to enter a void, but many holes are necessary if extensive modeling is required. Therefore, small modeling

tasks, local in extent, can be most efficiently accomplished with borehole sensors, while extensive models might be more cheaply acquired using portal explorers. Showing that each of the three configuration classes has superior efficiency for certain modeling tasks further motivates developing configurations in each class.

The early configurations distinguished, conceived and evaluated in this research are not yet refined or proven in the marketplace. They will be optimized, like all conceptions, through the test of time. Demonstrating configuration concepts and classifying configurations with respect to void access and mobility are purviews of this thesis. Developing optimality and demonstrating efficiency are not goals of this thesis. Arguments for efficiency are not comprehensive as detailed data about performance and cost related to specific configurations are incomplete.

Mobility combined with sensor range allows greater modeling coverage from a single entry than sensor range alone. For example, consider two robot system, one with a 100m sensor range, but no mobility, and one with a 50m sensor range and 50m traverse range. In a long, non-branching tunnel both systems will have equivalent modeling coverage, although the data density of the mobile system will be superior due to improved range sensor incidence angles. The mobile system will require two sorties in opposite directions within the tunnel to provide modeling coverage comparable to the longer range static sensor. However, when considering a more mine-like structure, with a grid of corridors on 50m centers accessed in the middle of one corridor, the mobile systems advantages become more apparent. Consider six sorties starting from the same entry point. One each sortie a robot traverses 25m to one of the two adjacent intersections and then traverse 25m down one of the three available corridors, as shown in figure 4.12. Combined with the 50m sensor range, this results in 500m of linear coverage, versus only 200m of coverage possible with the 100m range static sensor. This improvement in coverage due to mobility is encoded as a branching factor in the equations for analyzing efficiency.

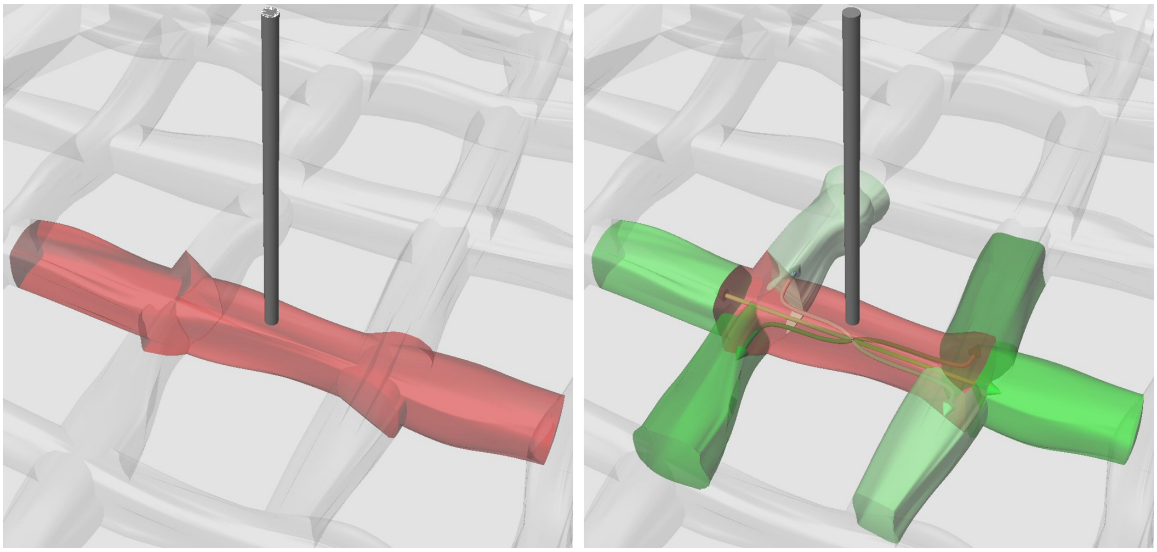

Figure 4.12: How a mobile robot branching factor improves modeling coverage from a single entry: Zero robot range, 2X sensor range (left); X robot range, X sensor range (right).

$$
\begin{aligned}
R_r &= \text{Robot Range}(m) && (4.1)\\
R_s &= \text{Sensor Range}(m) && (4.2)\\
D_e &= \text{Entry Diameter}(m) && (4.3)\\
L_e &= \text{Entry Length}(m) && (4.4)\\
\alpha &= \text{Scaling factor} \left(\frac{\$}{m^3} \right) && (4.5)\\
C_e &= \text{Entry Cost}(\$) && (4.6)\\
C_f &= \text{Fixed Cost}(\$) && (4.7)\\
E_m &= \text{Model Extent}(\$) && (4.8)\\
B_m &= \text{Model Branching Factor (dimensionless)} && (4.9)\\
C_m &= \text{Model Cost}(\$) && (4.10)\\
C_e &\propto L_e D_e^2 && (4.11)\\
C_m &\propto C_e \lfloor \frac{E_m}{B_m R_r + R_s} \rfloor && (4.12)\\
C_m &\propto L_e D_e^2 \lfloor \frac{E_m}{B_m R_r + R_s} \rfloor && (4.13)\\
C_m &= C_f + \alpha L_e D_e^2 \lfloor \frac{E_m}{B_m R_r + R_s} \rfloor && (4.14)
\end{aligned}
$$

Graphing cost to model versus model extent provides a meaningful view of configuration

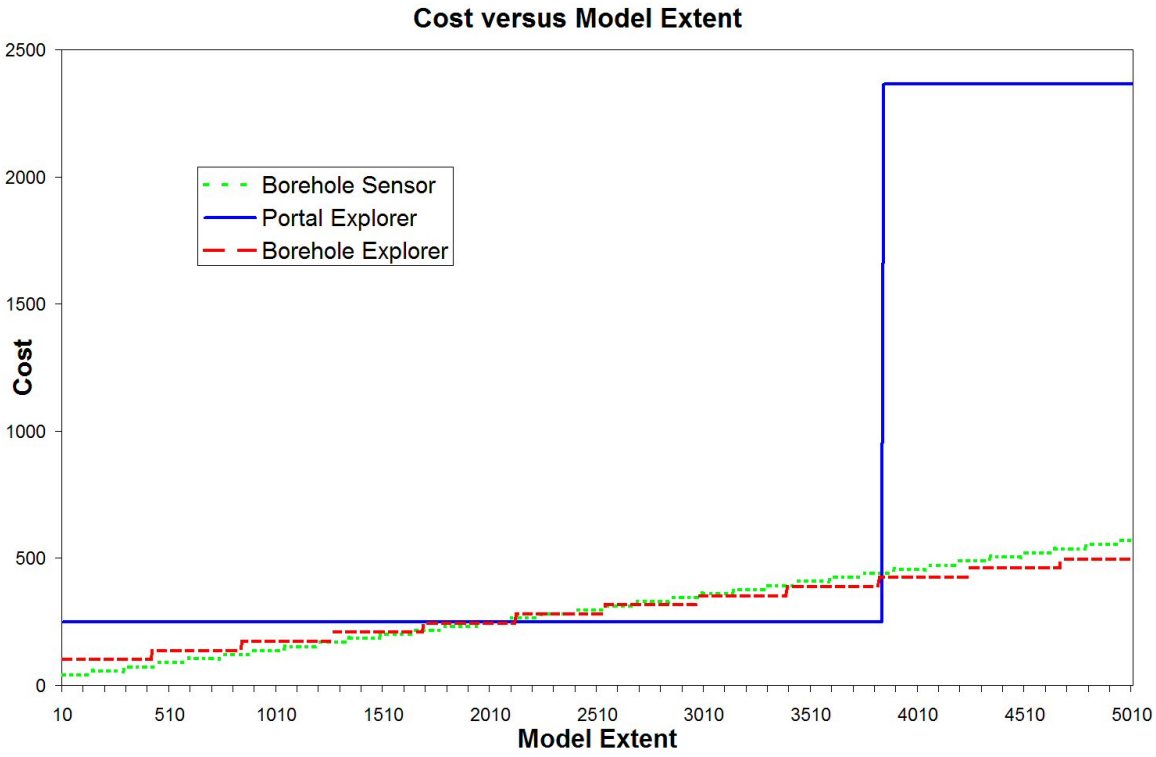

Figure 4.13: Graph showing configuration class cost versus modeling task extent.

class efficiency (figure 4.13). Model cost is composed of fixed, one-time cost and entry cost. One-time costs include such quantities as robot cost, mobilization fees, and insurance. Entry cost is proportional to volume of material removed in constructing an entry. Hence, deeper, larger diameter borehole entries are more expensive. Entry cost also depends on model extent versus robot composite range. If a modeling task requires more than one entry to complete, the entry cost will increase in increments. For the purposes of this analysis, it is assumed that a portal entry already exists and can be utilized in modeling.

The graph of cost versus model extent shows three distinct regions. Modeling tasks small in extent will be most efficiently performed with borehole sensors, as a single hole can ofter provide sufficient access to provide all required data. Larger modeling tasks are most effectively achieved with portal explorers covering large swaths of area. The limit to this region is the range of a portal explorer. Once an additional entry is required, the cost associated with providing an entry large enough to allow portal explorer access becomes prohibitive. For extensive modeling of constrained access spaces borehole explorers are the most useful configuration class. Each borehole drilled can support multiple sorties and extensive mapping is borehole explorers are employed. These three regions on the efficiency graph shift and merge based on one-time and entry costs, but provide yet another motivation

for the development of three configuration classes to attempt the span of subterranean modeling objectives.

Chapter 5

Existing Implementations

The primary contribution of this thesis lies in the exploration of the configuration space of subterranean modeling robots. While many configurations were conceived, implemented and tested as part of this work, some existing configurations, tailored to alternate environments, have been successfully deployed by others into subterranean environments to model voids. This section distinguishes these existing configurations and explains how they fit into the configuration space regions detailed in the introduction.

5.1 Unconstrained Entry, Zero Mobility

Large, human-operated sensors can be utilized in unconstrained entries, such as mine portals and other human accessible passageways. These sensors can be fragile given human oversight for their care, and they can be expensive because the care precludes common damage and loss. Featuring unmatched resolution and range, these sensors are commonly used to generate as-built three dimensional models in industrial, archealogical and museum settings. This class of modeling device is fully evolved and commercially available.

The unconstrained nature of void entries allows this class of device to utilize the largest components, thus enabling superior sensors, stiffer and more precise mechanisms, and rugged, heavy components. The superior nature of the sensors is demonstrated by the range and accuracy of these devices. The precision of the mechanism is reflected by the angular resolutions of these scanners shown in the same table.

Zero mobility implies that the sensor can only achieve vantage points for modeling that are accessible and safe for human presence. This limitation means that this class of robot configuration is not suitable for modeling the interior of abandoned mines, underground structures with unstable structure, or sewers and tanks which might contain noxious gases. However, modeling of the portal area of such inaccessible voids (the structure visible from the portal) is still possible. While systems for deploying these sensors into subterranean

voids can be conceived, deployment mechanisms like booms and pantographs are limited to tens of feet and hence don't impact arguments based on easy human access to the scanners.

Human access to scanners additionally implies that power and communications can be tethered to outside sources, so a scanner can consume unlimited electrical and processing power. Furthermore, an unreliable sensor can still be used to obtain models, as device failure, which can be easily corrected with human intervention, no longer carries consequences beyond increased data acquisition times. Without requirements on size, weight, power, or robustness, these devices feature the most advanced scanners and sensors of all the possible configurations for subterranean modeling.

Commercially available 3D modeling sensors feature angular resolution which is superior to all other scanners discussed in this thesis, but at prices larger than the cost of some complete subterranean modeling robots.

Unconstrained entry, zero mobility devices excel in modeling voids that are observable from a single vantage point. Modeling from multiple vantage points is also possible. Composite models are often created by stitching together range data from multiple vantage points as demonstrated for tasks such as modeling famous sculptures, creating as-built models of complex chemical plants, and monitoring open-pit mine high walls. Devices from companies such as a Leica have been used to model a portion of an underground tunnel that was also modeled using a custom built mobile subterranean modeling robot. The two different modeling methods produced comparable models and hence verified both the modeling procedures of both configurations.

5.2 Unconstrained Entry, Limited Mobility

Modeling of enclosed volumes has been demonstrated for non-hostile indoor environments numerous times. Robot configurations that have demonstrated successful indoor modeling can be applied to the task of subterranean modeling. Additionally, existing outdoor robots have demonstrated their ability to handle terrain comparable to that encountered in the harshest subterranean spaces. However, the generally more difficult conditions within voids means that the mobility of indoor systems is often insufficient, while the size and environmental compatibility of many outdoor systems limits what regions of a subterranean void are accessible. Hence these existing systems are classified as having limited mobility.

The class of robot configurations with limited mobility deployed through unconstrained entries encompasses three different types of configurations. First, constrained entry, limited mobility configurations (discussed in depth in chapter 8) deployed through unconstrained entries belong to the class of unconstrained entry, limited mobility configurations. Second, configurations designed for subterranean modeling, but with limited range due to design trade-offs made for ancillary considerations belong to the class of unconstrained entry,

limited mobility configurations but are superseded by configuration discussed in chapter 6 and therefore not further discussed. Third, configurations originally tailored toward solving problems other than subterranean modeling that have subsequently been retasked are unconstrained entry, limited mobility configurations. The remainder of this section deals with the third option and examines existing robot configurations and their applicability and adaptability to subterranean modeling.

Limited mobility robot configurations that operate through portals and other unconstrained entries have some lack of mobility, lack of range, or incompatibility with the subterranean environment. Insurmountable obstacles, steep slopes, and impassable standing water constrain mobility. Energetic constraints limit the maximum robot range. The operational environment of many subterranean voids limits the application of technology common in surface robots. The lack of GPS requires sensing configurations capable of solving the SLAM problem. Ventilation and explosion concerns limit the use of combustion based power sources. While many of these environmental concerns are factors across all configurations, they only limit robot mobility in the cases where existing robot configurations are retrofitted or retasked to solving the subterranean modeling problem.

5.2.1 Limited Sensing

Many possible configurations fall into the unconstrained entry, limited mobility classification, but only configurations that have been used for subterranean modeling tasks are discussed here. Mobile robots from simple two wheel differential drive robots through general purpose six wheel mobility platforms have been retasked towards subterranean modeling. While none of these robots performed as well as those configured and designed around subterranean modeling, they were instrumental in developing and demonstrating the potential of subterranean modeling.

The first robot to perform subterranean modeling was Terregator in 1985 [Champeny-Bares et al., 1991]. Terregator, as seen in figure 5.1 was a general purpose mobility platform and was used over the course of many years to follow sidewalks, map mines, and test new sensors. Terregator provided the first robotically generated subterranean maps of portions of the NIOSH Safety Research Coal Mine using a sonar ring, and later a laser scanner. However, while Terregator employed cutting edge technology of its day, limited computing, lack of miniaturization and the energetics requiring a gasoline generator made the system unsuitable for long range subterranean modeling. Additionally, at the time the concept of providing high resolution survey-quality models of subterranean voids was not fully envisioned. This limitation of vision perhaps placed an artificial bound on the true effectiveness of Terregator in the realm of subterranean modeling.

Figure 5.1: The Terregator mobile robots climbing steps(left) and gathering data in a coal mine (right).

5.2.2 Limited Mobility

The first robot to perform subterranean modeling using SLAM was a two wheeled differentially driven robot based on the Pioneer platform, shown in figure 5.2. Utilizing scanning laser range finders and software designed to map the corridors inside a building, this robot acquired data within the NIOSH research mine and demonstrated the feasibility of relying on SLAM without any GPS or inertial information in order to produce accurate and detailed subterranean maps. However, the data obtained from this Pioneer was acquired in a special paved portion of a research mine. The Pioneer's small, underpowered wheels are only designed for indoor operations and are unable to provide locomotion in muddy or even moderately difficult terrain.

5.2.3 Limited Applicability

Topographer, a modified H1 HUMMER outfitted with laser range finders and high-end inertial sensing conducted perhaps the fastest subterranean modeling to date (figure 5.3). Driving through a highway tunnel at highway speed, Topographer captured slices of laser data that were tagged with IMU readings. The IMU was part of a position estimation system that included differential GPS, so prior to entering the tunnel the global position of the vehicle was accurately known. After exiting the tunnel GPS was reacquired and errors due to observed IMU drift were corrected. The quality of the inertial system and speed of the laser range finder enabled this modeling. Unfortunately, Topographer requires a human driver.

H1ghlander and SandStorm (which have autonomously subtended tunnels at lower

Figure 5.2: A Pioneer-based mobile robot acquiring laser data in a coal mine.

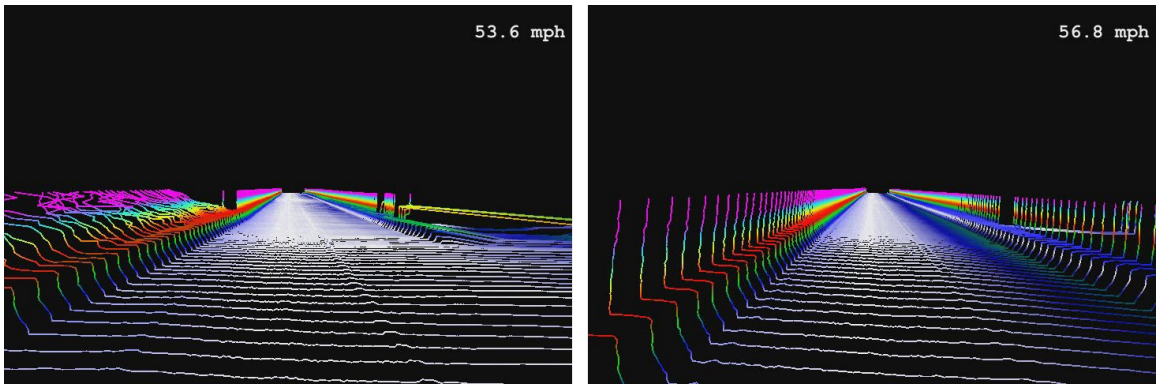

Figure 5.3: Models generated by Topographer: Approaching a highway tunnel (left), inside a tunnel with right hand access door (right).

speed) are fully automated robotic HUMMERs outfitted with nearly identical range and inertial sensors. While these vehicles can provide accurate subterranean models, they rely on expensive inertial measurement systems, utilize diesel engines, and are heavy and huge relative to many subterranean voids. Their size, more than any other factor, limits their utility within constrained access and constrained volume subterranean spaces. However, they are effective for mapping tunnels and underground spaces sized for traditional vehicles.

Chapter 6

Unconstrained Entry, Unlimited Mobility

Unconstrained entry means that the entry into an underground void is no more restrictive than the tightest pinch-point in the remainder of the void. However, this does not mean that the shape and size of a successful subterranean modeling robot is unconstrained. Subterranean environments are often cramped. For example, coal mines in the United States range in height from about two feet to about twelve feet with a median height of around five feet [Fotta et al.,]. Debris clutters the floor in both abandoned coal mines and breached bunkers. Objects project down from the ceiling, and walls can be partially collapsed. All of these limit the effective cross section through which a mobile robot must drive. For example, unconstrained entry is not applicable to a HUMMER sized outdoor robot because it is too large to enter and operate within most voids.

In the context of this thesis, unlimited mobility means that a robot is sized to fit and operate throughout its expected subterranean environment, equipped with locomotion sufficient to overcome expected obstacles, and range comparable to the extents of the underground environments in which it is expected to operate. For example, unlimited mobility is not a characteristic of small indoor robots as these configurations will be stopped by the first fallen roof bean encountered underground.

This section examines two mobile robots, Groundhog and CaveCrawler, specifically configured to produce subterranean models. These robots were designed to enter a void via portal and acquire range data while traversing difficult subterranean terrain.

6.1 Representative Implementation: Groundhog

Groundhog is a four wheel mobile subterranean mapping robot typifying unconstrained entry, unlimited mobility configurations and was built in direct response to the Quecreek

accident. After the accident, there was a desire to model the breach between the two mines as well as explore and model the abandoned Saxman mine, which was off limits to human inspectors. This simple high level objective, when combined with a few known conditions at the breach typifies unconstrained entry, unlimited mobility challenges and drove Groundhog's configuration design.

6.1.1 Groundhog Configuration Design

The Saxman breach (seen in figure 3.6) was created by the pressure of the water contained in the Saxman mine blowing out the barrier wall separating it from the Quecreek mine. As a result, the breach was ragged, smaller in cross section than the surrounding mine, and elevated off the floor of both mines. The smaller cross section limited Groundhog's height and width. The large step between the breach and the mine floors necessitated a large approach angle, large drive torques to climb the step, significant ground clearance to avoid high centering on the step, and a soft suspension to handle the step down.

The six foot wide by four foot tall opening limited Groundhog's size, particularly its height which was initial target at three feet. Additionally, the length of breach (roughly four feet) impacted Groundhog's drive configuration. Groundhog would either need tracks or a combination of wheel base and ground clearance sufficient to prevent high centering. The muddy, debris-strewn conditions on both sides of the breach required aggressive traction, thus favoring tracked or all wheel drive configurations.

Beyond determining the physical dimensions of the robot, the nature of the abandoned mine also influenced the design philosophy that went into constructing the robot. Under no conditions would humans be allowed to cross the breach into the Saxman mine. Therefore, the robot configuration had to be reliable. Reliability implies simplicity, so configurations such as walking robots or systems relying on untested custom built parts were not considered. Additionally, fewer actuators, fewer degrees of freedom, and simplified operations and controls were favored. The minimum number of actuated degrees of freedom necessary to navigate in a planar environment is two, hence Groundhog has only two degrees of locomotion freedom.

Coal mining regulations had a significant impact on the configuration of Groundhog. One of the major concerns in developing equipment for operations in coal mines is ventilation and the potential presence of explosive concentrations of methane. Diesel engines are sometimes used in mining, but not in areas with potential for methane accumulation. Only electrically powered machines are allowable in unventilated areas such as those down stream of the working face and the breach between the mines. Therefore Groundhog is electrically powered.

Current mining regulations recognize two primary means for safeguarding electro-mechanical

systems in order to prevent the possibility of explosion. Intrinsically safe devices don't possess sufficient energy to create a spark that could trigger an explosion. Therefore intrinsically safe is only useful for very low power sensors and electronics. Explosion proofing involves encasing potentially spark generating equipment within specially engineered enclosures that prevent an explosion inside the enclosure from propagating to the exterior environment. For Groundhog's primary actuators, explosion proof enclosures are the only option. Two approaches to explosion proof enclosures are possible. Using hydraulics to power the robot's locomotion allows all the high power electronic equipment to be placed in a single large central explosion proof enclosure. Pump motor, manifold, and switching electronics for the hydraulic valves can all be mounted inside the central enclosure, with only hydraulic power crossing the boundary of the enclosure. Using electric power requires explosion proof motors mounted externally to the central enclosure. While this shrinks the size of the enclosure, it also requires the appropriate connections and conduits be made to the electric motors in order to maintain the explosion proof rating. Ultimately it was decided to utilize hydraulics on Groundhog due to their inherent low speed high torque nature and the reduction in the number of connections and components that would have to be explosion proof.

Coal mining regulations dictate electric power, but do not specify whether that power should be carried on-board a device, or supplied externally via tether. Tethers offer one means to provide power to a robot as well as continuous communications. However, tethers must be effectively and actively managed in order to avoid either the robot or the environment destroying the tether. The two options for a tether and its associated deployment mechanism are to have a fixed tether reel at the point of deployment or mount the tether reel to the robot. Both options present problems. A base station tether reel requires the robot to maintain tension on the tether. While this is simple to accomplish over short, straight line distances, in the extensive and convoluted environment found in most mines it becomes difficult or impossible. The more corners the robot goes around, the higher the friction on the tether, until eventually the tether friction becomes greater than either the tractive force of the robot or the strength of the tether. Additionally, in order to prevent a robot from running over its own tether, a robot must coordinate its motions with its base station to ensure that the tether is respooled and managed. This coordination is difficult and prone to failure.

Robot mounted tethers largely eliminate the problem of cumulative frictional forces on the tether by deploying tether as needed based on robot motion. However, robot mounted tethers are limited in size and weight (and hence length) by available robot payload capacity and tip-over stability. As with the base station mounted tether, a system to prevent the robot from running over its own tether is also necessary. Disallowing certain robot maneuvers, such as multi-point turns, minimizes the risk of hitting the tether. Ultimately,

the configuration trade off is between the added complexity and motion restrictions a tether brings versus the ability to provide off-board power, processing, and control to a robot. In Groundhog's case, scenarios which both favored and disallowed tethers were envisioned from early on in the design process. Therefore Groundhog was designed to be self-sufficient with on-board power and processing, but with the tether supplying redundancy or additional functionality.

The nature of the desired modeling information also effected the configuration of Groundhog. Of most relevance was accurate modeling of the breach, which required high resolution, local 3D coverage. In addition to the breach, the true extent of the Saxman abandoned mine was also of interest. Despite the accident, mine operators planned to continue working the Quecreek mine after establishing a man made seal between it and the Saxman mine. However, the breach clearly indicated that the existing Saxman mine map was inaccurate. Continued coal mine on the Quecreek side would benefit from accurate knowledge of the extents of the Saxman mine. Therefore, Groundhog was designed to generate an accurate map of the perimeter of the Saxman mine near the breach using an edge based exploration pattern.

Interest also existed in modeling the portion of the Saxman mine between the breach and two borehole roughly six hundred feet from the breach drilled into the Saxman mine during the course of rescue efforts. This path based modeling task represented an intermediate task that require less refinement in Groundhog's hardware and software while still providing the potential to demonstrate subterranean modeling on a scale larger than the few feet of breach available. While much of Groundhog's original configuration was developed with this linear modeling task in mind, the more complex extent exploration task ultimately influenced Groundhog's configuration to a greater extent.

Effective exploration of the Saxman mine perimeter requires Groundhog to explore multiple dead end corridors. The comb-like mine structure favors configurations that can easily turn around or back up, systems with long range to fully explore the extents of the mine in a single mission, and agile platforms that can handle lots of maneuvering in tight spaces.

A robot can explore dead ends efficiently only if it can exit the dead end as easily as it enters. Retracing its path in reverse, turning in place, or perform a multi-point turn are all possible approaches to exiting a dead end, as shown in figure 6.1. Turning in place and multi-point turns incur energy penalties, as energy is wasted reorienting the robot. Additionally, locomotion configurations that can support point turns, such as tracks or differential steering, require significantly higher peak torque and power outputs in order to overcome the sliding friction inherent in such maneuvers. Therefore, energetically reversing is favored over turning.

From a sensing standpoint reversing is also favored over turning in place at a dead

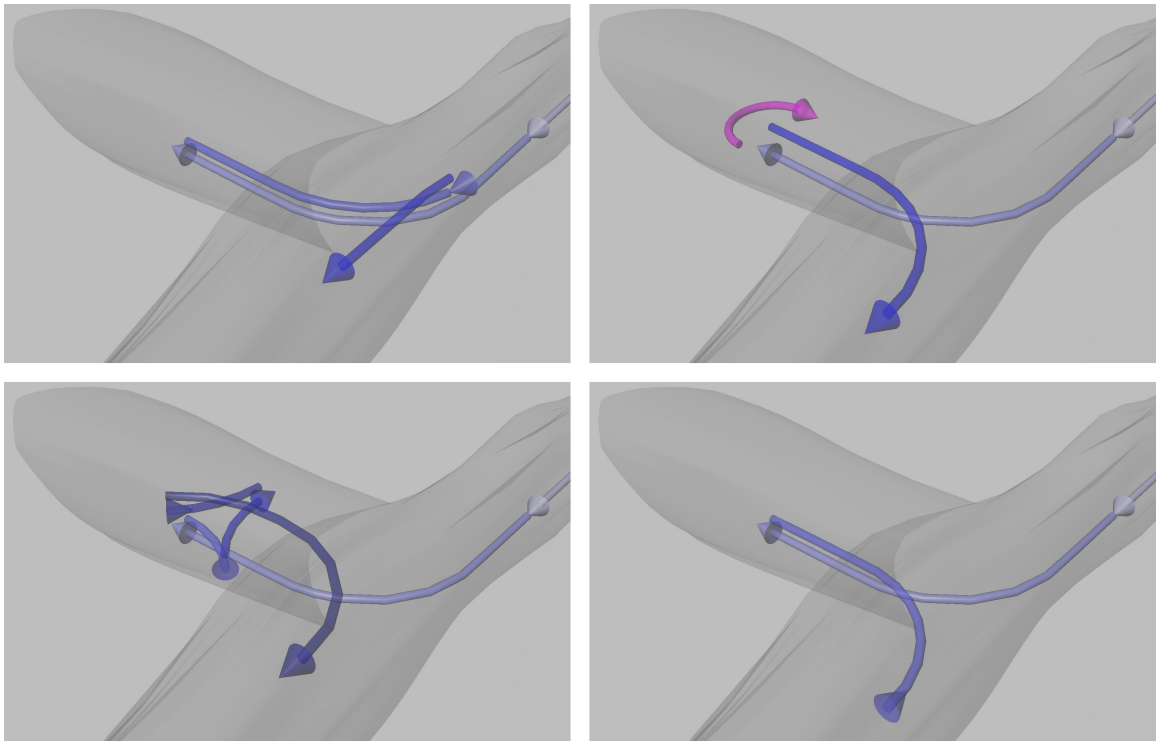

Figure 6.1: Ways for a robot to exit a dead end: Reverse course (top left), multi-point turn (bottom left), point turn (top right), reverse direction (bottom right).

end. Reversing can occur instantaneously and without changing sensor perspectives. This means that drift in inertial measurements are minimized because the time taken to turn is eliminated. Additionally, multi-point turns particularly, and to a lesser extent turning in place, involves significant robot motion, resulting in decreased localization accuracy as sensors and scanners are shaken about by the combination of rough terrain and high power maneuvers.

While reversing is preferred from an energetic and modeling standpoint, it does place additional requirements on a robot configuration. In an ideal world, reversing out of a dead end would simply be a matter of replaying locomotion inputs in reverse. However, this open loop approach is prone to failure in the muddy and debris strewn environment common to subterranean voids. Therefore, some close loop means of reversing is required. Two diverging approaches to providing the sensing and software needed to reverse are possible. The first approach uses the same sensors utilized in driving forward, but filtered through reversing software that models the robot and environment and that generates the correct motion commands based on forward looking sensors. The second approach is to use the same software for forward and reverse by configuring the robot mechanisms and sensors to be symmetric with respect to direction of motion. Reversing with only

forward looking sensors caries with it the implicit assumption that the environment being reversed over is unchanged. Additionally, having symmetric sensing implies either some level of redundancy in the sensing or continuous 360 sensing coverage around the robot. Because of these reasons, Groundhog was ultimately designed to be symmetric, although full sensing symmetry was not achieved prior to robot motions being placed under autonomous computer control.

Subterranean robot failures are to be expected given first generation robot prototypes operating in complex and dangerous subterranean environments. Robot failures underground are often unrecoverable, resulting in the total loss of the robot. Even severe failure modes, such as an autonomous vehicle rolling or an indoor robot falling down a flight of stairs do not result in the total loss of the robot and all its components. Robotic failures in the realms of space explorations, sub-sea operations, and nuclear cleanup provide analogs in terms of consequences of failure. But for subterranean robots to succeed they must be both reliable and economical, thus the space based approach of having redundant, expensive, and unique parts is generally not feasible. Therefore subterranean robots should be designed with loss in mind, utilizing cheap, easily available commercial and industrial components whenever possible.

Perhaps the most basic locomotion configuration choice, Groundhog is a wheeled rather than tracked vehicle. Both wheels and tracks are accepted mining industry technologies. However, wheeled systems in the size range dictated by the breach dimensions are much more commonly available than tracked systems. An explicitly steered wheeled vehicle also has lower peak power requirement compared to a differentially steered tracked vehicle. Finally, a tracked vehicle inherently has more components located lower down on a vehicle (suspension, track guides, etc.), thus exposing all of these element to gritty, abrasive, acidic mud present in mines. The decision to use wheels enable the extensive use of commercial components from all-terrain vehicles (ATVs) for Groundhog's primary locomotion mechanisms.

This section outlined the major elements of Groundhog's configuration, and the environmental, economic and physical constraints that led to these design choices. Additional insight into the relationship between the subterranean environment and robot configuration can be gained by examining the detailed mechanical, electrical and software system design of Groundhog. These topics are covered in the next three sections.

6.1.2 Groundhog Mechanical Design

Groundhog is a four wheeled vehicle. All four wheels are powered in order to provide traction for obstacle climbing even in deep mud. While traditional two wheeled Ackermann steering is sufficient for mobility in mines, Groundhog employs four wheel steering in addition to four

Figure 6.2: Groundhog's base frame with only original ATV components (left) and after additional structural enhancements (right).

wheel drive. Only steering two wheels results in a larger turning radius, which limits possible robot paths in constrained mine corridors, and also breaks the front/back symmetry of the robot, requiring custom software and procedures to handle reversing. Groundhog's driving and steering components come from two all-terrain vehicle front ends, severed from their rear halves and welded back to back. This results in a fully symmetric base frame (see figure 6.2). Additionally, it allows much of the drive train to utilize cheap, common, commercial components. In Groundhog the tires, wheels, hubs, steering knuckles, half shafts, A-arms, differentials, tie rods, and steering columns are all stock ATV components.

Most of the modifications to Groundhog's frame were made to reduce the overall system height. The suspension tie points were moved down and attached to box beams welded to the original frame. This lowered the vehicle profile and increased the initial compression of the coil-over shocks to compensate for the increased vehicle mass. Groundhog's suspension features four independent double A-arms with coil over shocks for roughly 6 inches of suspension travel. The steering column shaft, originally designed to connect to handlebars on the ATV, was chopped down and re-engineered to allow steering components to be placed directly onto Groundhog's base plate, as shown in figure 6.3. The two steering shafts were connected by a tie rod in a Z configuration, resulting in the front and back wheels always turning together. These modifications show how conventional configuration concepts are influenced by subterranean operations. The low ceiling height of many subterranean voids requires low robot height while maintaining mobility.

Groundhog's power source is a stack of deep cycle lead-acid batteries. Even though lead acid batteries are significantly inferior to modern lithium based batteries in both volumetric and mass energy density, they are cheap, robust, easily charged, familiar in mining and readily available, and hence consistent with Groundhog's design philosophy of utilizing cheap and readily available components. Groundhog uses a 24VDC system bus for compatibility with industrial motors, common laser line scanners, and many wide input range electronic devices. For a time, Groundhog's battery stack consisted of four series-connected

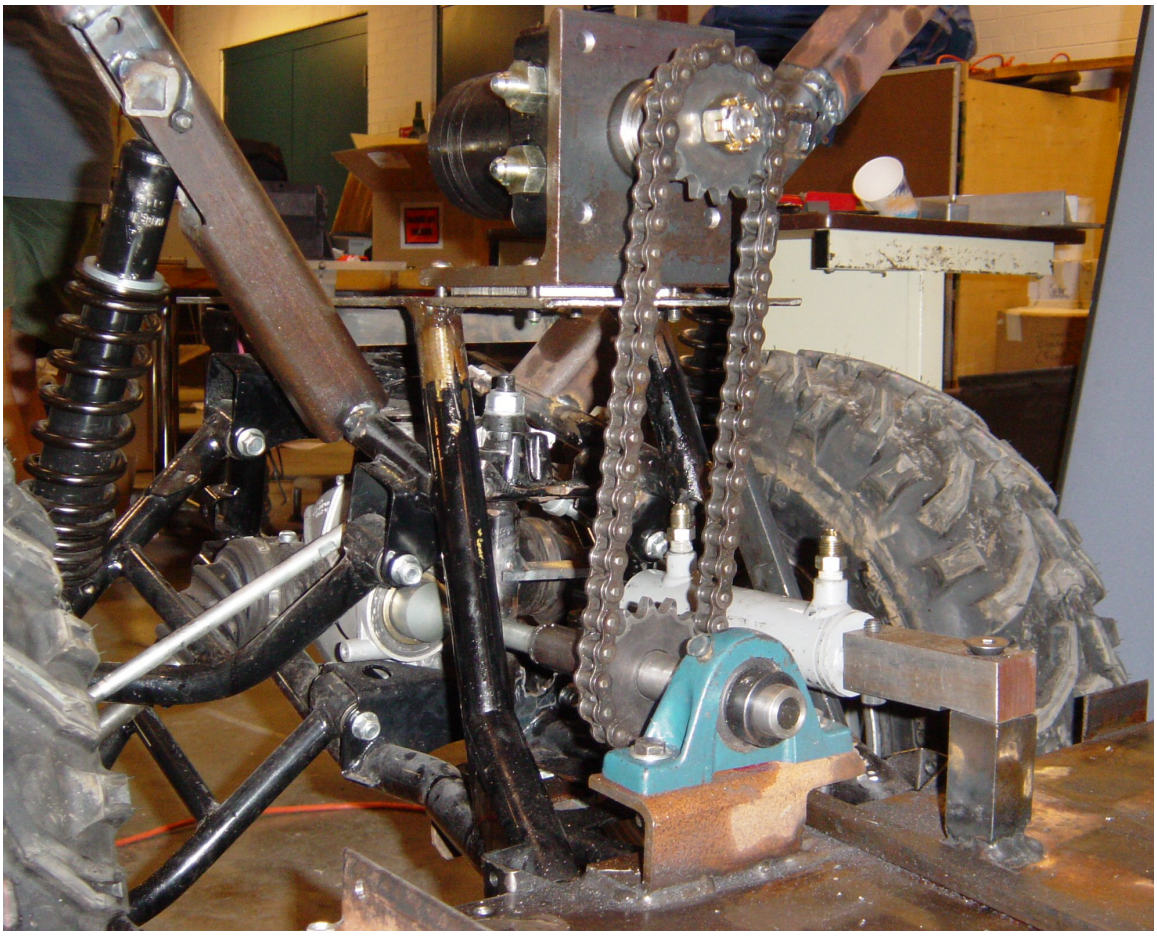

Figure 6.3: Details of Groundhog's steering, drive train and suspension.

6-volt, 150Ah batteries, providing a total of 3600 Watt-hours of energy. When more range (and hence battery capacity) was required, the system was switched to two parallel strings of three 8-volt, 120Ah batteries each, resulting in 5760 Whrs of energy.

Groundhog's hydraulic system consists of a pump motor, manifold, reservoir, two motors and one linear actuator, as shown in figure 6.4. The pump motor, housed inside the explosion proof enclosure, is a one horsepower 24VDC electric motor coupled to a hydraulic pump with two different flow rates selected by solenoid. The pump provides power to the hydraulic manifold, a commercial component originally intended for use in scissor lifts. The manifold is outfitted with three position 24VDC solenoid actuated valves. Two series connected hydraulic motors drive Groundhog. Each motor is connected via chain to the input to one of the stock ATV differentials. The series connection ensures that both pair of wheels will always be driven at the same angular rate, although torque sharing between the wheels is indeterminate. A double acting hydraulic piston, coupled to the steering tie rod

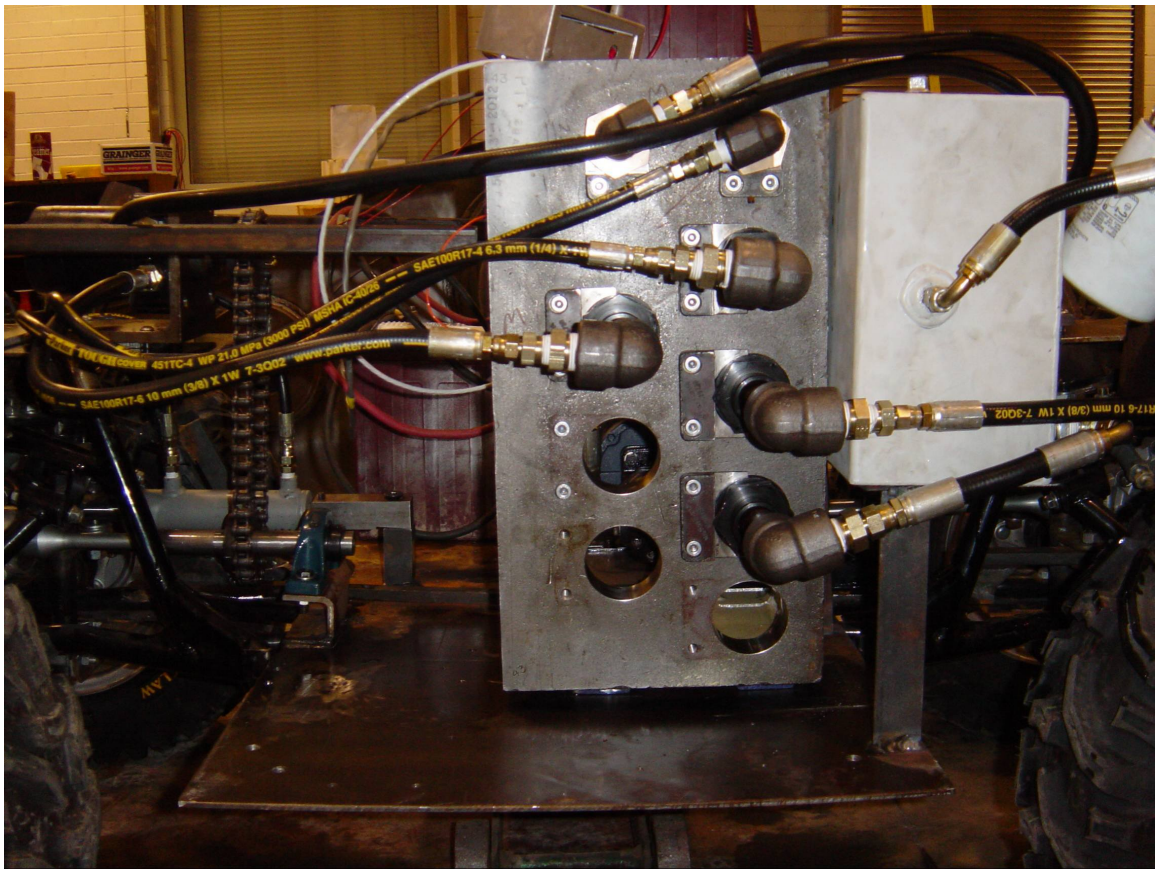

Figure 6.4: Details of Groundhog's hydraulic power system.

linking the front and back steering shafts provides steering angle control. The reservoir is mounted externally to the explosion proof for easy change out or addition of hydraulic fluid. A filter installed on the return lines prevents hydraulic fluid contamination. All hydraulic lines are ported through the explosion proof enclosure following common practice in mining equipment.

6.1.3 Groundhog Electrical and Sensing Design

Groundhog is operated from a single centralized computer which provides everything from low-level control of hydraulic valves through high-level reasoning about how to explore subterranean voids. At the core a 933Mhz PC-104 form factor motherboard provides processing, ethernet connectivity, a few serial ports, and keyboard, video and mouse connections for debugging purposes. Integrated on top of this motherboard in the PC-104 stack is a mixed signal card for providing digital and analog inputs and outputs, an expansions card for holding additional Flash based solid state memory for data logging, and a quad serial port card to provide additional control and communications to the modeling sensors. A

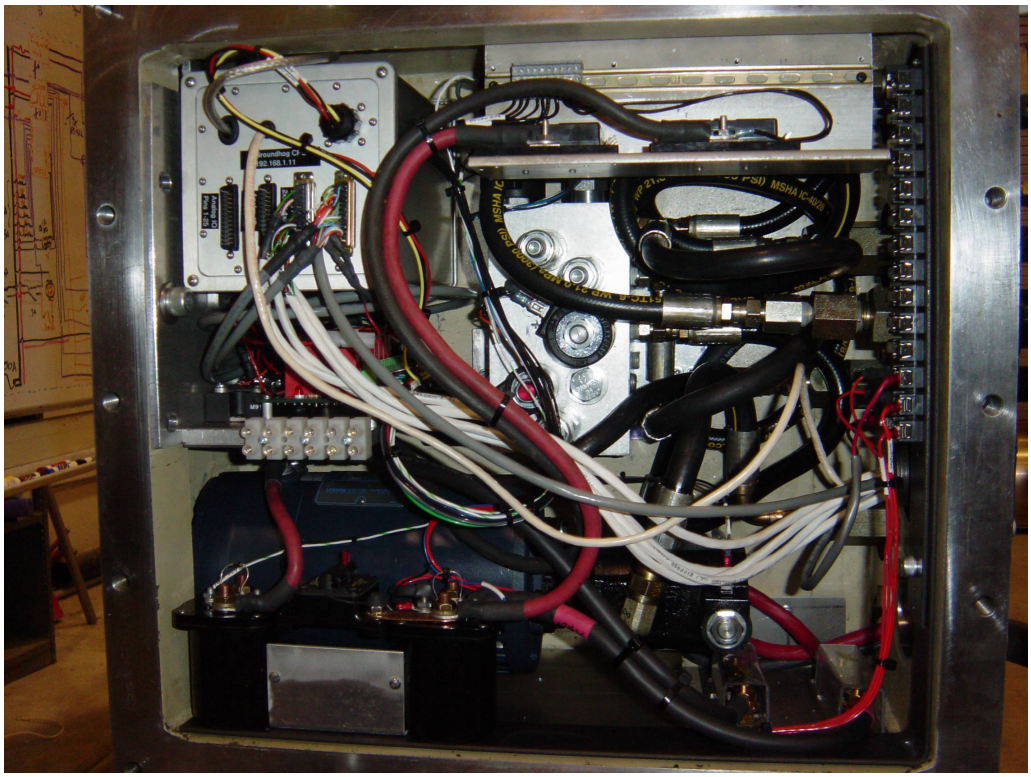

Figure 6.5: Groundhog's primary electrical components housed inside the explosion proof enclosure.

power conversion card connected to the primary power bus provides clean power to the computer and its peripherals.

Groundhog operates from a single 24VDC bus, fused at 100 Amps. The hydraulic pump motor draws directly off the primary bus. More delicate electronic components, such as the computer, sensors, and communications equipment, are powered from DC/DC converters to shield these sensitive devices from pump motor noise. Control signals flow from the computer, through the mixed signal card and to a bank of optically isolated solid state relays. These relays are wired to the solenoids in the hydraulic manifold mounted valves, as shown in figure 6.5.

Some inputs to the computer include redundant gas sensors, a potentiometer, a proximity sensor, a pressure sensor, and float switches. Gas sensors detect concentrations of oxygen, methane, and hydrogen sulfide. Concentation levels are reported to Groundhog via serial port. These measurements determine if the atmosphere is explosive, and can be used to either turn Groundhog back from entering more explosive areas or to shutdown Groundhog in case the environment is already too explosive. The potentiometer records steering angle and is calibrated as part of the Groundhog's initialization procedure. The

Figure 6.6: Two scanner configurations used for modeling: One of a symmetric pair of nodding line scanners (left), two non-symmetric fixed line scanners (right).

proximity sensor is used to count the teeth on the sprocket driving the differential input shaft, thus providing feedback as to Groundhog's speed. The pressure sensor is inline with the hydraulic motors and used to detect anomalies in the hydraulic system such as pump failure or overrunning of the motors due to Groundhog going down a steep slope. Eight float switches mounted at the corners of Groundhog's central base plate detect water level and are used to prevent Groundhog from going into dangerously deep water.

Groundhog's sensing incorporates industrial scanning laser range finders. These devices provide a minimum angular increment of one quarter degree and have a 180 degree field of view. The modeling sensor configuration on Groundhog evolved over time, as shown in figure 6.6. Originally, two fixed 2D sensors were used, one looking forward providing scan matching for localization, and one pointed upward, providing modeling data that was later integrated into models using the localization information from the forward looking sensor. While this configuration combined with human teleoperation control produced some of the models shown in section 6.1.5, it is insufficient for autonomous robot motion. In order for Groundhog to plan its own path it has to build terrain models. A three dimensional sensor, constructed by nodding a line scanner, gathers 3D information about the environment. The same data acquired for path planning is also used for model generation. A symmetric pair of these nodding scanners maintains robot symmetry while providing more complete modeling coverage.

Usage of scan matching to provide localization information was experimentally shown to be effective early in Groundhog's development. As a result, investment in expensive inertial sensing was limited, as full pose recovery from only inertial data was not required. A fiber optic gyro provides additional yaw rate information in order to improve scan matching and localization when Groundhog is cornering. A two axis tilt sensor provides continues estimates of gravity, thus correcting angular misalignments between 3D scans and compen-

sation for the yaw gyro to account for the non-level nature of the robot.

Groundhog's operations and void models are based on laser range data from the sensors outlined above. On-board cameras provide a means of viewing the environment and are useful for human monitoring, control, and understanding of the subterranean spaces Groundhog explores. A low light camera with an infrared LED ring provide video imagery, while a digital video recorder stores the video on board. A wireless transmitter operating in the 900Mhz range transmits the video signal back to the base station while the robot remains in line of sight range. Groundhog also provides an 802.11b wireless ethernet connection for line of sight operations.

The base station, while technically not part of Groundhog, is part of the operational setup used during deployment. The base station consists of an 802.11b wireless ethernet access point for data and robot control, a directional Yagi antenna for video reception, a camera aimed into the portal to externally monitor Groundhog's operations when in range, power conversion and fusing, and a battery. The base station allows operations to be conducted on monitored from a position away from the portal, which can be necessary for safety, comfort, or convenience reasons. A diagram showing a typical operation setup for Groundhog is shown in figure 6.7.

Beyond the electrical system outlined above, Groundhog also has some unique systems specifically designed to improve reliability. A hardware feature on the motherboard provides a watchdog timer. Once the watchdog is enabled, it starts counting down from its initial register value. If the register is not reloaded by a software process running on Groundhog's processor the watchdog will expire and generate a pulse output. This pulse is routed into a flip flop which is wired to temporarily interrupt the power supply to the processor, resulting in a hard reset of the computer. This provides a hardware failsafe to software induced errors by forcing a reset of the entire system.

6.1.4 Groundhog Software Design

Groundhog's autonomy utilizes the traditional sense-plan-act loop. Groundhog senses the environment in front of it through a 3D laser scan, acquired by the forward looking tilting laser scanner. This 3D scan is used for both modeling and planning. The 3D point cloud scan is converted into voxels, then condensed into a 2D cost map. Unlike outdoor robots, the cost map must take into account both obstacles on the floor and those projecting down from the ceiling. A minimum distance between a floor voxel and ceiling voxel detects pinch points, while gradients between adjacent voxels detects obstacles. Once the 2D cost map is obtained, an A* based planner plans a path to a point five meters ahead of the robot and centered in the corridor. The planner convolves a representation of Groundhog with the 2D cost map at each point in the plan to determine paths that minimize the amount of

Figure 6.7: Operational setup used by Groundhog during field experiments.

obstacles the wheels encounter. This allows the robot to either drive around or drive over obstacles and have been found to be very effective when railroad tracks are present in a mine. Once a path is determined, Groundhog follows it by utilizing laser based localization information to track the 2D path. Groundhog's local path planning processes is outlined in figure 6.8.

The baseline operating procedure above only works if the corridor in front of the robot is unobstructed. During nominal operations, the goal point selected for the A* planner is 5 meters in front of the robot, centered between the walls of the sensed corridor. If no plan exists which can achieve this goal, the goal is moved half the distance back toward the robot, and the planning procedure repeated. Once the goal is within 1 meter of the robot, the robot decides no path forward is possible and shifts from exploration to egress.

Perhaps the most critical function of Groundhog's software is that of always providing a way to safely exit an underground void. Unlike many robotic applications, how a subterranean modeling robot plans paths, interacts with its environment, and performs exploration change drastically dependent on the high level objective of the robot. When

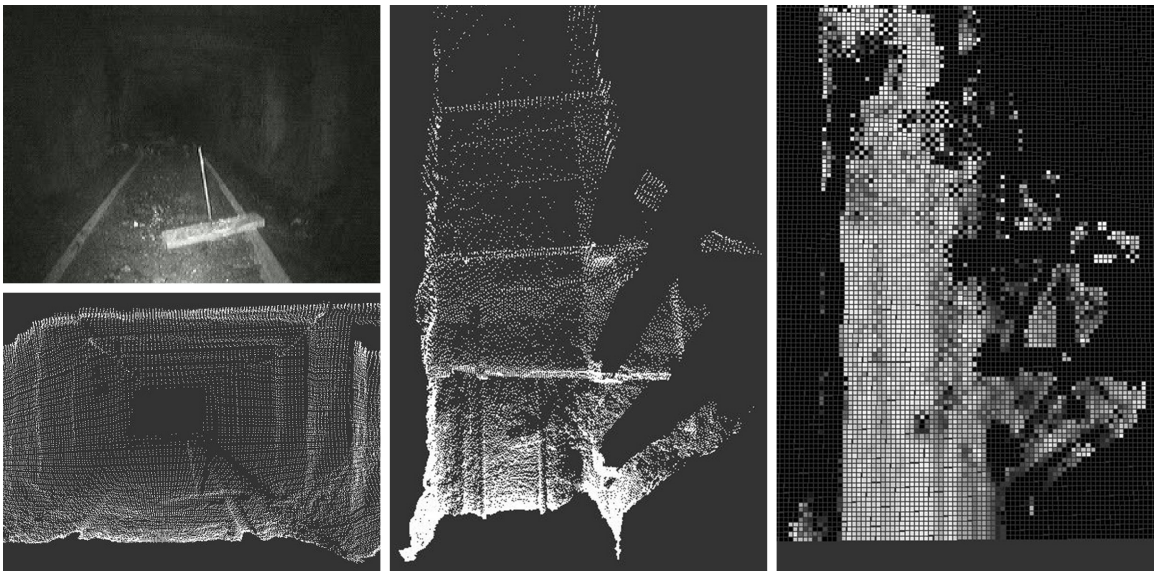

Figure 6.8: Low light video image of fallen roof bolt (top left), perspective view of 3D range scan (bottom left), top view of 3D range scan (center), resulting 2D cost map (right).

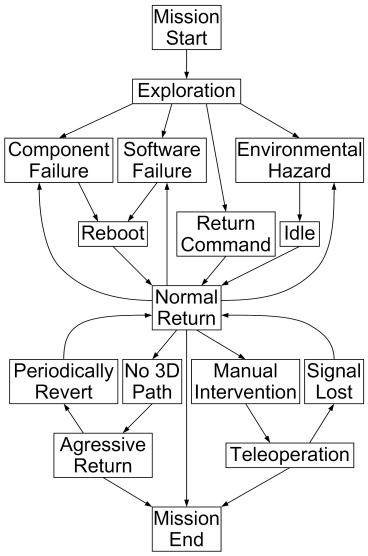

Figure 6.9: Flow chart showing Groundhog's logic for escaping from subterranean voids.

Groundhog is exiting a void, it will take risks and traverse paths that would be deemed untraversable during exploration. A chart representing the contingency options Groundhog explores during egress is shown in figure 6.9.

The procedure outlined above only works for corridors. Once a tee or intersection is encountered, the robots behavior becomes undefined. In order to explore more complex

subterranean spaces Groundhog employs topological maps. A topological map of an underground mine consists of a series of nodes representing intersections with interconnects between nodes representing corridors. The resulting graph can be used for high level path planning, such as planning a route back to the portal. Creating such topological maps during exploration requires a means of effectively identifying intersections. Further, uniquely identify intersections eliminates many of the problems associated with closing large loops, thus improving localization. Loop closure improves localization, modeling and mapping.

Beyond the software required to automate Groundhog's modeling of subterranean voids, there is also a requirement to keep Groundhog functioning in spite of software failures in any of the modules described above. A separate process monitors all of the tasks running on Groundhog, and if any of these processes hangs, terminates unexpectedly or is otherwise identified to be failing, the monitor process restarts the entire software system. Failure of the monitoring process is handled through the hardware watchdog resetting the entire system. This system is still susceptible to an infinite loop or resetting and restarting, but can still be considered to be single fault tolerant.

6.1.5 Groundhog Field Experiments

Groundhog has conducted numerous mapping and modeling experiments in active and abandoned mines. Several of these experiments are outlined below to highlight the interaction between field experiments and configuration evolution and to demonstrate Groundhog's modeling prowess. Field experiments, perhaps more than any other factor, drove Groundhog's configuration to what it is today.

Groundhog's first major modeling experiment occurred at a mine in Burgettstown, Pennsylvania. A surface mining operation had intersected an underground mine that had been long abandoned. The mine operator was interested in examining the exposed portions of the abandoned mine (shown in figure 6.10 and aligning the corridors that had been intersected with a prior hand drawn map. The easy access and visibility from down in the surface mine made this an ideal first test of Groundhog's modeling capability.

At the time, Groundhog was configured as a teleoperated, tethered data gatherer. The primary modeling sensor configuration consisted of two statically mounted line scanners, as shown in figure 6.11. The operator controlled the robot based on forward looking video and streaming data coming from the horizontal laser line scanner. No one is allowed into abandoned mines or near highwalls due to safety concerns. Tethers for data and video were managed by hand. Mapping and model building were all done off-line and off-board the robot.

Despite the unrefined system and crude operational procedures, Groundhog acquired the data necessary to generate some impressive three dimensional models and two dimensional

Figure 6.10: Florence Mine, showing exposed abandoned mine corridors protruding from the surface mining high wall.

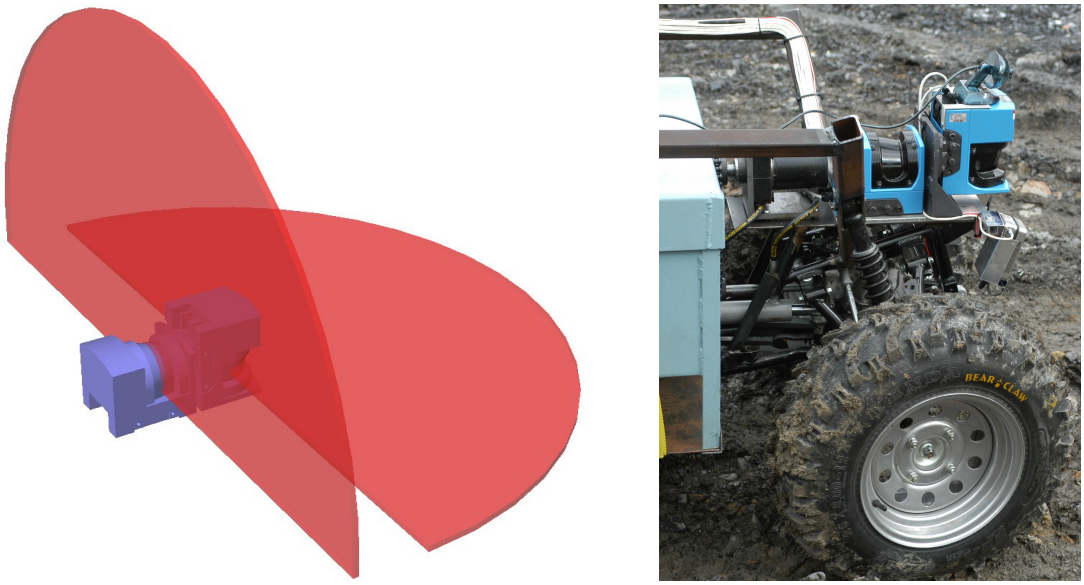

Figure 6.11: Groundhog's initial sensing configuration as conceived (left) and implemented (right).

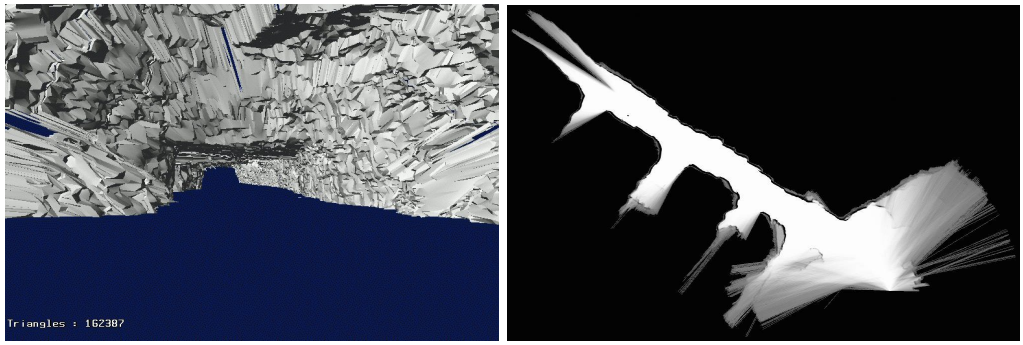

Figure 6.12: 3D shaded polygon model of the Florence mine (left). 2D map of the same area (right).

Figure 6.13: Groundhog at its deepest submergence in the Florence mine.

maps, as shown in figure 6.12. Correlation with video imagery confirmed the accurate mapping of such features as a support pole deep within the abandoned mine. Groundhog's progress into the abandoned mine corridor was limited to only one hundred feet by the water level and downward sloping mine floor. At its furthest extent Groundhog's deck plate, which usually rides roughly a foot above the ground, was completely submerged, as shown in figure 6.13.

The Florence experiment demonstrated the potential for robotic modeling of subter-

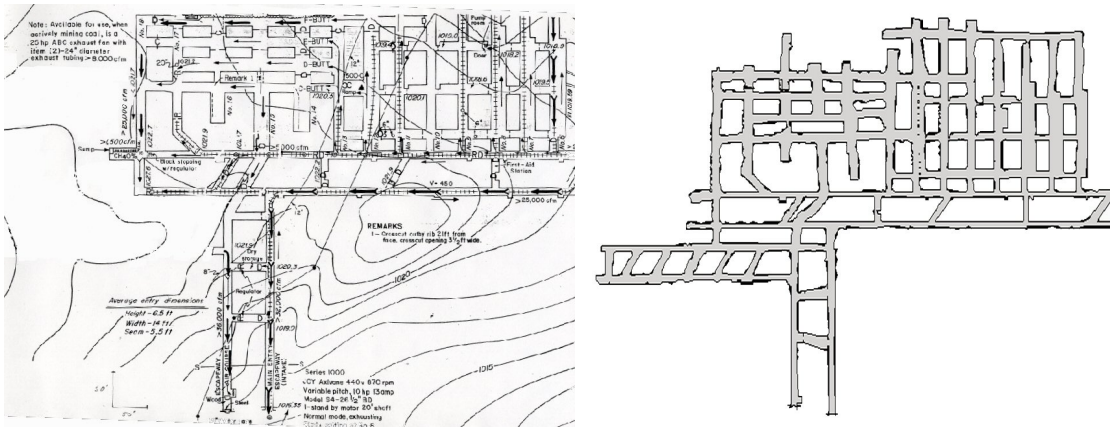

Figure 6.14: The Bruceton research mine: idealized hand surveyed mine map (left), robotically generated subterranean map (right)

ranean voids, but was not designed to traverse long range or to create extensive models. An endurance test at the Bruceton Safety Research Coal Mine pursued objectives of long range and extensive modeling. Due to the small size of the Bruceton mine (relative to most coal mines), it was possible to execute a full coverage modeling task. Samples of the resulting maps and models are shown in figure 6.14. The comparison between human and robot generated maps both verified the correctness of the robotic modeling techniques employed by Groundhog and demonstrated the superiority of machine generated maps for providing detailed non-idealized representations of subterranean voids. This is most clearly demonstrated by the diagonal cross cut in the center of the map, idealized as an angled cut but actually a wide s-curve. The Bruceton experiments also demonstrated Groundhog's linear range to be greater than a mile, but insufficient for its next planned abandoned mine exploration.

In preparation for exploring the Matthies abandoned mine, Groundhog's software was upgraded from teleoperation to autonomous corridor exploration and its energy storage (and therefore range) was improved by roughly 50%. The conditions and objectives of exploring the Matthies mine were previously described in section 3.3.3. The area of interest within the Matthies mine could be thought of as a tunning fork shape, with a short passageway leading from the downstream portal to a Y-intersection with both branches of the Y proceeding in parallel for a long distance to the two side-by-side upstream portals. Modeling experiments were conducted starting at all three avialable portals.

Groundhog conducted a total of nine experiments at the Mathies abandoned mine. A composite map of all data acquired across the experiments is shown in figure 6.15. Table 6.1 summarizes the experiments, what modeling they achieved, and how each experiment ended. The Mathies experiments demonstrated autonomous subterranean modeling for the

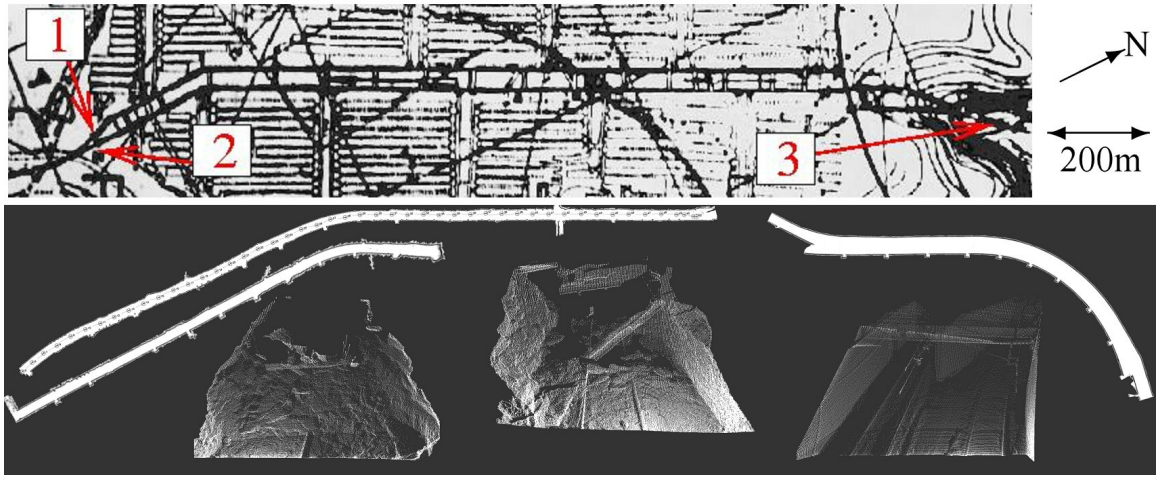

Figure 6.15: Composite robotically-generated map of the Mathies abandoned mine.

first time.

#	Goal	Portal	Mapping	Termination	Egress
1	500m	1	130m	Fallen roof beam	Manual
2	100m	2	100m	Mission Complete	Autonomous
3	100m	1	100m	Mission Complete	Autonomous
4	100m	3	60m	Submergence	Autonomous
5	500m	2	140m	Roof-fall	Autonomous
6	100m	3	20m	Navigation Malfunction	Manual
7	100m	3	10m	Navigation Malfunction	Manual
8	500m	3	230m	Fallen cable	Mixed

Table 6.1: Groundhog modeling experiments in the Mathies mine.

From a configuration stand point, the Mathies experiments were useful in that they demonstrated many of the higher level autonomy functions necessary to succeed underground. Groundhog correctly turned back when it encountered a fallen roof beam and a near complete collapse of the roof. It also navigated around a fallen roof bolt, choose to turn back rather than ensnare itself in a low hanging cable, and proceeded with caution when driving in water of unknown depth. Some of the obstacles encountered are shown in image 6.16. Some limitations in Groundhog's configuration were also revealed, such as the open differentials resulting in loss of traction on steep side slopes.

One interesting behavior observed during the Mathies experiments was the interaction between laser range finders, water, and path planning. The downstream portal at Mathies contained a mixture of standing or slowly flowing water, mud, and some raised railroad tracks, as shown in figure 6.17. The inconsistencies in modeling related to water resulted in Groundhog's path planning software driving from one mound of mud to the next while

Figure 6.16: Obstacles encountered by Groundhog during the exploration of the Mathies abandoned mine. Fallen roof beam (left) and complete roof collapse (right)

Figure 6.17: Wet, muddy conditions found at the downstream Mathies portal.

avoiding the areas with standing water. The resulting slaloming path resulted in Groundhog driving perpendicular to the direction of the tunnel more often than parallel. To understand what caused this behavior and correct it requires a deeper understanding of the interactions between lasers and water.

One of four things can occur when a laser beam intersects water or a transparent surface

such as glass. One, the laser can reflect off the surface of the water, providing range information to the top of the water. Two, the laser can penetrate the water, reflecting off and ranging to the bottom of the puddle or pond. Three, the laser can reflect off the surface of water, returning a range to some other part of the void as if the water surface was a mirror. And four, the laser can glance off the water surface and fail to provide any return. The first two cases return measurements with some validity, although not of the same surface, while the last two cases are incorrect data that should be discarded. Which of these conditions occurs depends on the incidence angle between laser and water, the conditions in the water, and the type of laser range finder used. The larger the incidence angle, the more likely the laser is to bounce off the top of the water and return an invalid measurement. The more sediment in or movement of the water the more likely a return from the surface of the water rather than to the soil beneath the water.

Laser range finders can be classified as first return, last return, or multiple return. First return laser range finders measure the time between the start of the transmitted laser pulse and the first time the reflected return energy goes above a certain threshold. Last return lasers use the same start point, but measure to the last reflected pulse within a certain time frame. Multiple return systems either produce a reading every time the return energy crosses the threshold or the entire analog record of reflected laser light intensity. Consider the example of a laser looking straight down into a pool of water. A first return system will only record the reflection of the top of the pool, a last return system will only return the reflection of the bottom of the pool, while a multiple return system will provide both data points and possible secondary reflections.

Groundhog's fixed, low vantage point means that incidence angle increases the further out into the environment the scanner is aimed. Groundhog's lasers are first return range finders and where operating in very muddy often flowing water at the Mathies mine. The combination of these attributes resulted in Groundhog seeing the first few feet of ground in front of it as a combination of clumps of mud and surfaces of puddles, but further out Groundhog observed huge holes in the environment caused by glancing range finder pulses. This resulted in Groundhog planning an effective path within the first few feet of its position, then directing the path to the largest protruding soil mound in the region lacking data. Repeated sensing and planning cycles slowly advanced Groundhog's position, but in a weaving course going from hillock to hillock.

Software upgrades during the course of the nine Mathies experiments gave Groundhog the ability to infer the existence of a water level from the spurious data points obtained from glancing laser pulses. In effect, a plane roughly level with the known location of the bottom of Groundhog's wheels is fitted to data points that correspond to reading from the top of the water's surface. Spurious data points significantly below this plane or non-returned range finding pulses are projected back to the plane, filling in missing data and allowing

Groundhog to plan and effective and efficient course that drive through rather than around puddles in the environment. The obvious problem with this approach is that the robot now has the potential to drive through the interpolated region and into a fatal hole covered by the water. No effective solution for safe planning with water has yet been developed, as laser and radar are usually blocked by water and sonar is only effective if it is submerged in the water, requiring a deployable sensor of significant complexity.

The Mathies experiments exhibited core components of corridor exploration, but did not demonstrate exploration of multiply connected room and pillar topologies, common in many mines. Two separate experiments demonstrated the component technologies necessary to produce extensive three dimensional models of subterranean voids with complex geometries. Additional experimentation in the Bruceton mine demonstrated topological map exploration, while a field experiment at the Dakota coal mine in West Virginia demonstrated large model integration.

Experimentation at the Bruceton mine resulted in 200 meters of mapped and modeled mine with all intersections identified and fully explored. These experiments demonstrated critical technologies such as identifying intersections, finding unique representations of intersections, and fully exploring a network of corridors. Side corridors to intersections were explored by driving fully past an intersection, then reversing direction to back into an offshoot. This allows Groundhog to image both sides of the crosscut before committing to entering it, enabling better detection of deadends not worth exploring and improved path planning for taking the corner.

In prior experiments with Groundhog the primary data deliverable of interest consisted of maps of abandoned mines. Models are needed when three dimensional feature are of interest, or when accurate mine heights are desired. In the Dakota mine field experiments mine height was the desired deliverable. Dakota was used as a surrogate for mines such as the Inez abandoned underground coal mine in Kentucky. This mine was located underneath a large reservoir used in processing coal extracted from nearby mining operations. The reservoir contained roughly 250 million gallons of water and 31 million gallons of very small particles of coal called fines. The reservoir breached the abandoned coal mine underneath it, dumping its contents into two different water sheds and destroying numerous homes. Subsequent investigation determined the cause of the failure was insufficient material thickness between the reservoir bottom and underground mine roof due to an inaccurate understanding of mine roof elevation.

Groundhog's field experiment at the Dakota mine obtained mine roof elevation by generating 3D subterranean models from robot acquired data. During two days of operations Groundhog traversed over a mile of mine corridors and acquired 584 3D scans. The individual scans were integrated using two different approaches. Initial estimates of scan location were derived from localization information provided by 2D scan matching. The path taken

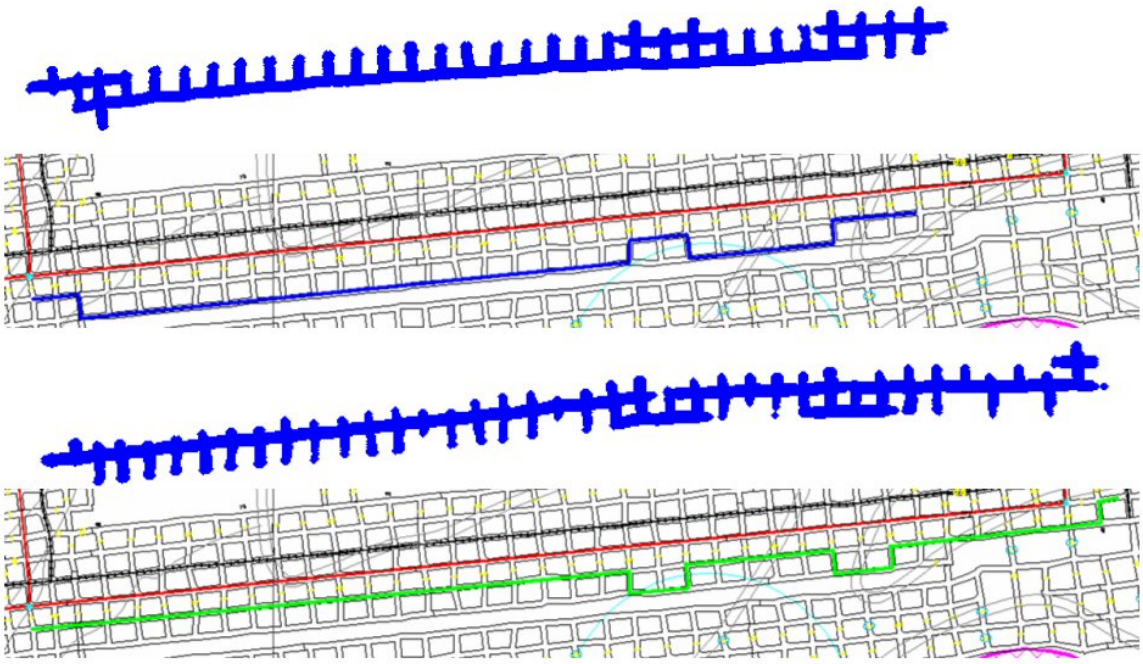

Figure 6.18: Robot acquired maps obtained by following the blue and green paths shown on prior maps.

by Groundhog and the resulting robotically generated maps are shown in figure 6.18. 3D models were integrated either pairwise or by integrating a scan with the composite model of every other scan which contained overlapping data (as determined from estimated scan positions). A comparison of these two approaches is shown in figure 6.19. Clearly utilizing more data during scan integration is beneficial and decreases the drift rate from 6.5% to 1.5%.

Groundhog's field experiments demonstrated many firsts: First robotically acquired three dimensional subterranean models, first autonomous mapping of subterranean voids, and first models based on topological exploration of subterranean voids. These experiments also demonstrated both what parts of Groundhog's configuration succeeded in subterranean environments, and which failed. These lessons learned were applied to the configuration design of a second generation unconstrained entry, unlimited mobility robot, dubbed CaveCrawler.

6.2 Representative Implementation: CaveCrawler

CaveCrawler's drive train, primary mapping sensors, power system, approach to localization, height and weight are all configured to provide superior performance to that possible

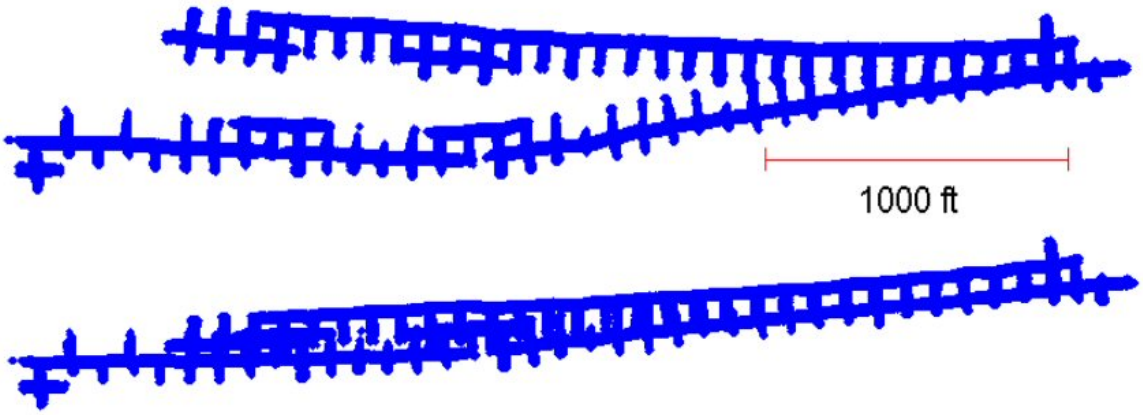

Figure 6.19: Pairwise model integration (top), scan to composite local model integration (bottom).

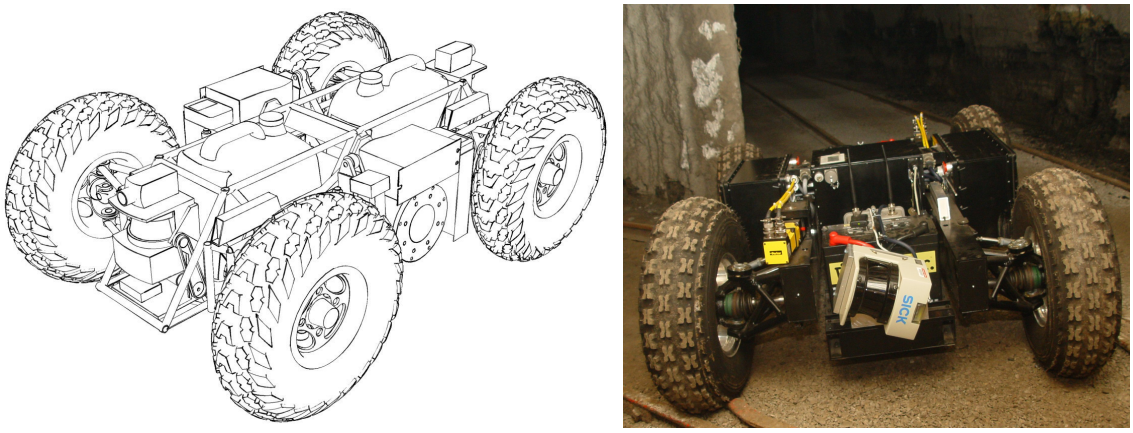

Figure 6.20: CaveCrawler original concept (left) and final implementation (right).

with Groundhog. CaveCrawler, seen in figure 6.20, shares Groundhog's general footprint but differs in locomotion, sensing, and software.

6.2.1 CaveCrawler Configuration Design

Many of the configuration decisions made during the development of Groundhog were reapplied to CaveCrawler. Much of CaveCrawler's configuration was determined by observations of Groundhog's performance during field experiments. Driving and sensing symmetry was one of the best configuration decisions made in developing Groundhog. Locomotion symmetry enabled easy, efficient exploration of dead-ends, while sensing symmetry allowed full modeling of cross corridors prior to entry. CaveCrawler's configuration features the same front/back symmetry found in Groundhog plus additional symmetry about the center of the

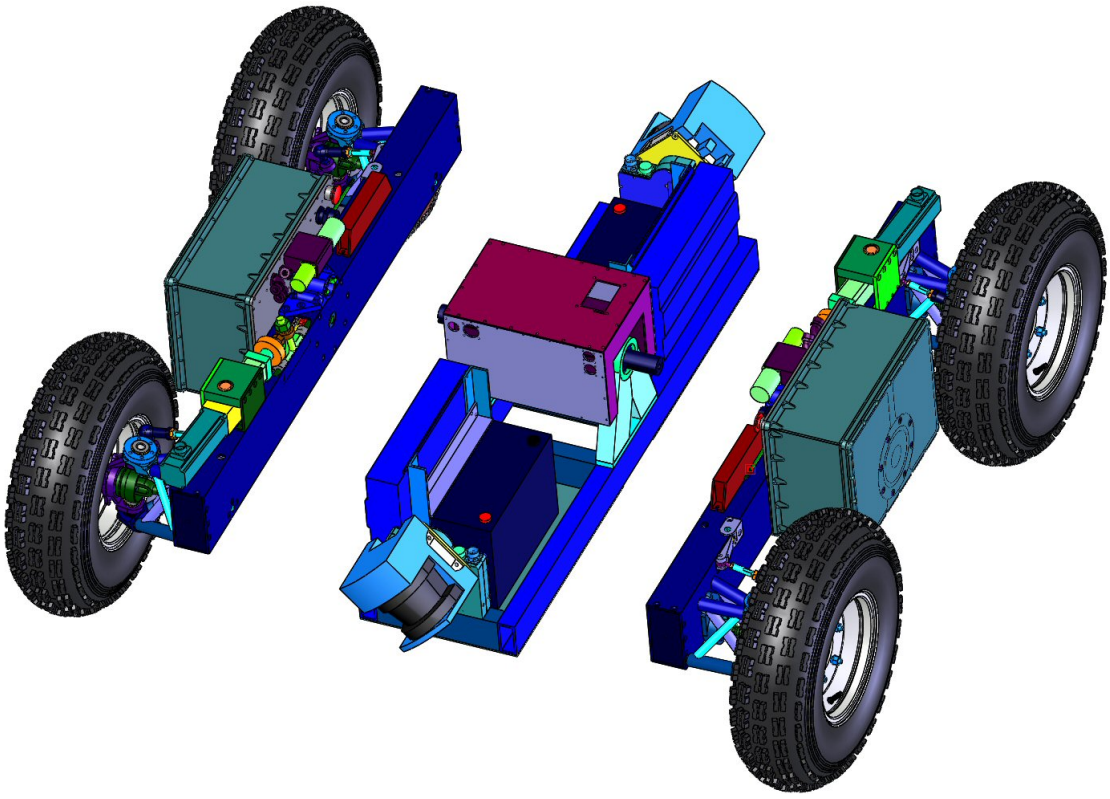

Figure 6.21: CaveCrawler's three part design.

robot. Groundhog's rapid development (initial concept to first field experiment in under sixty days) resulted in almost every part being unique with no interchangeability between components. CaveCrawler employs a three part chassis, with all locomotion components contained in two identical side frames, as shown in figure 6.21. This core configuration decision resulted in improved manufacturability through significantly decreased unique part count. Unique components such as the central computer, primary electronics bay and power source are all housed in the center section.

Groundhog's core configuration decisions resulted in a simple platform incapable of fine control. Speeds were limited to two discrete steps, steering angle feedback was crude, and Groundhog's only response to rolling down hill was an immediate, uncontrolled skidding stop. To overcome these limitations, CaveCrawler is configured to provide continuously variable speed control, and continuously variable steering angle control. CaveCrawler achieves this improved control by employing closed loop feedback controllers and components instead of Groundhog's open loop, bang-bang control.

CaveCrawler also improves upon Groundhog's sensing configuration. Experimentation with Groundhog demonstrated that more 3D data improves scan matching and model quality be reducing drift. Therefore, CaveCrawler is configured to provide nearly complete

modeling coverage with more uniform point density. In addition to improving sensing coverage, CaveCrawler's sensor configuration also decreases the time required to take a scan and enables dynamic detection of obstacles, something not possible with Groundhog.

The combination of variable, finely controllable robot speed with continuously scanning sensors enables CaveCrawler to operate in a fundamentally different regime from Groundhog. Both systems have configurations that can operate under the sense-plan-act model, where sensing and planning occur while the robot is stationary, and acting utilizes a highly reduced set of sensors and algorithms. CaveCrawler's configuration enables continuous motion, where the sense-plan-act loop occurs in a high frequency loop. CaveCrawler's continuously rotating sensors can be down sampled to provide 2D scans similar to those used by Groundhog during its act phase while also providing information about upcoming obstacles. Variable speed control combined with more flexible software allows CaveCrawler to deliberately and methodically pick its way through dense debris, while rapidly speeding over uncluttered terrain. Every minute a robot spends idling consumes power that could be better spent exploring, thus continuous dynamic exploration and modeling effectively increases range. Also, while many subterranean modeling tasks are unconstrained with respect to time, some, such as miner rescue, are heavily time dependent.

In addition to acquiring models faster, CaveCrawler is also configured to provide higher accuracy models through improved localization. Improved inertial sensing provides more information for determining pose. More importantly, CaveCrawler utilizes survey-based localization whenever possible. A 360 degree prism mounted to a swing arm on-board CaveCrawler allows for very accurate robot positions to be acquired by using an external survey instrument, as outlined in 3.4.1. The survey information also eliminates the need for other global referencing methods and virtually eliminates drift in the resulting models.

At the core, CaveCrawler and Groundhog were designed and constructed under fundamental different configuration design philosophies. Groundhog was developed to quickly demonstrate 'Firsts', proving to the world that subterranean modeling robots were both possible and useful. Groundhog's design emphasizes minimal solutions and simplicity at the cost of flexibility and precision. CaveCrawler was engineered to provide a platform for developing new configurations, algorithms, and operational approaches to solving subterranean modeling problems. CaveCrawler's design emphasizes flexibility, modularity and controllability at the cost of simplicity. Both philosophies have advantages and drawbacks, and both provide additional insight into effective subterranean modeling. CaveCrawler's configuration and its relationship to subterranean modeling capabilities is outlined below, as Groundhog's was above.

6.2.2 CaveCrawler Mechanical Design

CaveCrawler is a four wheeled device, with all four wheels both powered and steered. CaveCrawler's three part construction results in a left sideframe, center section and right sideframe. Each sideframe drives and steers a pair of wheels. Each pair of wheels has a single motor providing drive torque and a linear actuator controlling steering angle. The system of tie rods and crank arms coupling the two wheels to the linear actuator ensure that they are always on a circle centered along the line perpendicular to the direction of motion and intersecting the center of the robot. The single motor guarantees both wheels rotate at the same speed and always provide traction, even if one wheel is not in contact with the ground. The only coupling between the motors and linear actuators of the two sideframes come from the central computer. In effect, the computer acts as an intelligent electronic differential, ensuring that steering angles and motor speeds result in smooth, no-slip robot motions.

CaveCrawler's configuration allows two methods of changing direction, Ackerman steering or differential steering. For maximally efficient operations, CaveCrawler's steering and drive should be configured as a dual Ackerman steered system, identical to Groundhog's configuration. Differential steering, common in tracked vehicles, consists of driving the two sides of a vehicle at different speeds. This creates additional frictional forces not present in an Ackerman configuration as the motion of the contact patch is not parallel with the track of tire treads. Differential steering also allows for turning in place by driving the two sides of a vehicle in opposite directions. On CaveCrawler, steering the wheels outward until all wheels are aligned with a single circle concentric with the robot center would result in a differential point turn with no energy lost to drag. Geometric and mechanical constraints limit the maximum outward steering angle of CaveCrawler's wheels to that shown in figure 6.22, resulting in reduced but not eliminated power loss to friction. The final configuration of interest is to steer all the wheels to their maximum inward extent, thus providing maximum static drag. This configuration is useful as a zero energy parking brake on hills and for making descents down loose rubble piles more controlled by effectively increasing the control surface available for braking and directing the robot.

Each pair of wheels is powered by a high pole count industrial servo motor. The winding configuration of the motor results in high torque at low speed in a large diameter, short stack configuration as opposed to the more common small diameter long stack servo motors. The motor is coupled through a custom 3.5:1 planetary gear stage to a coaxial double sprocket. Each wheel is coupled to the double sprocket through a chain drive, providing a final gear reduction of 41:15, as shown in figure 6.23. The direct coupling of both wheels to a single motor eliminates the possibility of loss of traction, thus prevent a condition like that observed in Groundhog where a steep side slope combined with open differentials resulted

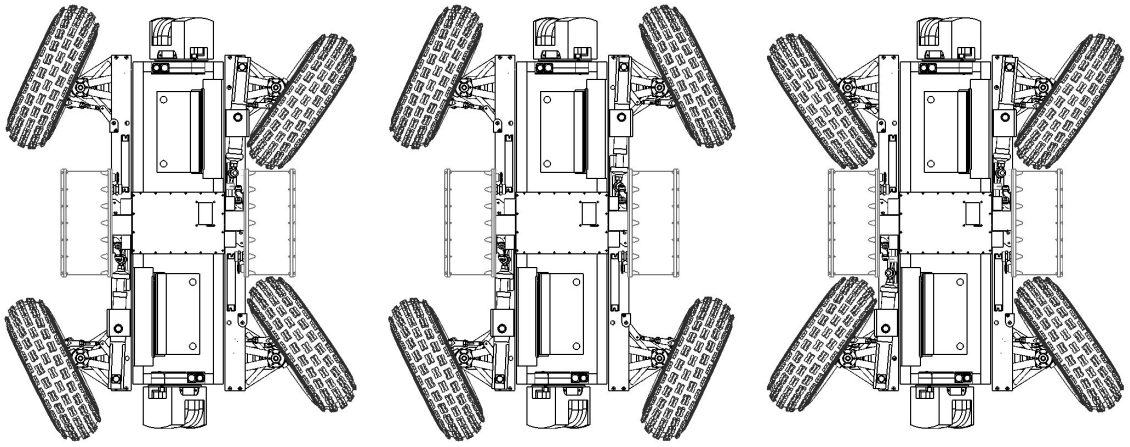

Figure 6.22: CaveCrawler's possible steering configurations: All wheel Ackerman steering (left), point turn (center), parking brake (right)

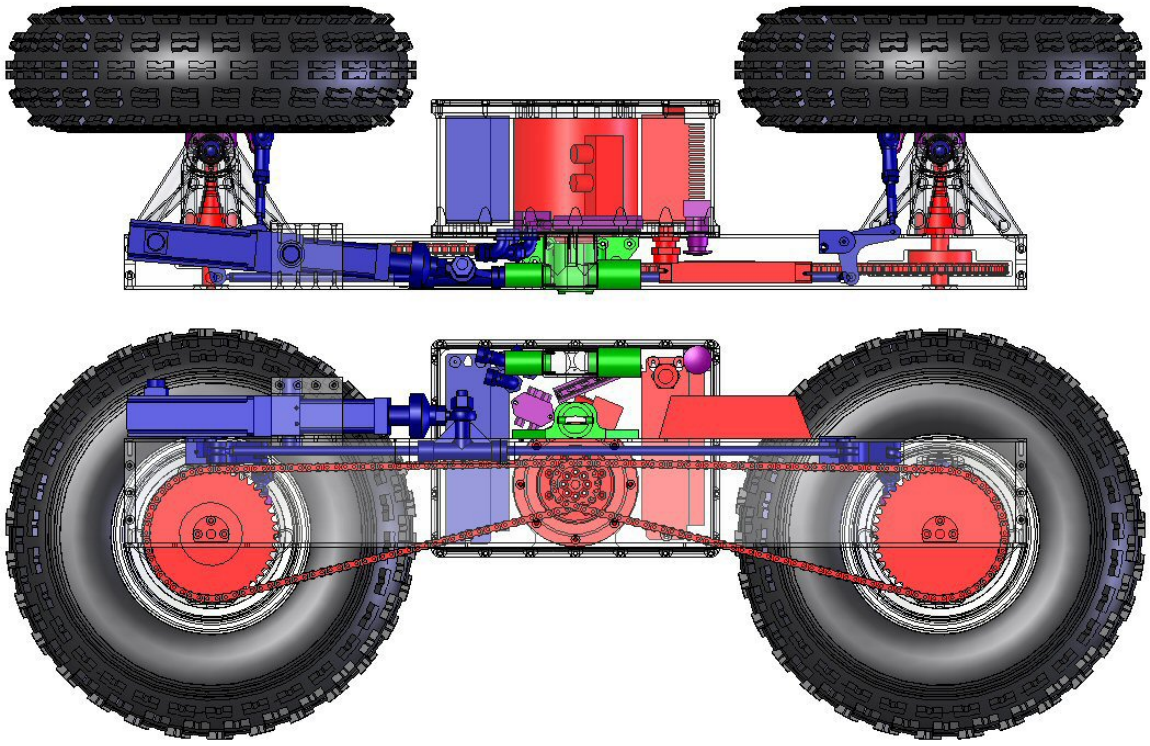

Figure 6.23: CaveCrawler's locomotion actuators and their coupling to the wheels (shown for one sideframe).

in a loss of tractive power.

CaveCrawler's suspension system also improves traction and mobility through increased compliance with the terrain. CaveCrawler's suspension is entirely passive. Passive suspen-

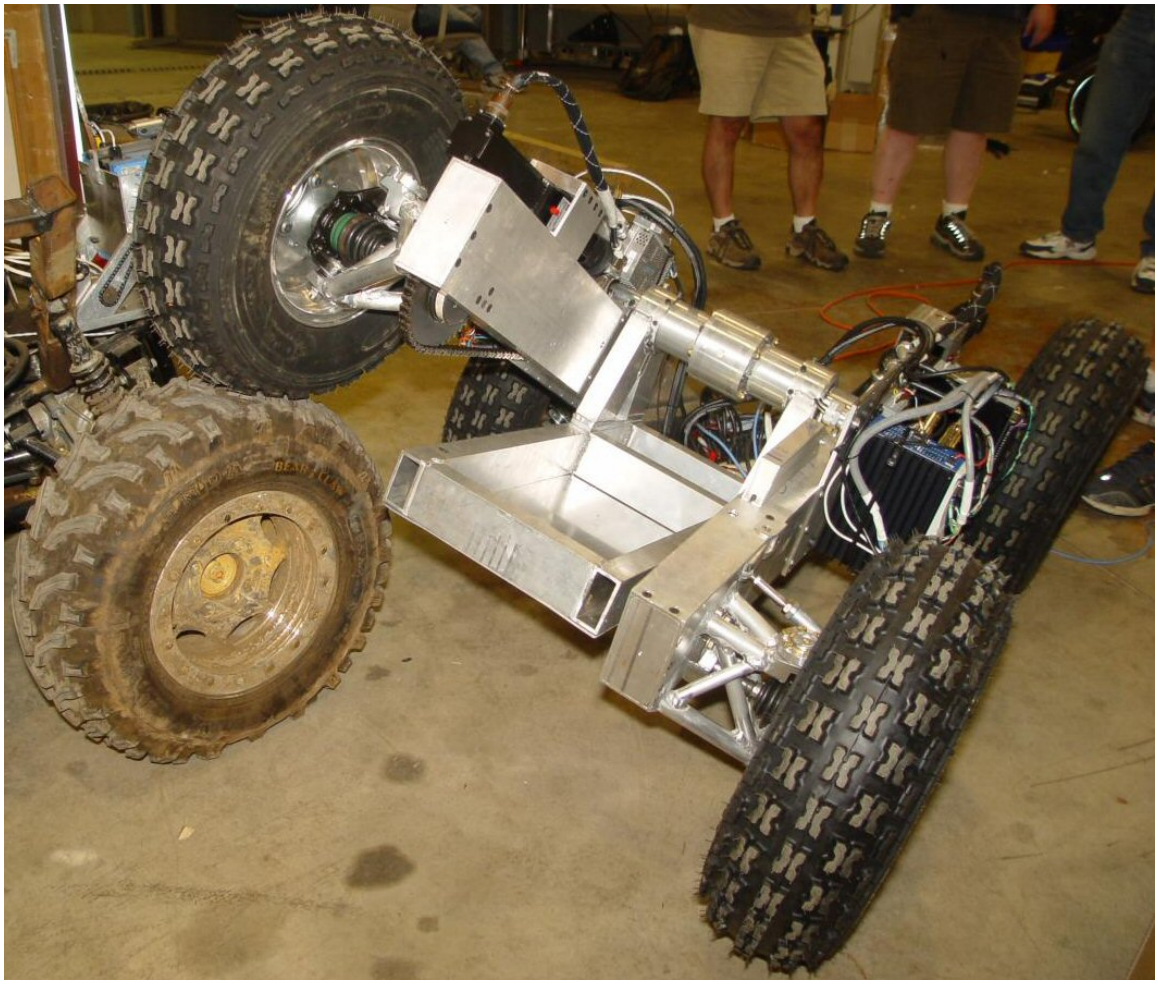

Figure 6.24: CaveCrawler demonstrating its body averaging suspension while driving one wheel up over Groundhog's tire.

sions have shown their effectiveness in applications such as planetary exploration robots, but have serious drawbacks once machine speeds move out of the quasi-static realm and into the dynamic range. CaveCrawler's top speed, while higher than Groundhog and many space robots, is still low enough to allow a purely passive suspension. CaveCrawler employs a body averaging suspension which keeps the center section at half the pitch angle between the two side frames. This allows for large excursions of a single wheel while traversing an obstacle while maintaining overall system ground clearance, as shown in figure 6.24.

CaveCrawler's suspension was designed both to provide compliance and to reduce CaveCrawler's height to less than two feet tall. On multiple occasions Groundhog turned back from a corridor due to low hanging cables. While this is the correct decision for Groundhog's software to make, it is necessitated by Groundhog's height. Groundhog's height was driven by required suspension tie points and the large central electronics enclosure. Cave-

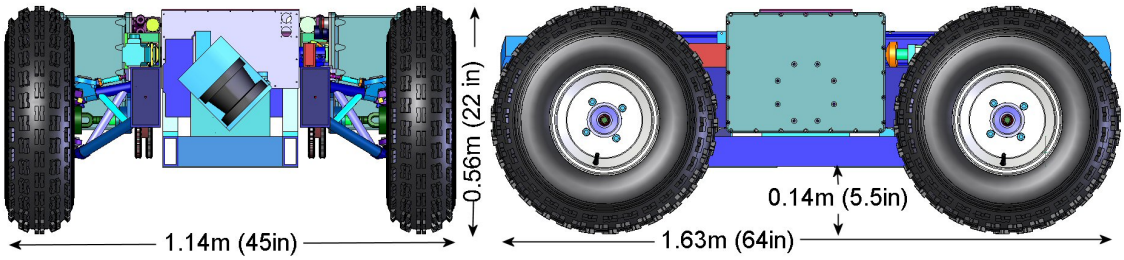

Figure 6.25: CaveCrawler's low profile enables operations in thin seam mines.

Crawler's purely passive suspension removes the need for elevated suspension tie points, while its distributed electronics enclosures eliminate the need for large centralized volumes. Ultimately, CaveCrawler's height is determined by its wheels, as shown in figure 6.25.

In addition to height, CaveCrawler improves upon Groundhog's configuration in a number of areas. CaveCrawler is less than half as heavy as Groundhog, providing more range for the same energy expenditure. CaveCrawler's wheel base is a few inches wider than Groundhog's, but close to a foot shorter. This improves the capability for point turns and improves break-over angle. CaveCrawler top speed is roughly five times Groundhog's 1/3 mile per hour operating speed and CaveCrawler's turning radius (when Ackerman steering) is less than Groundhog's. CaveCrawler's unique design did require some unique sacrifices relative to Groundhog. Ground clearance was reduced slightly due to required component volumes conflicting with CaveCrawler's reduced height. Additionally, CaveCrawler's total energy storage is less than half that of Groundhog. Finally, CaveCrawler is more expensive due to higher quality components and more custom components as compared to Groundhog. A summary comparing Groundhog and CaveCrawler is shown in table 6.2.

Characteristic	Groundhog	CaveCrawler
Height, m (in)	0.91 (36)	0.61 (24)
Length, m (in)	2.03 (80)	1.63 (64)
Width, m (in)	1.07 (42)	1.14 (45)
Clearance, m (in)	0.23 (9)	0.13 (5)
Turning Radius, m (in)	2.13 (84)	1.52 (60)
Point Turn	No	Yes
Degrees of Freedom	2	4
Top Speed, m/s (mph)	0.45 (1)	0.91 (2)
Mass, kg (lbs)	726 (1600)	340 (750)
Energy (kWh)	6960	2400
Sensing, scanning	3D	3D
Sensing, moving	2D	3D

Table 6.2: Comparison of Groundhog and CaveCrawler figures of merit.

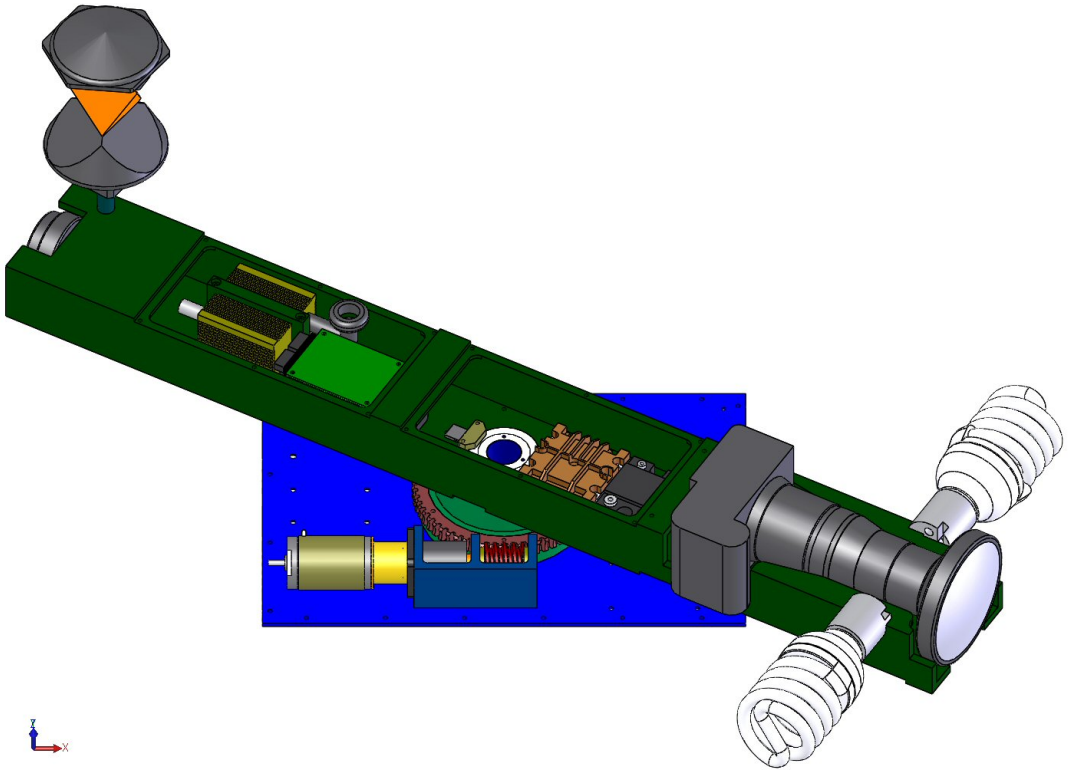

Figure 6.26: CaveCrawler's swing arm, used for survey-based localization.

CaveCrawler was designed with reproducability and manufacturability in mind but at increased financial cost. Groundhog's approach to dealing with eventual robot loss was to minimize platform cost. CaveCrawler minimizes the time required for rebuilding by utilizing symmetry to reduce unique part counts and employing standard components with multiple manufacturers. CaveCrawler's construction is almost entirely out of aluminum to minimize weight and to facilitate machining, versus Groundhog's primarily steel construction.

CaveCrawler has an actuated sensor platform designed to enable survey-based localization that Groundhog does not have. A rotary arm on top of CaveCrawler moves a survey target prism in a large circle. Multiple survey reading to the prism determine the center of the circle, and hence the position of the center of the robot. The swing arm is actuated by a worm gear driven by a servo control brushed DC motor, as shown in figure 6.26. Accurate encoding of angular position at the output allows each survey reading to be tagged, thus providing robot yaw information. The swing arm is also a useful base on which to mount additional still or video imagery components, as its full circle sweep provides complete coverage for secondary imagery.

6.2.3 CaveCrawler Electrical and Sensing Design

Two sealed absorbent glass mat (AGM) lead acid batteries provide 2400 watt-hours of energy to power all of CaveCrawler's systems. These batteries do not vent, and they are certified for air flight, to facilitate rapid deployment of CaveCrawler via air. The parallel wired batteries run a core 12VDC bus. Computing, communications, control and sensing all run from independent DC/DC converters at voltages from 12VDC to 48VDC. The locomotion subsystem is powered by a 120VAC bus generated by an on-board 2kW continuous, 3kW peak inverter. The high voltage AC bus makes possible the use of stock industrial components, reduces required power wiring gage (thus greatly simplifying wire routing), allows wall powered operations when in the lab and provides an on-board source for powering laptops or other equipment useful for field operations and debugging.

Control of drive motors and steering linear actuators is accomplished using stock industrial servo drives wired for 120VAC input. These drives integrate a motor amplifier and closed loop controller and can be controlled via serial port or with analog inputs and provide complete torque, velocity or position loop control without any additional external components. Due to rate limitations on the serial port, CaveCrawler utilizes analog signals generated from analog to digital converter which are part of the PC104 form factor computing stack. An integrated resolver on the drive motor in each side frame provides position and velocity feedback for linear motion, while an integrated potentiometer on the acme screw based linear actuators provides wheel angular position encoding. Both amplifiers are housed in the sideframes, with connections to the center section routed through one data and one power connector. This location maintains the division of locomotion and processing/sensing functionality between sideframes and the center section.

CaveCrawler's central electronics bay houses voltage converters, solid state relays, Universal Serial Bus (USB) hubs, an Ethernet switch, and custom electronics for distributing power and control. CaveCrawler's electronics are designed to be modular. Multiple USB, Ethernet and power connectors are available, allowing rapid change out of sensors and other payloads, as compared to Groundhog's more monolithic approach. CaveCrawler's central computer is a 933Mhz PC104 motherboard which enables and disables each of the four amplifiers (two drive, two steering) as well as individual power supplies to sensors and primary AC power flowing into the sideframes. Additionally, the mixed signal I/O card on CaveCrawler controls the central AC inverter, allowing for all AC power to be shut down when the robot is stationary, thus reducing wasted power. An emergency stop circuit is also wired in series with computer controlled solid state relay which distributes AC power to the sideframes. Feedback from the emergency stop circuit, amplifier fault lines, and inverter control lines allows CaveCrawler a much finer degree of control over its power system than was possible on Groundhog.

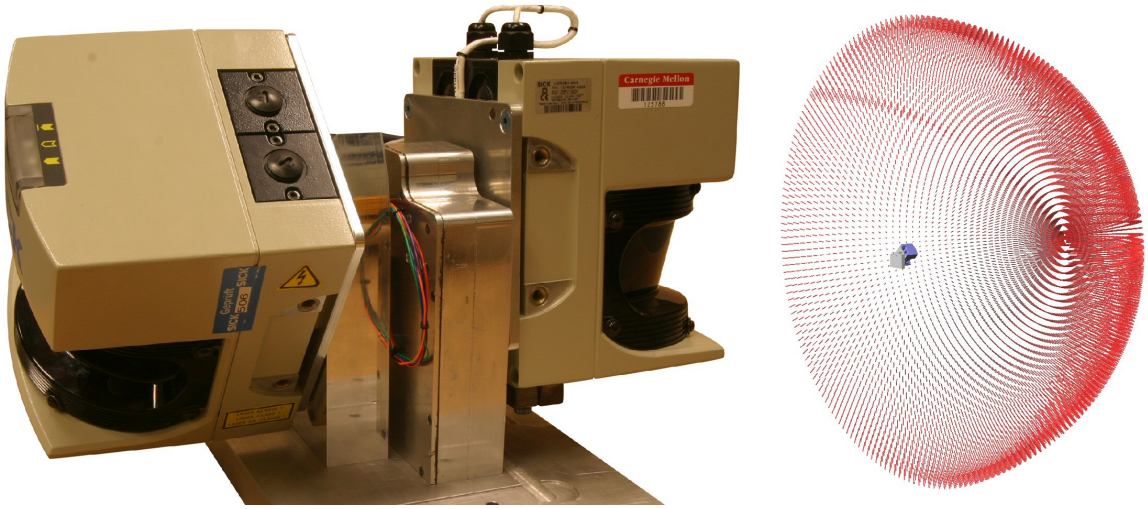

Figure 6.27: Two generations of spinning laser range finder hardware(left). Spinning laser data density distribution (right).

CaveCrawler sensor configuration is specifically designed for creating superior subterranean models. Groundhog's sensing was driven primarily by the requirements of simple scanner configuration and effective path planning. 3D models of the ground in front of the robot were necessary to enable exploration, and were also useful in modeling. CaveCrawler's sensors were designed for modeling enclosed 3D spaces and effectiveness in acquiring data necessary for navigation.

CaveCrawler's primary modeling sensor, shown in figure 6.27, continuously rotates an industrial laser line scanner. Continuous rotation necessitates a slip ring for data and power transmission, but in return reduces required scanner power, results in uniform scan density, and improves dynamic obstacle detection. By balancing the rotating component about the axis of rotation, the system only consumes the power necessary to overcome the friction in the mechanism. Oscillating scanners require significant peak power to rapidly accelerate/decelerate the mass of the scanner at the edges of the field of view. Thus continuous rotation reduces both peak power requirements and the average energy drain associated with running a scanner. The industrial laser line scanner has a relatively large, fixed angular increment. When combined with the slower rotation axis this results in a foveal or spotlight effect with higher data densities near the axis of rotation. This distribution is beneficial for modeling corridors, as it results in a more uniform data density down the length of a modeled corridor, as shown in figure 6.28. It is also useful in navigation, providing maximum understanding of obstacles directly in front of or behind a robot. Dynamic obstacle detection, either due to obstacle motion or robot motion, is improved by having every line scan sample directly in front of the robot and by having two scans of terrain in

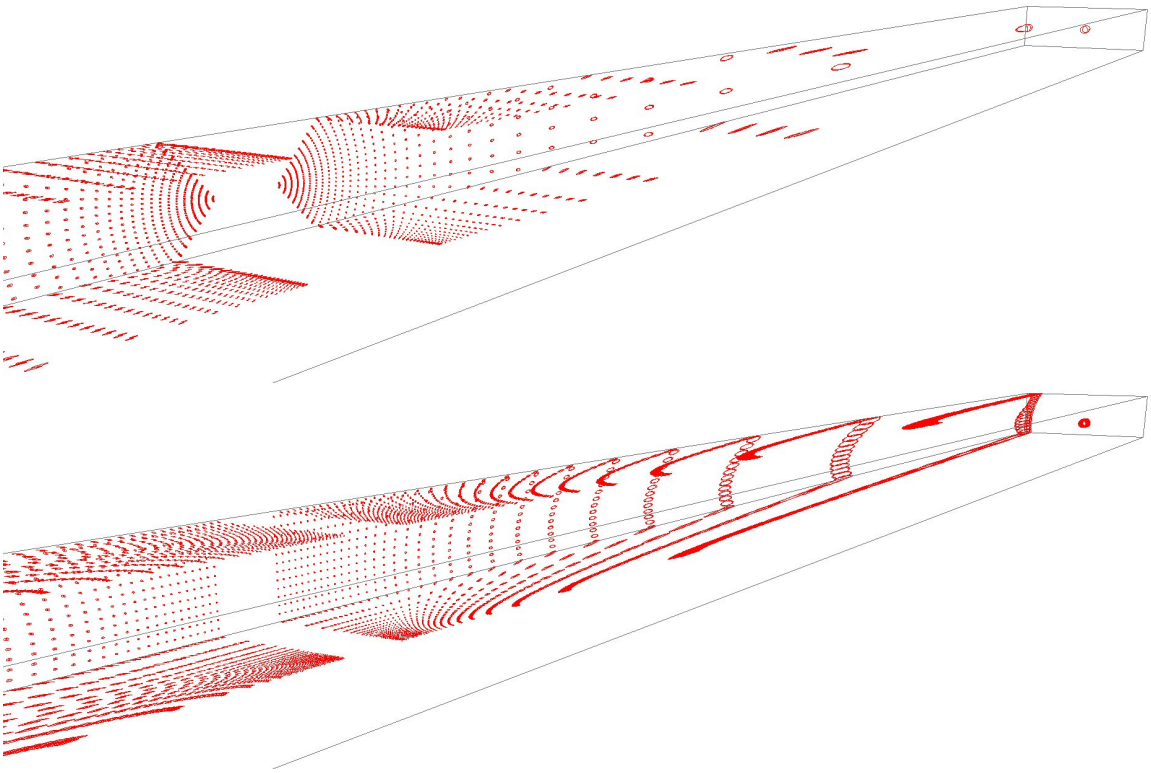

Figure 6.28: Scanner data densities looking down a corridor: Groundhog (top) and CaveCrawler (bottom)

front of the robot for each full line scanner rotation.

The spinning laser modeling sensors and CaveCrawler's survey prism swing arm are both rotated by DC motors controlled via serial motor controllers. In addition to the survey prism, the swing arm also carries a high resolution digital camera fitted with a 183 degree fisheye lens and a low-light wireless video camera. The still camera provides imagery with comparable coverage to the spinning laser modeling sensor. Color data from the camera can be fused with range data from the laser to yield texture mapped polygonal subterranean models, as shown in figure 6.29. The video camera provides full coverage viewing around the robot. When the robot is within wireless transmission range of a base station, this improves teleoperation control relative to that possible on Groundhog.

CaveCrawler's inertial sensing consists of one yaw rate measuring fiber optic gyro and a six axis MEMS based IMU. The fiber optic gyro provides superior sensing to that available with MEMS technology. CaveCrawler relies on its inertial sensing primarily in determining the direction of the gravity vector when it is static, as position and yaw are determined by the survey instrument and swing arm when in range of a human operated survey instrument. When out of survey range, inertial sensing become more critical to robot navigation.

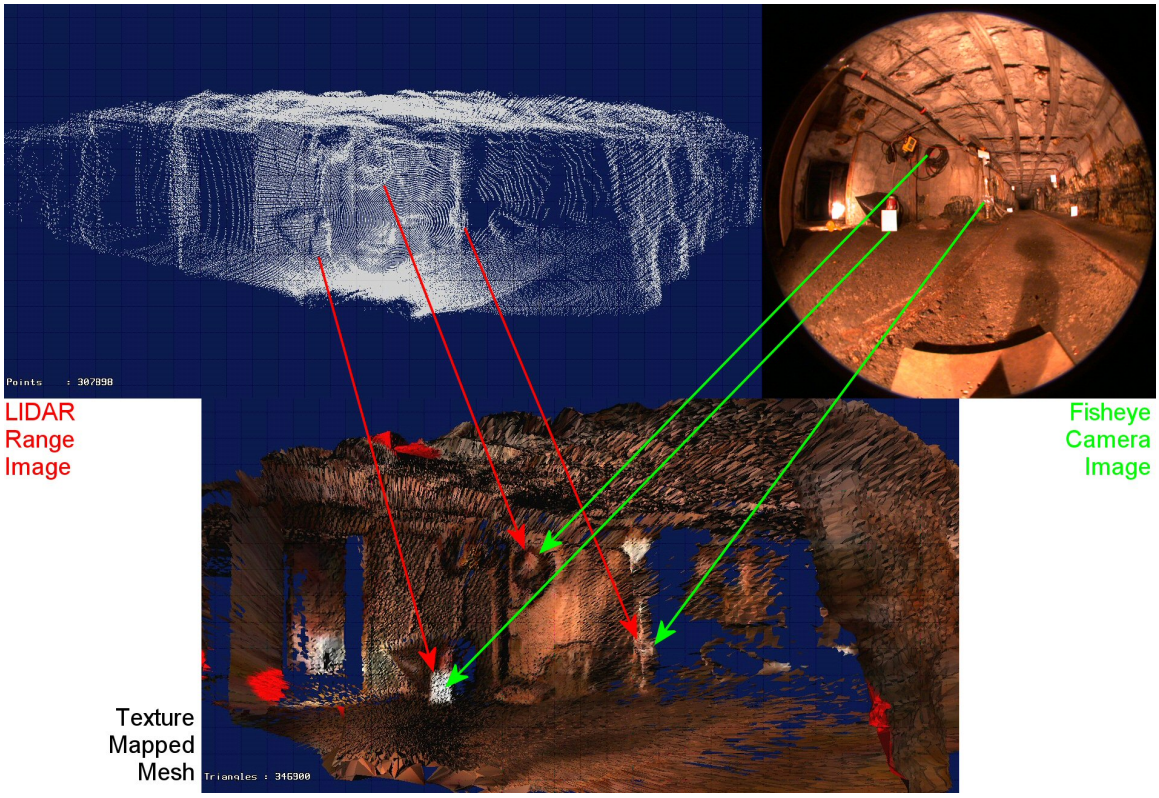

Figure 6.29: Texture mapped 3D mine model (bottom center) resulting from the fusion of a laser range point cloud (top left) and a corresponding fisheye camera image (top right).

6.2.4 CaveCrawler Software Design

CaveCrawler's software has to solve the same problems faced by Groundhog's and as a result has many of the same functional components, although the underlying code is often different. However, as with the rest of CaveCrawler's configuration, the software design is flexible and expandable, in contrast to Groundhog's rapidly implemented but rigid code base. The remainder of this section examines differences and new concepts introduced in CaveCrawler's approach to subterranean modeling and assumes an understanding of Groundhog's software, as detailed in section 6.1.4.

The core software problems facing both CaveCrawler and Groundhog are how to acquire the data necessary for modeling, how to explore subterranean voids, and how to successfully escape from voids after exploration is complete. Successful data acquisition is primarily a function of scanner design and successful interfacing with serial data streams and will not be further discussed. Exploration is driven by knowledge of void topology. Successful egress requires a combination of reliability and robustness, mutually exclusive objectives.

Subterranean exploration can be conducted with or without prior knowledge of void

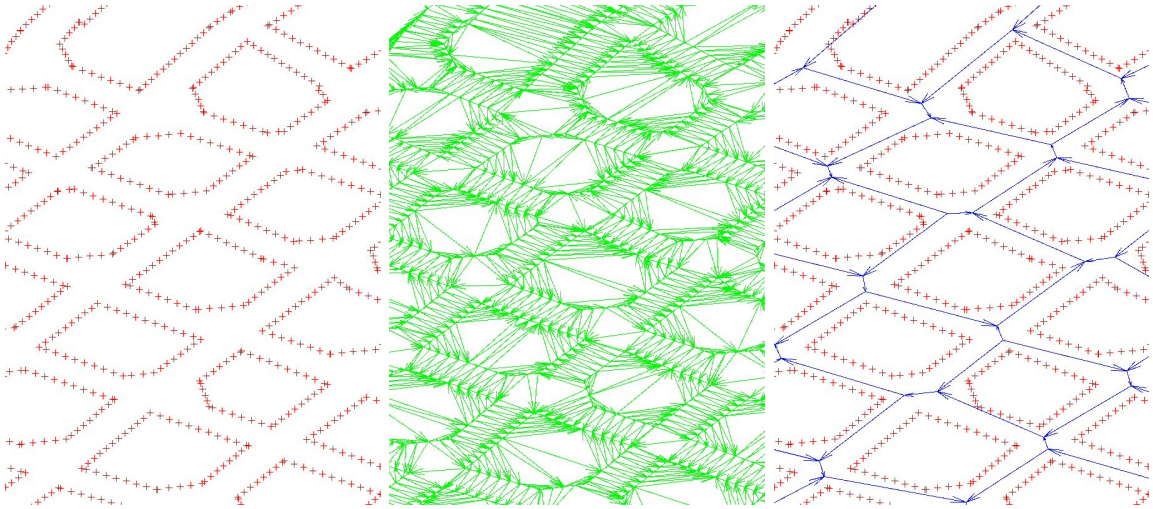

Figure 6.30: Discretized Prior mine map (left), Delaunay triangulization with intersection detection (center), resulting graph (right).

structure. Prior knowledge is unlikely to provide specific, detailed information about local void geometry, but can often provide a first estimate of void topology. Knowing void topology prior to the robot first entering a void allows an initial exploration plan to be developed off-board the robot. This can speed modeling and exploration by providing a shorter path into and out of the void than is possible when exploring blindly. CaveCrawler is configured to take advantage of prior knowledge in the form of electronic mine maps. The same algorithms that are used on board the robot to classify intersections, Delaunay triangulization, can be applied to void maps. The triangulization identifies intersections in the electronic prior map by finding triangles with three long edges, as shown in figure 6.30. Intersections are connected and the topology of the space is encoded using the Voronoi diagram of the space.

Voronoi diagrams are the complement to the Delaunay triangulization, consisting of piecewise linear paths formed by connecting triangle centroids. The key property of a Voronoi diagram is that each point in the diagram is equidistance between nearby obstacles. Therefore, a robot following the Voronoi diagram will drive down the middle of corridors. CaveCrawler uses Voronoi diagrams to determine its desired path through a mine when driving continuously, as shown in figure 6.31. Native use of Delaunay triangulization during navigation eliminates the redundant work done on Groundhog for intersection detection, classification, and unique identification by natively performing these functions.

Reliable egress from voids is a core requirement of any subterranean modeling software system. Relying on the same software used to explore deeper into a void to successfully escape is prone to failure. For example, a narrow passageway or low roof that appeared just

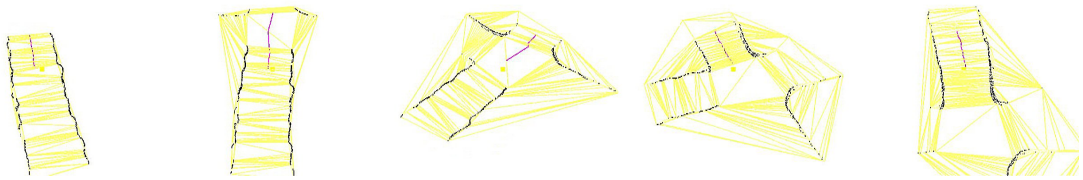

Figure 6.31: CaveCrawler's path planning system turning left at a four-way intersection.

large enough to traverse when exploring the void might appear too small due to different scan locations or slight changes in the environment when exiting the void. Constraints and constants need to be relaxed when existing a void. However, while most subterranean environments can be thought of as static, it is possible for such a space to undergo change during a modeling experiment, and the most likely cause of change is the modeling robot disturbing something in the environment. For example, Groundhog drove over a half fallen cable during one of its experiments. Video from the interaction shows the cable tensing and springing back as Groundhog's wheels forced it temporarily to the ground. This kind of incidental interaction between robot and void is impossible to avoid, and therefore must be accounted for when designing software for effective egress.

Groundhog solves the problem of adaptability to unknown conditions through a deterministic sequence of operational modes, with each mode taking greater risks until conditions for the original standard operating mode are reestablished. CaveCrawler approaches the same problem by having a flexible software framework which allows easy rewiring between various processing components. Loosely built around the ModUtils modular software framework, CaveCrawler responds to a failure of its default mode of operation by rewiring components such as local planners, cost map generators and goal selectors. A different, dynamically selected software configuration has the potential to reroute CaveCrawler around trouble areas. This concept is still in development and has yet to be demonstrated in the field at the time of writing, but the underlying framework has been implemented on CaveCrawler and simulation has demonstrated the feasibility of the idea.

Much of CaveCrawler's core software has yet to be fully developed. Results have instead come from efforts directed toward fully utilizing CaveCrawler's survey locked positioning system first explained in section 3.4.1. CaveCrawler has created some of the most detailed, comprehensive subterranean models yet generated, as outlined in the next section. Software design on CaveCrawler is an ongoing endeavor, with many concepts under development but only a few field tested.

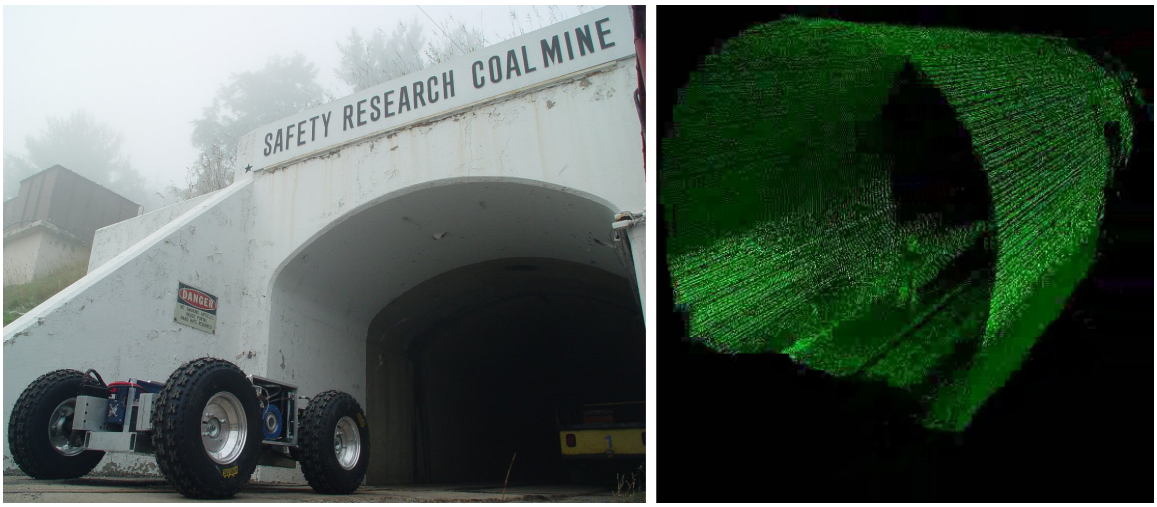

Figure 6.32: CaveCrawler entering the Bruceton mine for first underground testing (left). Spinning laser scanner train tunnel scan with human (right).

6.2.5 CaveCrawler Field experiments

The quality of models obtained in CaveCrawler's field experiments is superior to the quality of Groundhog's models. Integration of survey information enables higher quality models, and much of the developmental effort and field testing has been dedicated to demonstrating and improving survey locked modeling.

Early testing of CaveCrawler in the lab, at the Bruceton research coal mine, and on various rubble piles exhibited the sufficiency of CaveCrawler's locomotion configuration and demonstrated abilities such as traversing a flight of stairs, climbing a rubble pile and driving over fallen beams and up steep curbs. CaveCrawler's spinning scanners were also verified in a number of environments such as train tunnels, sewers, and mines. Some results from subsystem tests are shown in figure 6.32.

The first full system testing of CaveCrawler created a survey locked model of the high-bay space where much of CaveCrawler's development occurred. A robotic total station was leveled and initialized at the corner of two long corridors. CaveCrawler started near the survey instrument (as shown in figure 6.33) and acquired scan data at roughly ten foot increments down the length of both corridors. At the time CaveCrawler's prism spinning swing arm was still under development so two fixed locations on the robot were used to provide survey points. This only constrained the location of the robot along five degrees of freedom. The final degree of freedom was determined from the gravity vector reading supplied by the IMU. A total of X individual static scans, each consisting of data from the front and rearward looking lasers, were integrated into a single model. The resulting model generated is shown in figure 6.34.

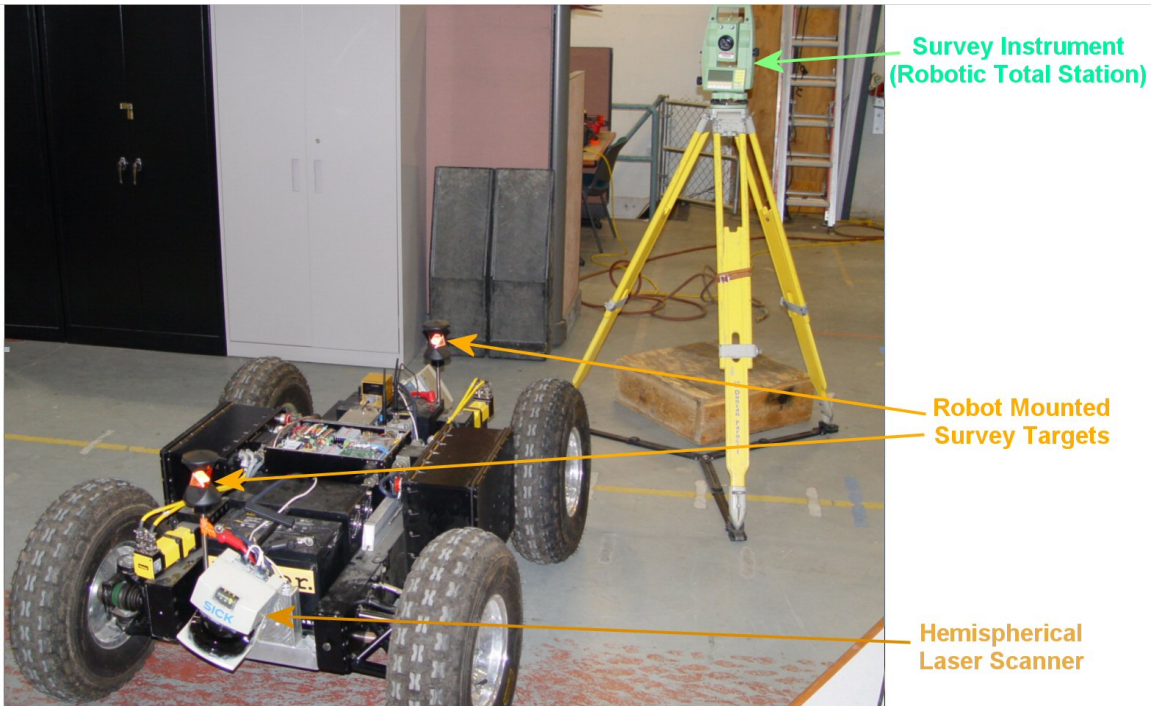

Figure 6.33: CaveCrawler with a statically mounted prism next to a robotic total station survey instrument

Once modeling was demonstrated in the lab, CaveCrawler was deployed into a limestone mine to acquire an accurate model around a few pillars. Traversing around a large obstacle such as a pillar in a mine requires moving both the robot and the survey station. Surveyors commonly perform traverses, which involve accurately surveying to a pole mounted prism located at the desired next location of the survey instrument. After the survey instrument is relocated, a backsight to the prism pole at the original location of the instrument is taken. This process, called carrying a traverse, allows readings taken from the new instrument location to be integrated into the original coordinate system with minimal errors.

Four survey instrument locations provided sufficient coverage to survey in each static scan location used in modeling the Limestone mine interior. Surveying combined with the improved data distribution of the spinning scanners allowed such precise alignment between scans that features such as a narrow steel cable hanging from the ceiling were clearly visible in individual scans and the integrated model. Some views of the resulting model are shown in figure 6.35, but to fully appreciate the accuracy and density of the resulting model requires three dimensional flythroughs which can not be conveyed in this format.

CaveCrawler conducts its most extensive modeling in a thousand-foot long horseshoe shaped tunnel. Not only was the distance longer than any prior or subsequent modeling experiments, the data density was higher as was the overlapping coverage of any given

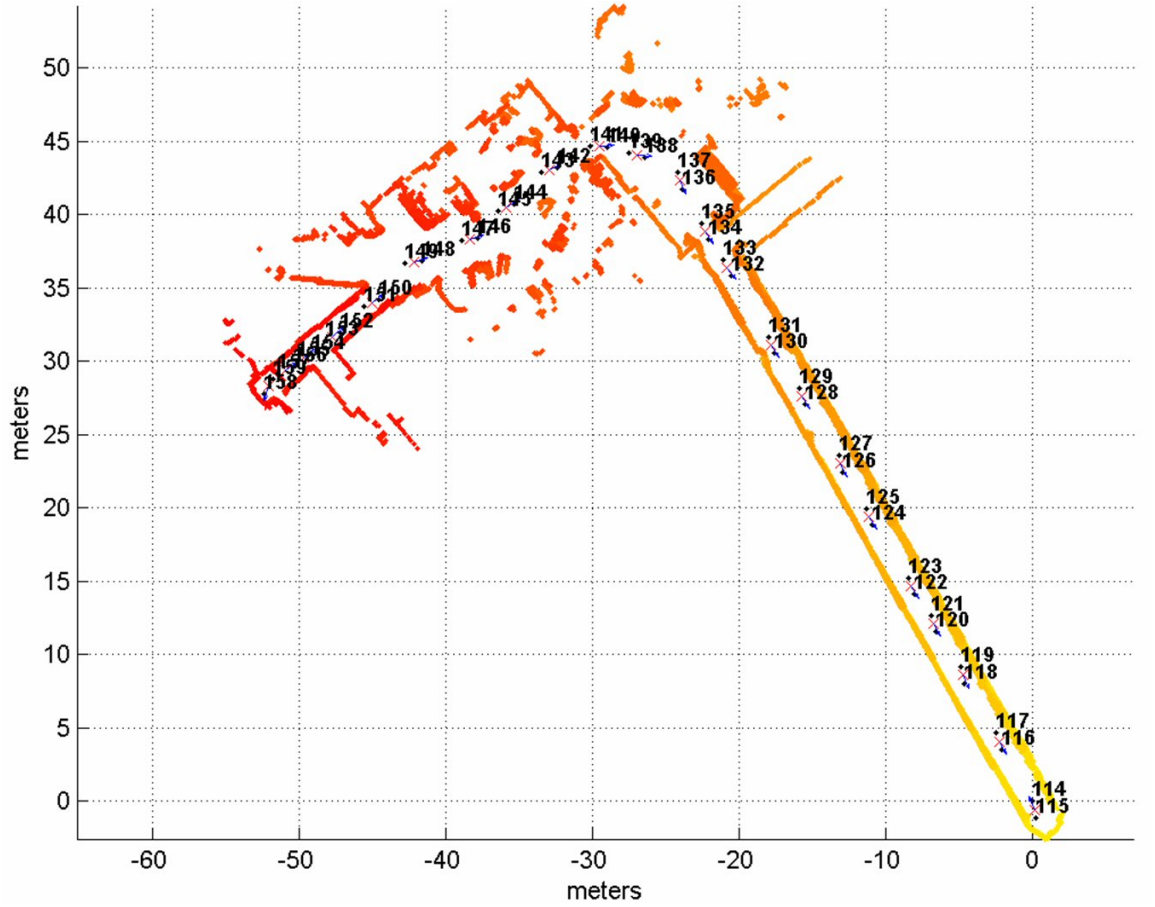

Figure 6.34: Composite high-bay model generated using survey locked subterranean modeling techniques.

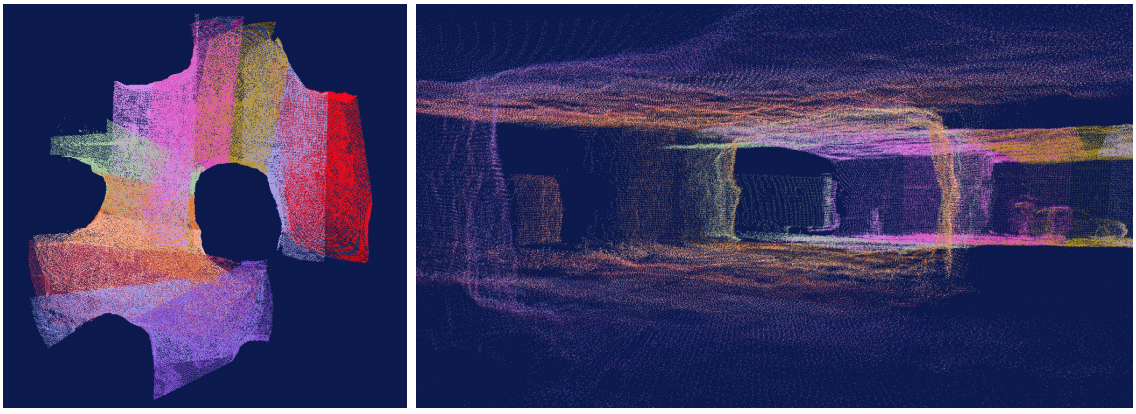

Figure 6.35: Overhead (left) and perspective (right) view of the limestone mine model. Each individual scan is represented by a unique color.

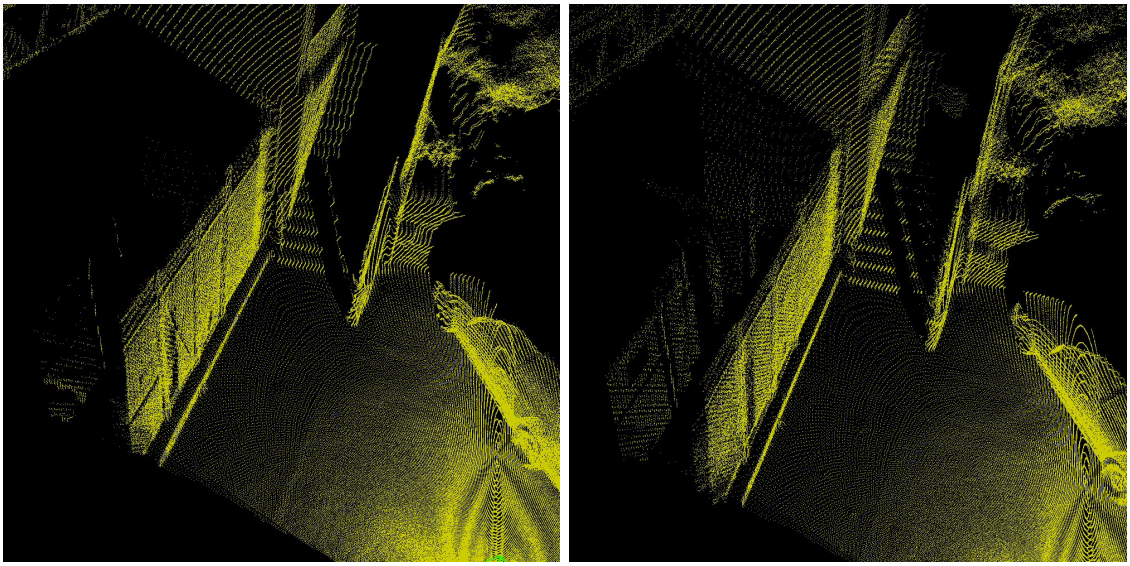

Figure 6.36: Modeling through chainlink: longer range, higher divergence laser (left), shorter range, lower divergence laser (right).

feature within the tunnel. Nearly a half-billion data points were acquired and integrated into models similar to those shown in figure 6.35.

One feature of interest in the horseshoe tunnel was that much of the roof was unsound with numerous cracks and fissures and prone to spalling, a process whereby pieces of rock spontenously break free of the surrounding strata. To protect workers within the tunnel from potential falling rock, chain link or similar netting was secured to the roof using roof bolts. Chain link presents an interesting challenge to modeling, as it is desired to model both the chain link and any void volume behind the chain link. Experiments conducted on a chain link enclosed power substation demonstrated a hidden characteristic of laser range finders that does not appear to have been previously reported in the literature. Scans taken from a fixed location but using two different types of lasers, a Sick LMS200 and Sick LMS291, are shown in figure 6.36. Idealization of laser beams as perfect cylinders is usually acceptable, but real world lasers diverge slightly, usually on the order of hundredths of a degree. For example, the LMS200 has a spot size of X inch at Y while at the same distance the LMS291's spot is Z inches in diameter. A large divergence laser's cone of energy will almost always have a reflection off the chain link. When coupled with a first return laser no data will be obtained from surfaces behind the chain link. Therefore, lower divergence lasers are necessary for modeling through netting and were used on CaveCrawler when modeling the horseshoe tunnel.

The horseshoe tunnel also presented a unique opportunity to test and refine procedures for modeling changes in subterranean voids between sequential surveys. CaveCrawler mod-

eled the void at two different times. Between the first and second modeling significant equipment and instrumentation was installed in the tunnel. Using nearest neighbor classification, as described in section 3.4.4, five models of interest were created of the same tunnel. The most obvious models are the original composite model M_1 and the second, later model M_2. Also of interest is the model containing only points that haven't changed between the two scans $M_{both} = M_1 \cap M_2$, the model containing points unique to the first scan $M_{old} = M_1 \cap \overline{M_2}$, and points unique to the second scan $M_{new} = M_2 \cap \overline{M_1}$. By color coding each of these three types of points, for example M_{both} = red, M_{old} = green and M_{new} = blue, changes made to the tunnel become immediately identifiable. A flythrough movie with the points colored as above was generated from the data gathered and was useful in visualizing tunnel changes, intentional and otherwise.

The models of the horseshoe tunnel were composed of a total of over 400 static scans. Each scan consisted of one complete revolution of the spinning lasers, lasting roughly sixty seconds resulting in double coverage of the area with over a million points per scanner, or two million points per station. Three complete traverses of the tunnel (two originally and one after instrumentation had been installed) provided triple coverage of all tunnel surfaces. In all, over a half-billion data points were obtained and integrated into composite models. The size of the data set necessitated down sampling and segmentation of the data in order to be manipulatable using standard 3D visualization programs on top end processors, however all data was available for review and study. The results represent the largest robotically generated subterranean models created to date. Cross sections along the length of the tunnel were also generated, with differences between traverses highlighted.

6.3 Conclusions

Unconstrained entry mobile robots provide the fewest restrictions and most options to robot configuration of any of the classes discussed in this thesis. As a result, configurations of this class have produced the largest, densest, and most accurate subterranean models to date, and will continue to do so into the future. Requiring portals for entry limits this configuration class to modeling objectives that include or are proximal to a portal. However, if this condition is met, no better method for providing accurate subterranean models exists.

The representative implementations outlined above, while perhaps not embodying optimal configurations for subterranean exploration and modeling, clearly demonstrate the challenges of the problem and two point solutions for successful unconstrained entry, unlimited mobility robot configurations. Groundhog and CaveCrawler accomplished many subterranean modeling task previously thought impossible, and demonstrated the feasibility, quality, and utility of robotic subterranean modeling when a portal is present.

Chapter 7

Constrained Entry, Zero Mobility

Many subterranean voids are only accessible through constrained entries such as boreholes in mines, manholes in sewers, or vents in underground facilities. Constrained entries disallow the use of the mobile and large static subterranean modeling robots discussed in the previous two sections, and require new and innovative approaches to acquiring subterranean range data. The simplest approach to miniaturization, and therefore apt starting point, is to configure robots with the minimum actuation necessary to acquire 3D data. Therefore, configurations which only provide scanner actuation and zero driving mobility are considered.

Just as with Groundhog, the original motivation for a zero mobility borehole deployed devices stemmed from the Quecreek accident. Groundhog was configured to take advantage of the entry into the Saxman mine provided by the breach. Another, lesser-known entry into Saxman had been established during the rescue effort through a pair of six inch diameter boreholes. These boreholes were drilled in order to observe mine water level and flow, but presented an ideal alternate method of gaining understanding about the conditions and extent of the Saxman mine (if sensors had been available).

The available boreholes, as shown in figure 7.1, provided the primary configuration constraints, specifically a 290 foot deep, sub-six-inch diameter borehole entering a mine with a seam height of roughly four feet. Therefore any successful robot configuration must compact into a small diameter package for deployment, support communications and deployment over hundreds of feet, and not require more than a few feet of void height for deployment. Given the lack of driving mobility, the ideal modeling conditions would be a scanner configuration with unlimited sensor range, a full sphere coverage pattern, and minuscule angular increments between adjacent data readings.

Utilization of borehole deployed static sensors requires development of new modes of operation in addition to device configurations. Mobile systems can only enter void through portals or through interconnections between voids. Hence mobile robots can't model non-

Figure 7.1: Boreholes into the Quecreek and Saxman mines: two 30" diameter rescue holes (left) and two 6" diameter communication holes (right).

contiguous voids. Borehole systems enter voids through boreholes, and new boreholes can be drilled where needed. Additionally, boreholes can be drilled on demand in many areas, often enabling entry into voids at positions of maximum interest or utility. While overlapping data sets taken from different holes can be aligned and correlated using the techniques outlines in section 3.4.2, even with ideal scanners distant boreholes combined with underground geometry will often result in separate, unconnected sets of model data.

While much field work is based on previously drilled boreholes at fixed and often non-optimal locations, sometimes subterranean modeling is conducted in cooperation with borehole drilling crews. When borehole placement is selectable, prior modeling data enables intelligent selection of borehole locations. Void model data can be analyzed to determine intersections between corridors, dead ends, and occluded areas. This analysis can result in borehole placements that maximize modeling data coverage. Therefore, in addition to mechanical, electrical and software configuration decisions, operational configurations must also be considered in designing an effective zero mobility constrained entry robot.

A series of devices conforming to the configuration philosophy outlined above was developed as part of this thesis work and their configuration and results from the field are discussed below. Development of devices led to the discovery of a wider array of applicable subterranean modeling tasks which in turn led to new configuration requirements. This cyclic process is traced through four generations of borehole deployed zero mobility sensors referred to as 'Ferrets'.

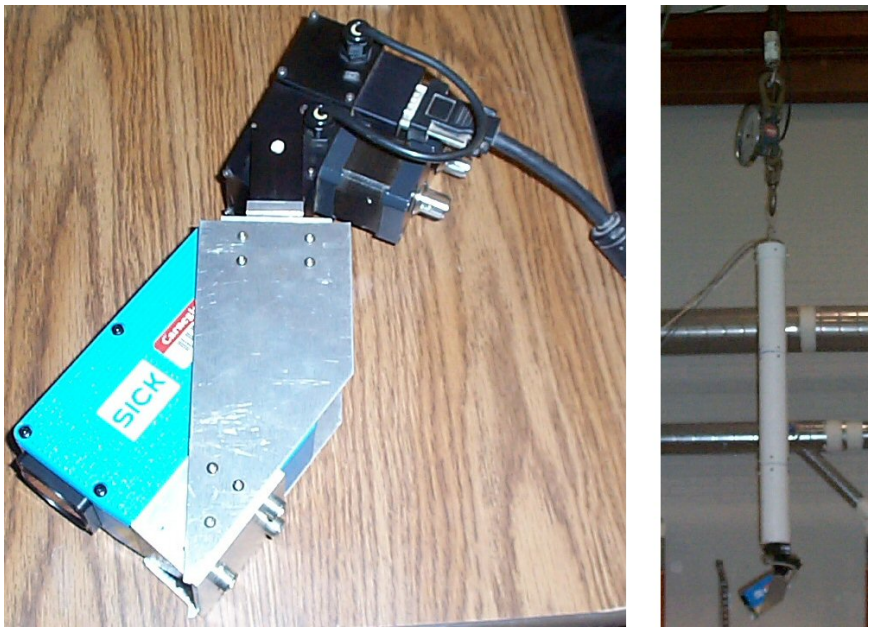

Figure 7.2: Core components of the original Ferret: Laser and pan/tilt unit (left), housing and deployment configuration (right).

7.1 Representative Implementation: Ferret

The original Ferret device achieved small size by using a small laser and a stacked component layout. The laser, a single point industrial laser range finder, was marginally effective but was hindered by limited range and low update rates. The laser was attached to a commercial pan and tilt stage based on small stepper motors. The laser was attached to the output of the tilt (or elevation) stage which was in turn attached to the output of the pan (or azimuth) stage, as shown in figure 7.2. This resulted in a cantilevered configuration which placed large torque requirements on the stepper motors. The cantilevered configuration also decreased scan accuracy by causing the center of mass of the entire device to shift during scanning. Electronics, computing and connections for deployment and communications tethers are housed and protected from the environment by a short segment of four inch PVC piping.

During deployment the laser is coaxial with the device and hole for compactness. After deploying into a void the laser can be actuated through 360 degrees of pan motion and 170 degrees of tilt motion. Scanning typically progresses by incrementing on angular increment in tilt, then sweeping through a complete 360 degrees in pan, then repeating this procedure with the pan rotation direction reversed until the entire volume is scanned. The stepper motors limit the angular resolution of the device to one degree in both axes. The stepper motors run open loop as space constraints precluding the addition of external angular

Figure 7.3: Graphical representation of figures of merit pertaining to the original Ferret: Scan time (left), coverage (center), entry diameter (right).

encoders. An initialization routine sets the zero points of each axis relative to the body of the device. Figure 7.3 graphically represents the coverage and data density possible with the original Ferret.

A microprocessor communicates with the laser and stepper motor controllers. Communications to the surface are routed over Ethernet, limiting the depth of deployment to the maximum Ethernet length segment of 100 meters. The 24VDC system bus is sourced at the surface and transmitted down a separate tether cable. Additional auxiliary sensing is limited to a single side facing infra-red proximity sensor used to detect when the body of the device has existed the bottom of the hole into the void. The on-board microprocessor packages and transmits data, while a laptop on the surface records and post processes the data, providing visualization and feedback about the void being scanned.

Ferret was first deployed into an underground void unexpectedly uncovered during the construction of a building. The extent of the void was unknown and no prior information about it existed. Ferret was deployed into the void and acquired the model shown in figure 7.4. This work validated the concept of zero mobility constrained entry configurations for subterranean modeling. While Ferret's high level configuration decisions were demonstrated to be effective, specifics of the mechanism and laser precluded the acquisition of highest quality subterranean void models and ultimately provided the impetus to create an improved version of the device, FerretII.

7.2 Representative Implementation: FerretII

FerretII's configuration was designed around improving the single weakest element in the original Ferret's design, the laser. The original small short range industrial laser was re-

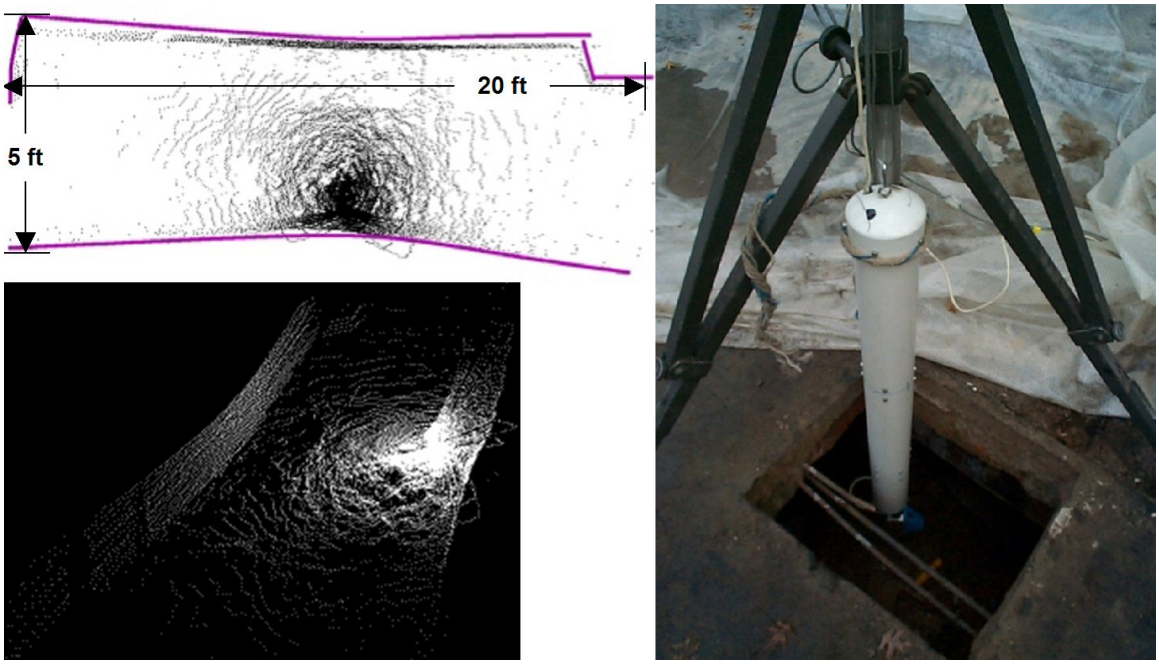

Figure 7.4: Ferret being deployed into an unknown underground void (right), resulting void map (top left) and void model (bottom left).

placed with a larger, much more powerful single point model originally intended for use in construction and surveying. While range was improved, the laser update rates were still low, requiring long scan times. Unfortunately, the improved laser's size drove the cross section of the entire device upwards, requiring a minimum of an eight inch diameter pipe for deployment. In addition to the laser, the scanning mechanism was also significantly improved. A comparison between figure 7.3 and 7.5 provides a summary of improvements in FerretII relative to the original Ferret.

Servo control of brushed DC motors outfitted with optical shaft encoders and planetary gearheads offers significant improvements in resolution, controllability and torque as compared to stepper motors. FerretII utilizes such motors in a configuration designed to minimize both required torques and disturbances during scanning. This is accomplished by rotating the large and heavy laser about its center of mass for both the pan and tilt axes thus eliminating the cantilevered laser mass found in the original Ferret's configurations. The final mechanical coupling of the pan stage is an internal gear. Mechanical support for the pan stage comes from a large bore thin section bearing. This combination of support and actuation of the pan stage allows power and communications wiring to be routed through the open center of the device, thus improving cable protection while decreasing cable movement. The tilt stage output is coupled through a chain drive to the centrally mounted pivot point attached to the housing surrounding the laser, as shown in figure 7.6.

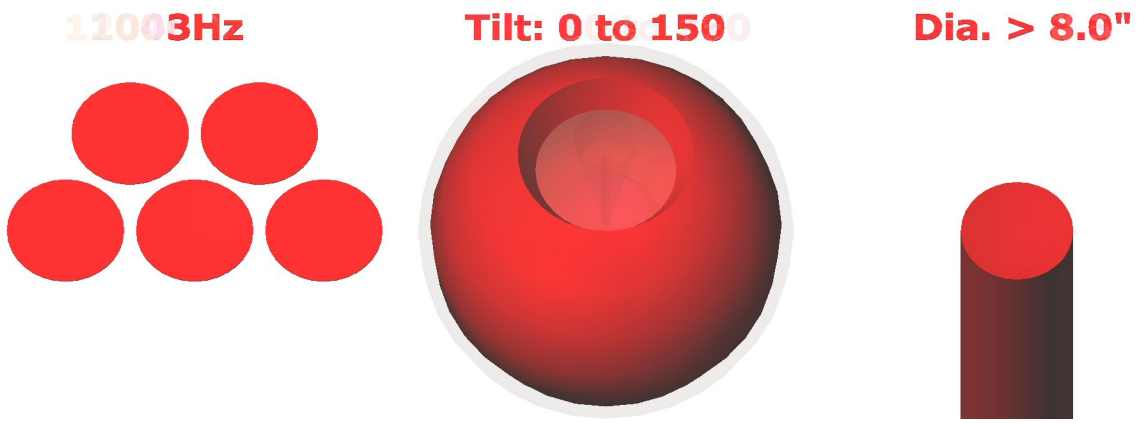

Figure 7.5: Graphical representation of figures of merit pertaining to FerretII. Note the improved range and coverage, but inferior access diameter and scan time.

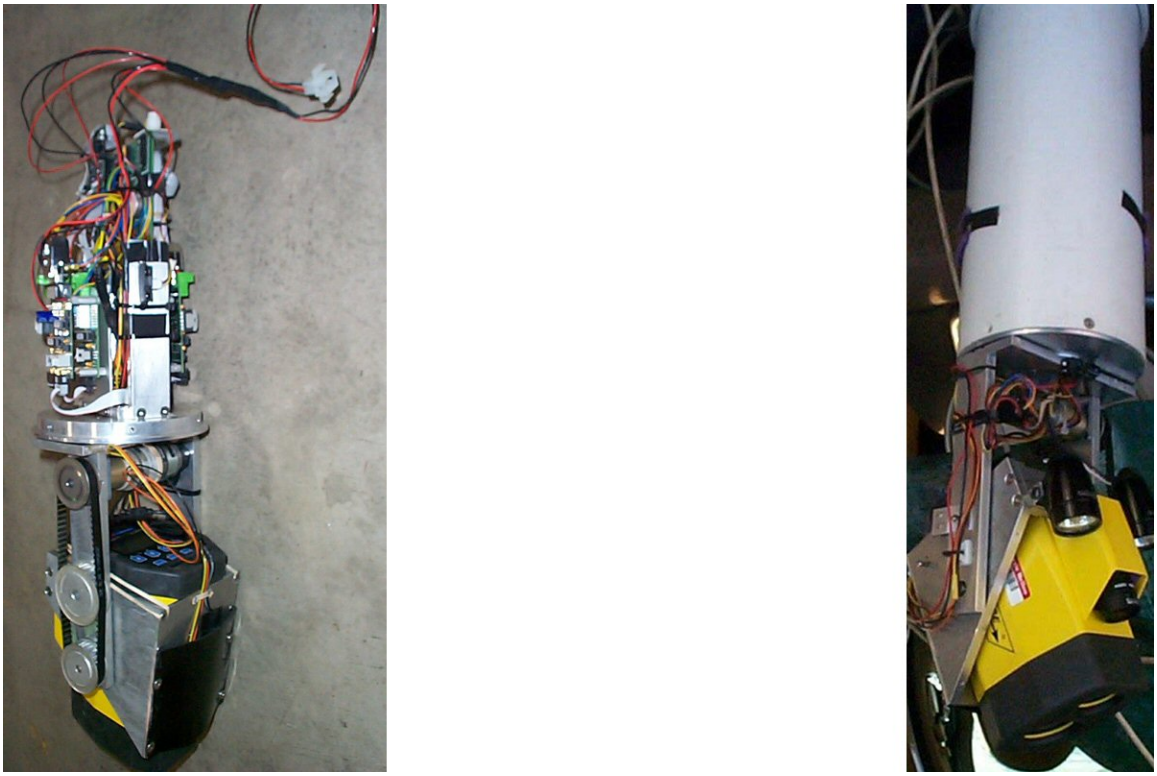

Figure 7.6: FerretII's configuration and primary mechanisms: unhoused (left) and ready for deployment (right).

FerretII also introduces improvements in processing and sensing as compared to the configuration of the original Ferret while retaining some configuration features such as tethered 24VDC power and Ethernet communications to the surface. An improved processor from the same family of microcontroller devices used in the original Ferret was integrated into the system, along with serially controlled closed loop servo controllers for the brushed motors. A low light video camera with infrared LED lighting was added for improved operator feedback, especially during deployments. Much of the configuration of borehole-deployed devices is driven by the diameter of the borehole and desired operations once the void is reached, but the borehole can often be the biggest obstacle to successful data acquisition.

Boreholes are idealized as continuous, uniform, clean, vertical cylinders connecting surface to subterranean void. In reality, the various strata of rock encountered during drilling tend to shift the drill bit from its vertical course. Additionally, unsound layers of rock are often encountered, resulting in fragments of rock partially constricting the hole or falling on top of void modeling devices as they are being lowered. Ground water also presents problems, with cascades of water sometimes entering the borehole at a specific depth. Even when ground water is avoided the difference in temperature and humidity between the usually warm (50 to 60 degrees Fahrenheit) and wet underground void and the surface can result in condensing humidity along the length of the borehole. Device optics are often fogged during deployment until thermal equilibrium between the device and the void is established. With water present, devices usually return from deployment coated in mud.

Some of the problems encountered in unlined boreholes can be ameliorated with casings. Casings are plastic or steel pipes that are driven into a borehole after drilling is complete. Casing provide a smoother inner bore for device deployment and limit the amount of loose material and water entering a hole. However, casings bring with them their own set of problems. Casings usually come in twenty-foot lengths, with flared interfaces between pipes. These interfaces result in distinct seams that can be problematic or even offsets between two casing sections given sufficient soil pressure. Soil stresses caused by the drilling can also deform plastic casings, but their lower cost has made them the most commonly used type of casing.

The combination of borehole and casing imperfections yields three critical observations relative to borehole device configuration. First, the effective hole diameter will always be less than the nominal diameter and devices should be designed accordingly. Secondly, devices should be designed to be robust to failing debris and to withstand forced insertions or extractions from deformed holes. Finally, observation of conditions in front of borehole devices during deployment can detect these problem conditions and enable judgments about hole feasibility to be made.

Many of these observations of borehole conditions were made during the field testing of

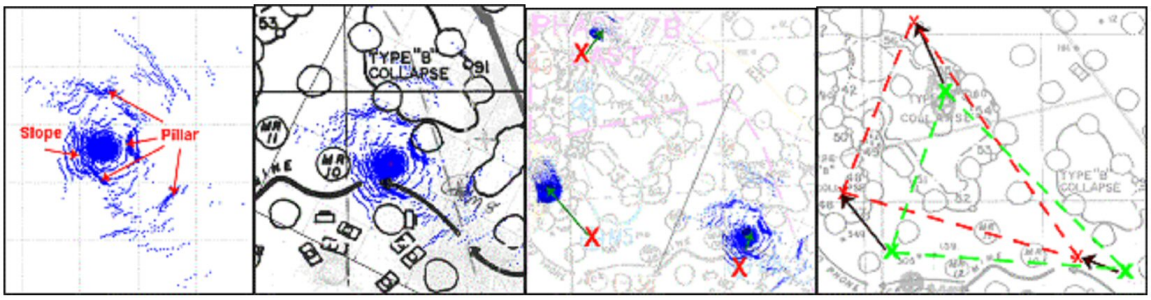

Figure 7.7: Prior map of a Kansas City limestone mine with FerretII acquired void model data overlaid.

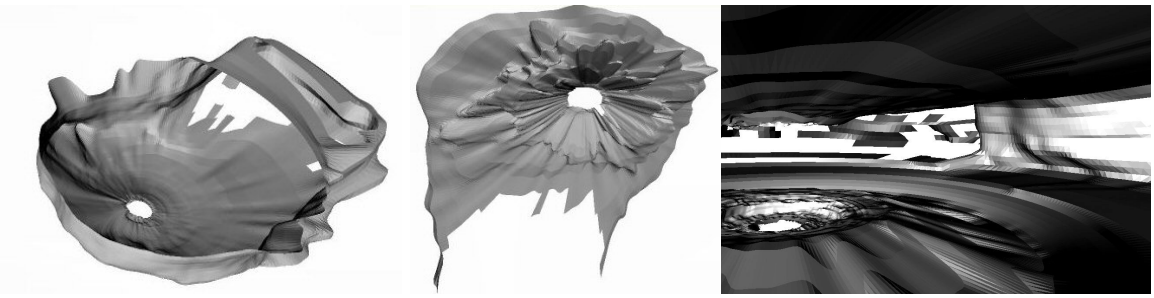

Figure 7.8: 3D shaded models of a Kansas City limestone mine acquired by FerretII from three different holes.

FerretII. FerretII's most extensive void mapping was related to the analysis of a partially backfilled limestone mine outside of Kansas City, as described in section 3.3.1. Deployment into three different holes resulting in three partially overlapping void models, as shown in figure 7.7. FerretII acquired X points per hole which were used to produce 2D maps, 3D point clouds and meshed polygon models as shown in figure 7.8. Another lesson about deployment via borehole-deployment can be seen from these models. The volcano-like cone directly beneath the borehole in each model is a result of cuttings from the drilling process falling into the void. This reduces the effective height of the void to less than nominal at the borehole.

FerretII has no onboard inertial sensing, therefore determining scan orientation in global coordinates is difficult or impossible. FerretII global position can be relatively accurately determine by surveying in the center of the borehole on the surface and measuring deployment depth relative to the top of the borehole. This provides model position information bounded by the accuracy of the borehole. Roll and pitch can be determined by the free hanging nature of FerretII. Gravity aligns the device and hence determines two axes of orientation. The final rotary axis, yaw, is the most difficult one to recover for FerretII. With the limestone mine in question yaw was recovered by rotating the individual hole models

about the yaw axis until the error between overlapping scans was minimized. The prior mine map was also useful for establishing initial estimates of yaw angle, although ultimately the surveyed data determined discrepancies in the mine map and its global referencing.

FerretII produced the first expansive, overlapping subterranean models acquired by borehole-deployed robot. The modeling conducted by FerretII was used to provide verification more safely and effectively than possible with humans and provided quantitative measurements about the extent to which the backfilling of the mine was successful, something not possible with borehole cameras or qualitative human observations. However, FerretII's large size required an eight inch or larger hole, mechanical configuration decisions made field servicing and repair difficult, and numerous fatal single point failures were unaddressed by the design. Therefore a new configuration compliant with six inch lined holes, featuring easier access for servicing, and designed to be robust to single point failures was deemed necessary.

7.3 Representative Implementation: FerretIII

FerretIII required a different laser in order to maintain the range of FerretII while providing access constraints compatibility comparable to the original Ferret. A large, long range, rapid update industrial single point laser was selected integrated into FerretIII. The general mechanical configuration of FerretII was reused in FerretII. The laser is balanced about a pivot point driven by a chain connected to a servo controlled DC motor. The pan stage utilizes large-bore thin section bearings and is driving by motor similar to that used in the tilt stage. One alteration in the mechanical configuration of FerretIII is in the location and mounting of the video camera. The size of FerretIIIs laser precluded mounting a camera against the laser. Instead, the camera is mounted to the same pan stage as the laser, but to an independent camera tilt stage. Therefore, pointing of the camera is independent of laser tilt angle. While the laser can not be tilted in the borehole due to its large size, the camera can be fully rotated, enabling improved visual observations of trouble spots within the borehole. This placement of the camera creates a small cone of occlusion directly below the device where the laser intersects the camera, as shown in figure 7.9.

One of the major changes in FerretIII's design is in the body of the device rather than the mechanical configuration. A square tube, while perhaps not intuitive given deployment into a round hole, greatly simplifies fabrication and allows the addition of numerous access panels (as shown in figure 7.10) for access to core electronics in the lab or the field. FerretIII's improved electrical system resides behind these access panels, with one panel covering a full fledged 600MHz PC104 form factor computer, another housing isolating DC/DC converters, one holding an IMU, and the final few enclosing batteries backups utilized in case of tether failure.

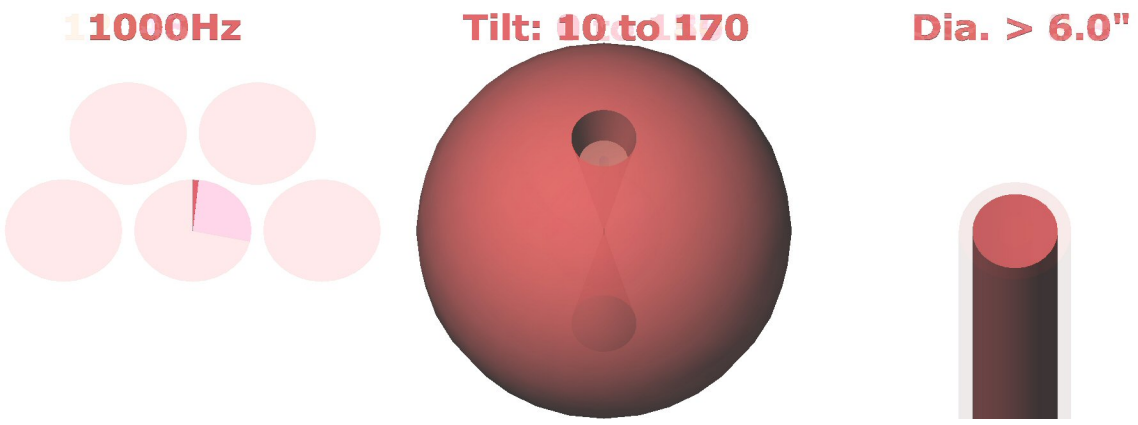

Figure 7.9: Graphical representation of figures of merit pertaining to FerretIII. Note the smaller required borehole size combined with wide range and rapid scan times.

FerretIII prevents robot loss through the use of redundant systems. Mechanical design results in the device potentially being in an unretrievable T configuration, as shown in figure 7.10. Therefore, single point failures could strand the robot in this fatal configuration. Mechanically, dual chain couplings between the tilt motor and laser mounted sprocket provide mechanical redundancy. Electrically, an on-board 12VDC sealed lead acid battery stack is charged by power delivered over the tether at a nominal 14VDC. Should the tether power lines or surface power supply fail in any way the device and all its function remain powered by the batteries. Similarly, the on-board computer monitors both the power and communications connections to the surface. If power or communications are lost, the downhole computer initiates a safing procedure which returns the device to a deployment configuration for safe retrieval. Finally, during early field testing failures with the computer were observed, resulting in the device freezing in an unextractable configuration. Four spare lines were rerouted to the surface to provide direct connections to the terminals of the tilt motor. A double throw double pole relay switches the motor terminals between output signals coming from the amplifier (nominal operations) and two spare lines (emergency retrieval). Triggering the relay applies power directly to the tilt motor, providing hard line tilt angle control from the surface independent of computer, communications, and servo controller functionality. These redundant systems enable FerretIII to survive and be retrieved from any single fault failure that should occur down the borehole.

Part of the configuration of a borehole-deployed device is the system used for lowering into the borehole. A steel cable running over a pulley to a manually controlled winch actuates device depth. Data and power tethers and a tape measures bound to the steel cable to provide depth measurements and communications with the device. Simple hand management of tethers was effective in the field, however room for improvement exists.

Figure 7.10: FerretIII's deployment (left) and unretrievable T (right) configurations

FerretIII uses a large tripod for easy access to the device immediately prior to deployment, as shown in figure 7.11. Additionally, for shallow holes less than 100 feet in depth FerretIII utilizes an alternate method of deployment. Sixteen foot sections of nesting, perforated square steel box beams provide an alternate means of mechanically supporting the device, yield an alternate measure of hole depth not influenced by cable flex, and most importantly project underground device orientation to surface. As the device is lowered, new sections of box beam are attached together with bolts. Careful alignment of marks on the ends of each tube result in both orientation and depth being preserved. When the device is at depth, a cross member is bolted to one side of the tube to provide both structural support and an exaggerated visual reference of underground orientation. Either the box beams or the steel cable can be used individually for deployment and retrieval, with both used in tandem being the most common configuration.

Beyond the improvements to the device and its deployment system, device control and model post processing algorithms were also improved. Figure 7.12 shows the Graphical User Interface (GUI) used by the operator during scanning and a model manipulation tool useful for analyzing data once a scan is complete. The GUI provides an operator the ability to change angular increments, scan sectors, and video camera behavior. The

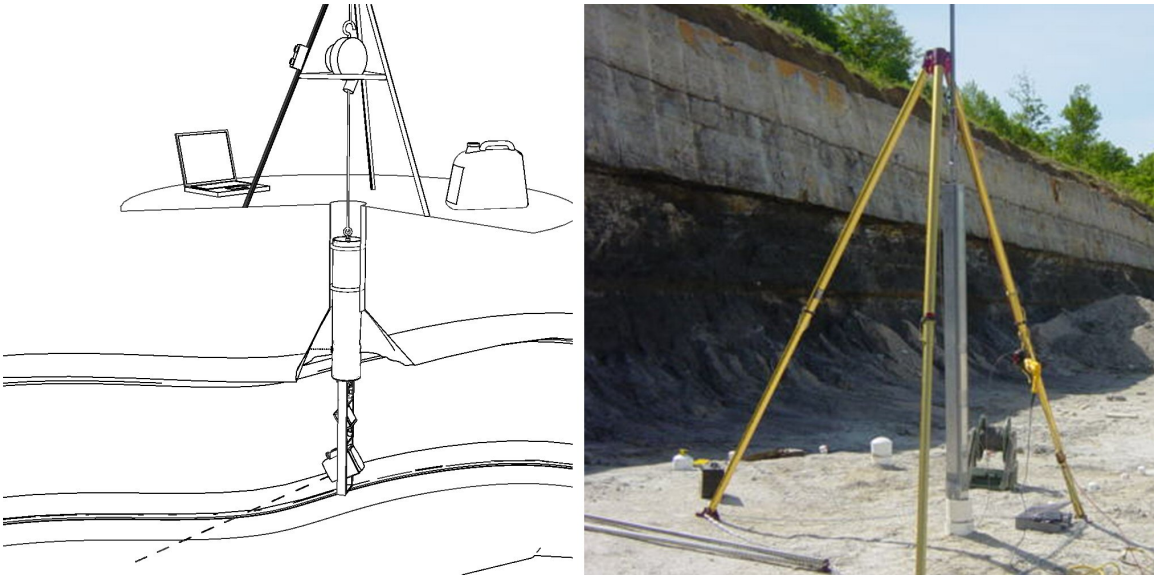

Figure 7.11: Deployment system concept (left) and implementation for FerretIII (right).

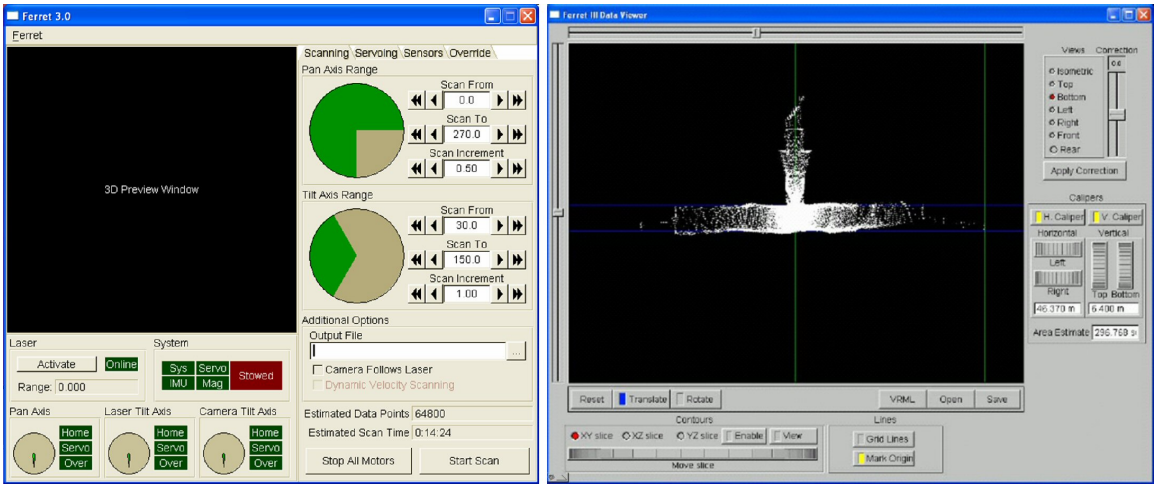

Figure 7.12: FerretIII's user interface (left) and model viewing and manipulation tool (right)

data viewer provides tools for extracting 2D cross sections from 3D data as well as the ability to take rapid measurements of models in the field. GUI design, while not critical for successful subterranean modeling, enables more efficient data collection and thus improves the cost performance of robotic subterranean modeling versus other methods and is therefore mentioned here.

FerretIII's GUI, fail safe systems and overall configuration have been tested by more hours of field work than any of the other static sensors described in this work. FerretIII acquired a model from a single borehole in a coal mine as part of its initial configuration and

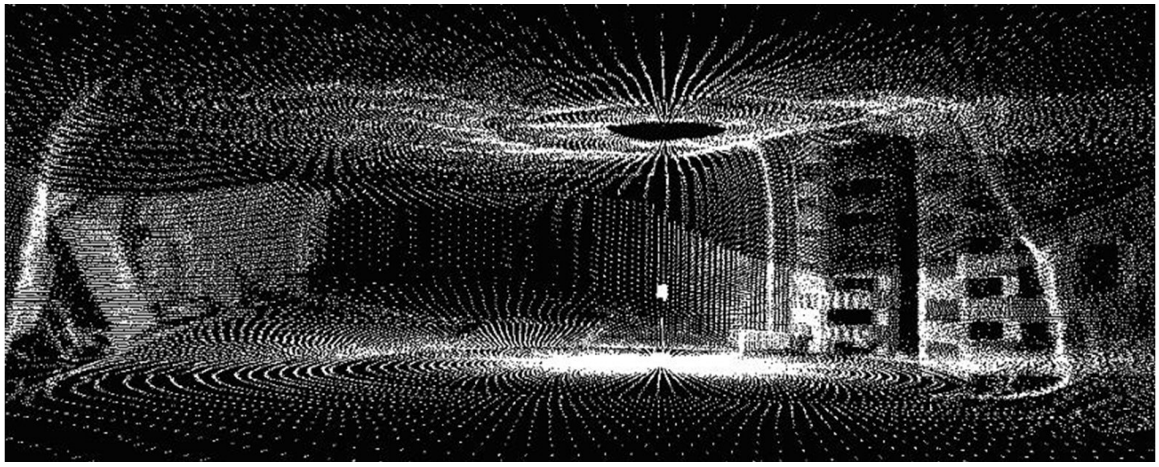

Figure 7.13: Coal mine void model obtained by FerretIII (note cribbing stacks).

functionality testing (see Figure 7.13). FerretIII has been deployed and used extensively in coal mines in West Virginia, collapsed limestone tunnels in Indiana and in industrial basements.

FerretIII was used to validate use of resistivity as a geophysical technique for detecting abandoned mines at Lots Branch in West Virginia. Accurate void models were required in order to validate and verify the assumptions and results of the geophysical survey. FerretIII acquired multiple subterranean void models and provided the data necessary to substantiate the geophysical technique. The shallow thirty feet deep holes combined with the use of box beams during deployment enabled accurate device alignments without the use of inertial measurement systems. This field mission also provides an anecdotal but representative look at the interaction between boreholes and successful void scans.

Figure 7.14 shows a series of boreholes drilled to determine void location and the modeling data acquired from successful boreholes. Of the over thirty labeled boreholes, only six were deemed acceptable for void mapping, primarily due to boreholes striking coal in the mine pillars rather than void spaces. Of the six potential holes, only four had sufficiently intact and undeformed boreholes for deployment. Of the four boreholes only one resulted in a void model of an intersection rather than of a corridor segment. While all models are of interest, models at intersections provide more data from a single hole and improve correlation between holes and with existing maps. If sensing and drilling operations had been coupled, then data obtained from FerretIII could have resulted in better borehole placement, but unfortunately drilling and void modeling operations were conducted separately.

Figure 7.15 shows some of the details and different views obtainable from referenced 3D model data. Numerical integration of point cloud models leads to estimates of void volumes. Cross sections can be aligned to the direction of an underground corridor, or

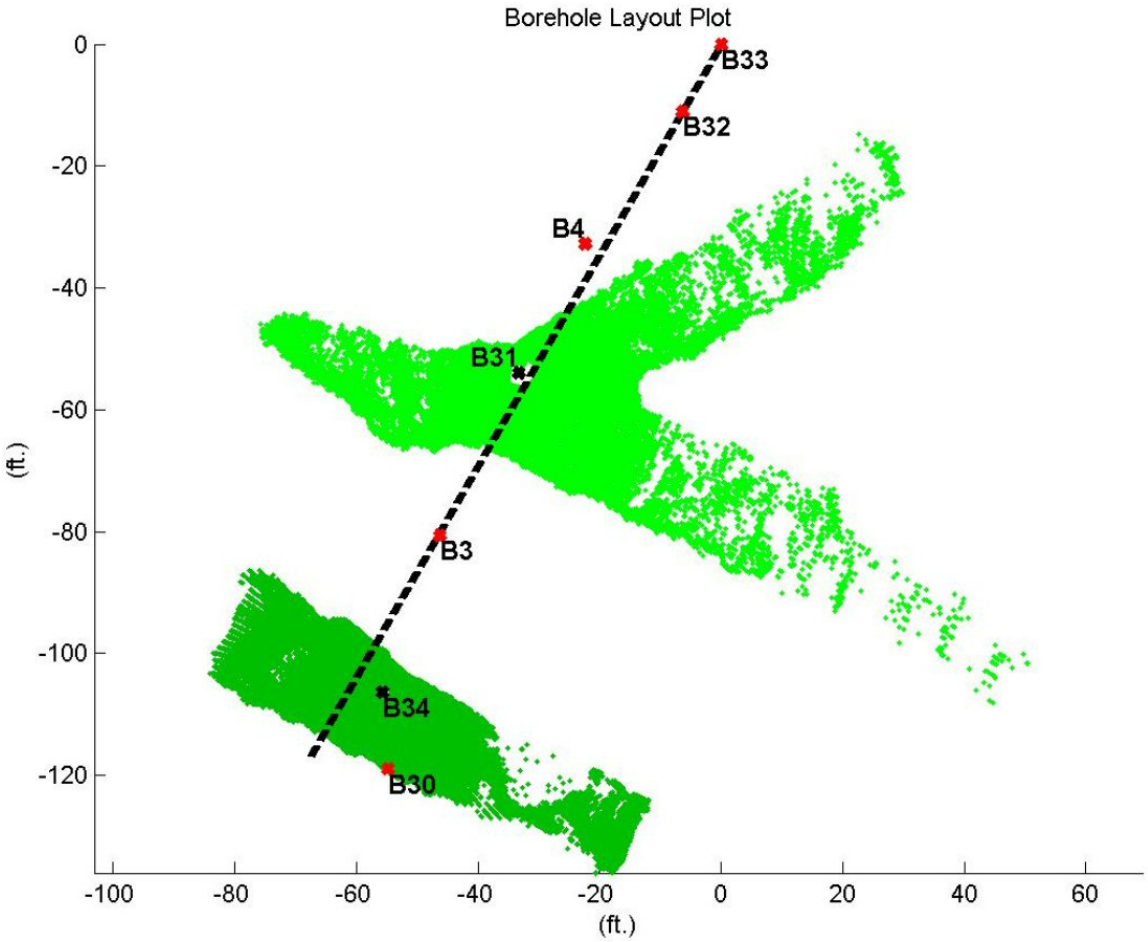

Figure 7.14: Boreholes and voids at the Lots Branch mine (note only 2 of 7 boreholes were acceptable for deployment).

to cardinal direction. Perspective views provide a better understanding of void structure. The strongest benefit from 3D void data, the ability to flythrough data and manipulate models in real-time in three dimensions, is impossible to convey within the limitations of this document, but is shown to be of high utility to geologists and mining engineers in the field.

In many respects, FerretIII represents an endpoint in the evolution of the configuration it embodies. At the time of writing, commercially available laser range finders with comparable range, update rate, and environmental robustness do not exist in smaller packages than the one integrated into FerretIII. Therefore a similar configuration that is compatible with smaller diameter boreholes is not possible. While minor incremental improvements in range, resolution or accuracy might be possible, they represent upgrades to the existing system rather than new configurations. However, by removing the laser from the down-hole

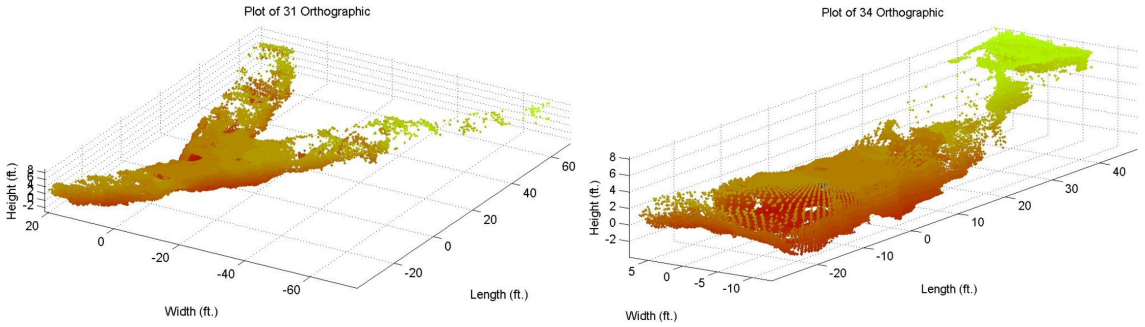

Figure 7.15: Details from some borehole scans at Lots Branch: Borehole 31 (left) and Borehole 34 (right).

portion of the device revolutionary rather than evolutionary configurations become possible, as described in the next section.

7.4 Representative Implementation: CoreHoleFerret

CoreHoleFerret's design enables superior void data acquisition through an improved concept of operations. CoreHoleFerret deploys down the inside of industry-standard hollow coring bits. Once the last core section has been extracted, a hole runs down the center of the drill steel from the surface to the void. Deployment down a coring bit utilizes the drill steel as a borehole liner, eliminating the need for installing costly casings. Deployment immediately following drilling through hole lined with drill steel (rather than plastic) eliminates the possibility that soil mechanics will pinch or deform the borehole into an unusable geometry. CoreHoleFerret also utilizes an elephant-trunk-like tilting motion, thus minimizing clear void height required for deployment. Cheaper, more consistent deployments combined with reduced void height requirements results in increased void data acquisition effectiveness.

CoreHoleFerret is deployable in a 2.5" diameter hole. This hole diameter is common to an HQ coring rig which is an industry designated bit size which generates 3.75" holes and 2.5" diameter core samples. Unlike conventional rotary hammer drilling which breaks up and removes all of the material in a hole, coring removes a small annulus of material through drilling while the central core is removed intact. Because of the smaller diameter holes, coring rigs are much smaller, quieter, and more maneuverable than rotary hammer rigs. Therefore coring rigs can place boreholes in locations not possible with rotary hammer rigs. There smaller size, quieter operations, and limited build up of cuttings around the hole also make coring rigs ideal for operations in residential or other inhabited regions.

Coring rigs also offer a unique method of deploying into boreholes. After the coring rig breaches an underground void the last core sample and sample retrieval mechanism can be removed from the drill stem. The wire line hoist running down the now hollow center of

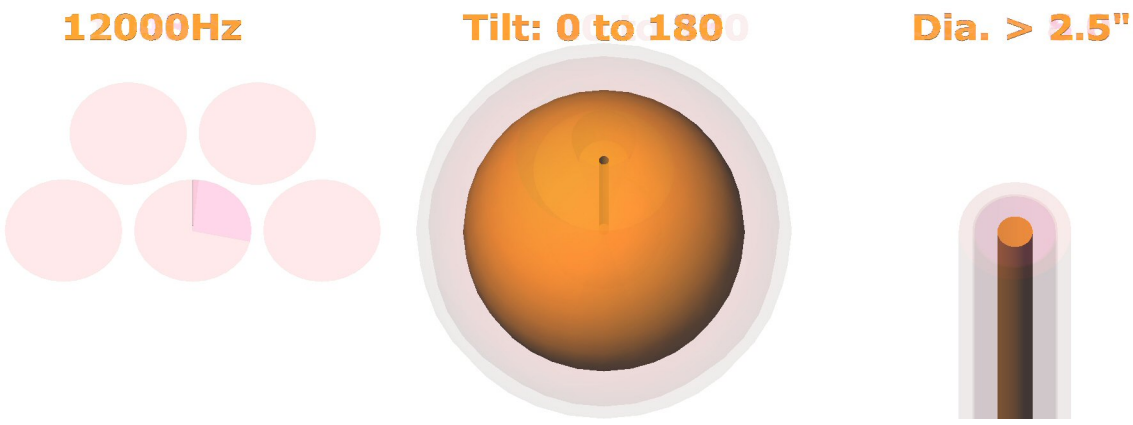

Figure 7.16: Graphical representation of figures of merit pertaining to CoreHoleFerret. Note the small borehole size, more complete coverage and faster scan times relative to previous generations of Ferret.

the drill stem can then be used to lower CoreHoleFerret. CoreHoleFerret is configured to be compatible with deployment down the interior of the drill stem, with an outer diameter of just 2-7/16". This small size also allows deployment down cased or uncased holes after the HQ drill stem is removed.

As with the first three generations of Ferret devices, the laser range finder selected determines much of the rest of the configuration. CoreHoleFerret employs a fiber coupled laser range finder. This allows the large, expensive and delicate optics, electronics and processing associated with the range finder to remain on the surface while a pair of fiber optic cables transmit and receive laser pulses between the surface and the device underground. The device deployed into the subterranean void requires passive collimating and focusing optics compatible with fiber optic terminations, but does not need any active range finding components. Removing the largest and most expensive component of the system from the downhole portion of the device allows radically different configurations and reduced borehole diameters. The commercial fiber coupled laser system utilized, combined with the mechanical design detailed below, results in the coverage and resolution shown in figure 7.16.

While the fiber optic coupling of the laser range finder makes CoreHoleFerret's configuration possible, it also requires the scanning mechanism to prevent sharp bends in the fiber. Fiber optic lines have minimum bend radii, as excessively tight curvatures result in loss of signal or mechanical failure of the glass fiber. The centrally pivoted tilt joint used in prior Ferret configurations is not configured to maintain a minimum bend radius, requiring either a large and flexible service loop or a pass through along the axis of pivot. Both options would result in the minimum bend radius of the fiber being violated. Therefore,

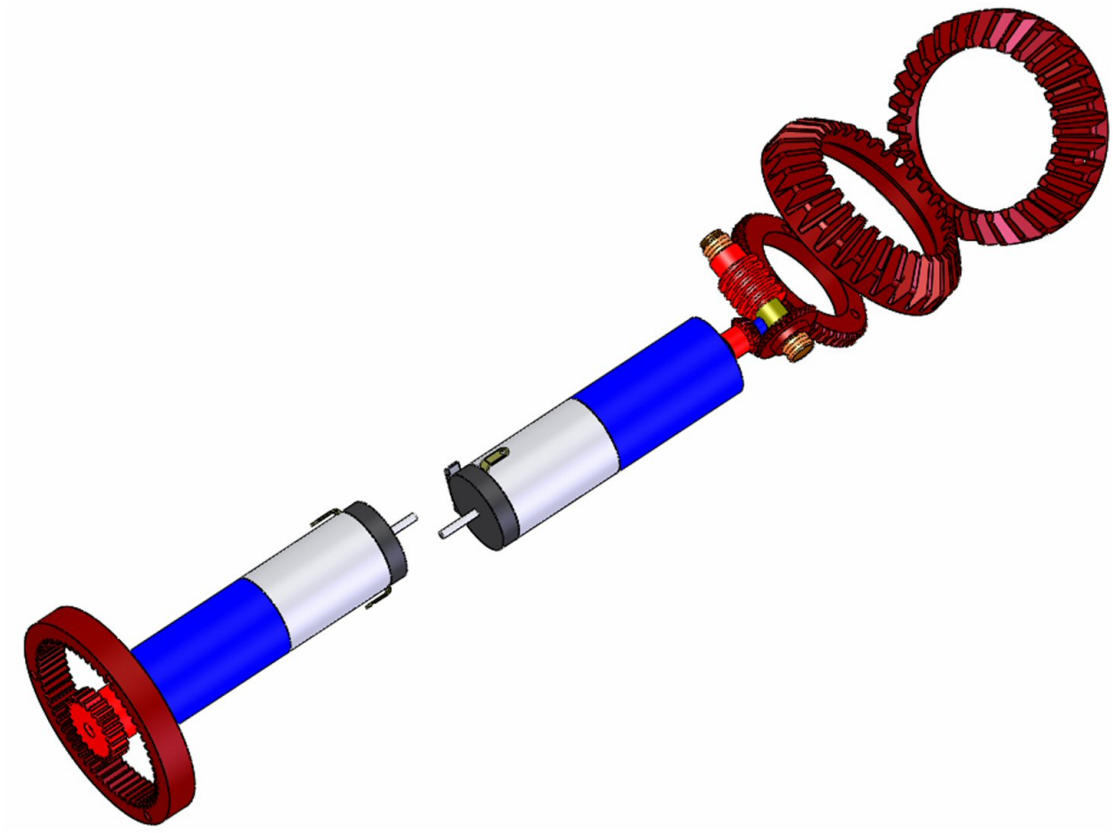

Figure 7.17: CoreHoleFerret gear trains. Note the miter gear closest to the motors is fixed relative to the motors.

CoreHoleFerret's tilt joint was redesigned to provide maximum possible bend radius while also improving device sealing and conforming to the small device diameter.

The mechanical transmissions employed by CoreHoleFerret are shown in figure 7.17. The pan axis is the same configuration used in the prior two generations of Ferret, but with the addition of a dynamic rotary lip seal to water and dust proof the joint. The tilt axis is based on a mechanical coupling through a pair of miter gears. The tilt gearmotor drives a worm through 2:1 bevel gear reduction stage. The worm rotates a worm gear mounted at forty five degrees to the primary axis of the device. The worm gear is coupled to the middle joint. If the lasers optics were mounted to the middle joint they would sweep out a ninety degree cone as the worm gear rotated, which when coupled with the pan axis would result in hemispherical coverage below the device. Instead, the optics are mounted to a third segment with a miter gear rigidly attached. A mating miter gear is rigidly mounted to the primary housing of the device. As the middle joint is actuated the miter gears force the optics joint to rotate relative to the middle joint at the same rate of rotation as that between the middle joint and the base joint, as shown in figure 7.18. This results in the tilt

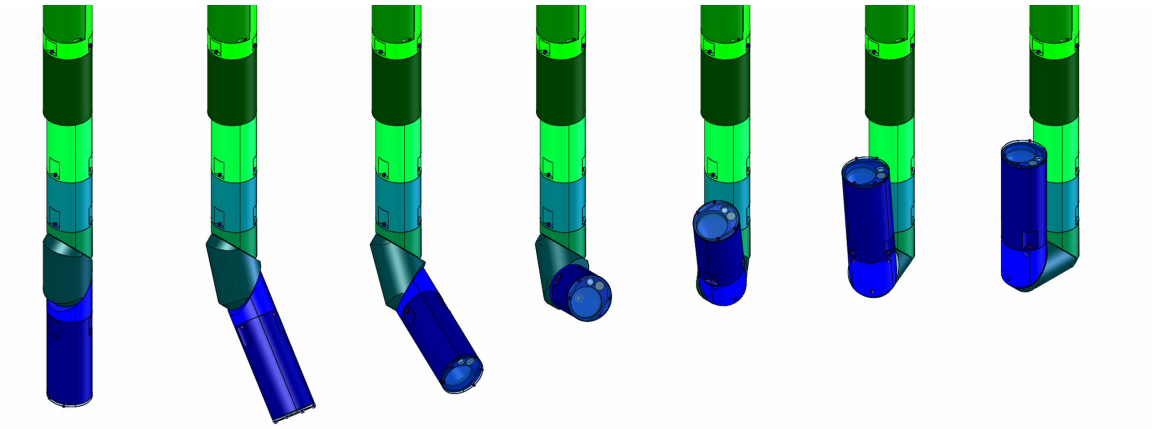

Figure 7.18: Series of images showing actuation of the elephant-trunk-like tilt axis (note the pan axis is static throughout).

axis going through a full 180 degrees motion, resulting in nearly full spherical coverage.

The configuration shown results in a fully sealed device which maintains a large bend radius throughout its full range of motion, as shown in the section views of figure 7.19. While the gear train is as simple as possible given the angles and orientations over which power must be transmitted, the cumulative gear backlash is significant. Therefore angular position is encoded directly at the joint in addition to the motor mounted encoders. While the motor mounted encoders are standard incremental optical encoders generating quadrature signals, the joint encoders are virtual absolute encoders. Virtual absolute encoders use a pseudo random serially encoded index tracks to provide absolute positioning information from any starting location after rotating through a few degrees. Encoding the index track means that the encoder disk and read head are no larger than a comparable resolution incremental encoder, a critical feature for the tight volumetric constraints of 3"Ferret. Absolute encoding also eliminates error prone and time consuming homing procedures used by previous generations of Ferret, which were only outfitted with incremental encoders.

In addition to improved angular encoding, CoreHoleFerret of necessity has superior orientation sensing to prior configurations. CoreHoleFerret's small size and large interior volumes for components results in a very light weight device. The cantilevered mass of the optics and middle joint through much of the tilt joint range results in an off axis device center of mass. Rotation of the pan axis results in device wobble, similar to that observed in the original Ferret configuration. This problem can be corrected by rigidly bracing the device against the end of a borehole or through better sensing of device body orientation. Bracing against the end of a borehole is highly dependent on borehole geometry and not feasible in tall voids where the device needs to be tens of feet away from the end of a borehole. Therefore improved sensing was selected as a means to record and eliminate

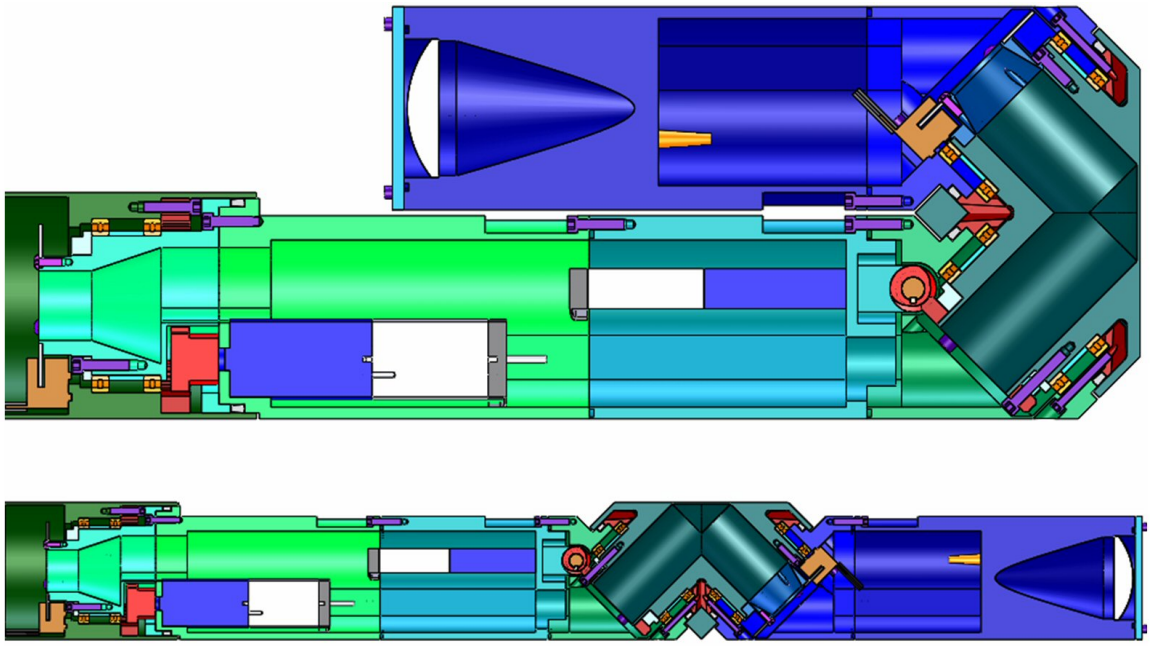

Figure 7.19: Section views showing 3"Ferret at 180 degrees tilt (top) and 0 degrees tilt (bottom).

wobble based inaccuracies.

CoreHoleFerret senses orientation through a six axis MEMS inertial system and through an analog liquid level sensor. The inertial system is part of an integrated device which combines magnetometer data with gyro data through a filter to produce smooth, disturbance rejecting yaw angle measurements. Actual changes in device orientation will be recorded by both magnetometer and gyro, while disturbances in only one of the mechanical or magnetic domains will be rejected by the filter. Additionally, the fixed nature of the Earth's magnetic field can be used to correct gyro drift. The analog tilt sensor takes capacitive measurements of a conventional level bubble, providing an alternate source for roll and pitch information with different error characteristics than those seen in MEMS IMUs. These two systems combined provide the data necessary to account for angular offsets of the body of CoreHoleFerret.

Almost all the various angular and orientation sensors onboard CoreHoleFerret are connected to a single microprocessor (as shown in figure 7.20) that collects and packages the data for transmission. While very small form factor computers exist, the required analog, digital, and serial interfaces for communicating CoreHoleFerret's sensors are more commonly found on microprocessors. Additionally, CoreHoleFerret utilizes a custom printed circuit board for power distribution, sensor interfacing and ease of wiring, so the additional electrical integration effort associated with adding a microprocessor is minimal. Finally,

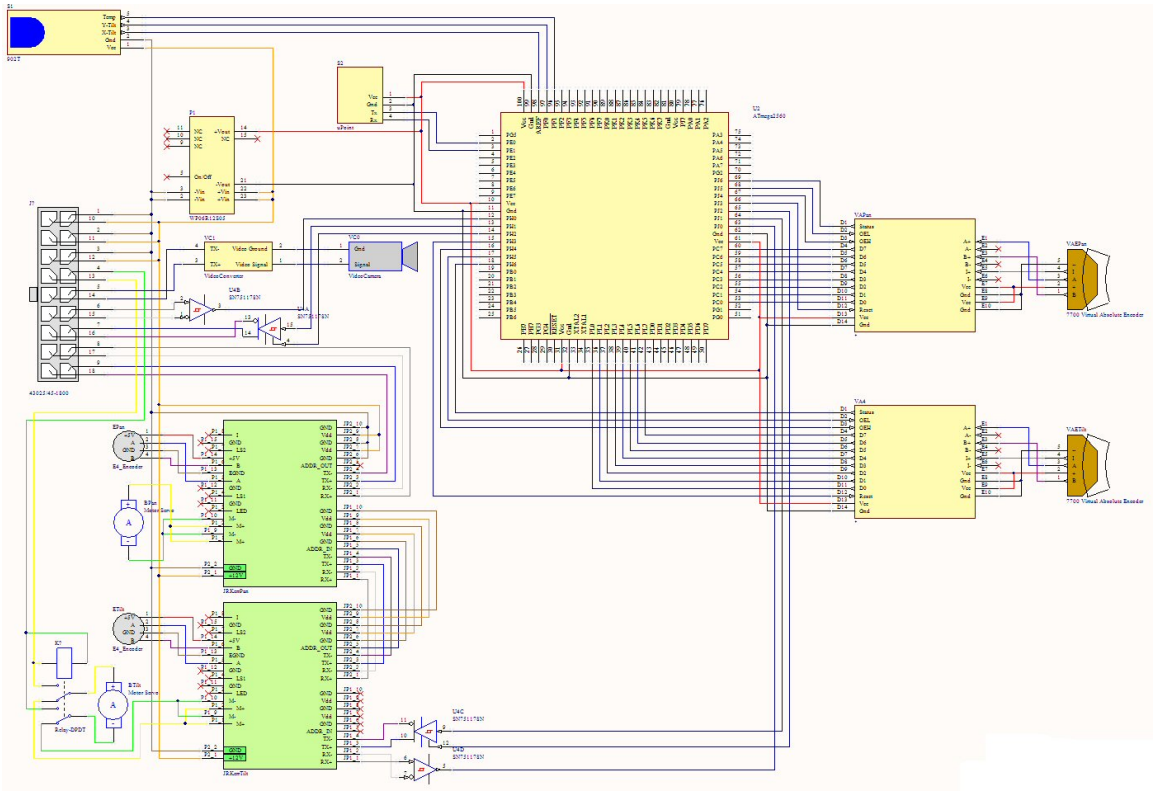

Figure 7.20: Schematics of the CoreHoleFerret down-hole system. From left to right, motor controllers, down-hole microprocessor, virtual absolute encoders.

the downhole microprocessor serves to extend the fault tolerance concepts introduced in FerretIII to include redundant communications lines and operational procedures.

CoreHoleFerret's two motors are controlled by servo amplifiers connected to a shared RS485 serial bus, referred to as the motor bus. The servo controllers are only connected to the motor mounted encoders and the motors themselves. Commands on the motor bus originate from a control computer on the surface during nominal operations. However, the downhole microprocessor is also connected to the motor bus. During nominal operations it only listens to motor command traffic. If the microprocessor fails to observe communications on the motor bus for five seconds, it assumes that something has failed in the serial bus lines or in the surface control computer. The microprocessor responds by commanding the motor controllers to return to preset home positions corresponding to the retrieval configuration of the device.

The microprocessor also has its own RS422 connection to the surface control computer. During nominal operations, this connection is used to communicate supplementary data such as device orientation and absolute joint angles. Having the microprocessor connected to two independent serial buses provides redundant communications and enables rerouting

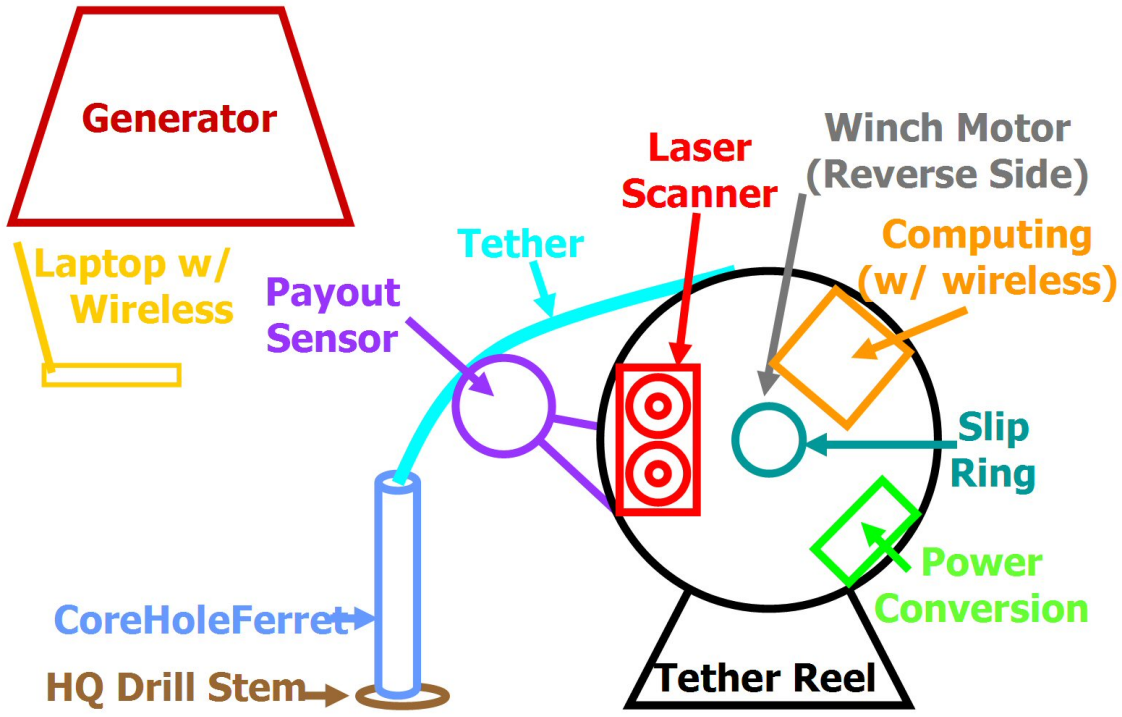

Figure 7.21: Conceptual view of complete CoreHoleFerret system.

of data and control should either communication channel fail. In addition to redundant communications, power is supplied on three pairs of conductors, allowing any one pair to fail without compromising current carrying capacity and any two pairs to fail without cutting all power to the device. Additional spare lines are also routed to provide relay controlled direct connections to the terminals of the tilt motor, as in FerretIII.

CoreHoleFerret's design necessitates a more advanced and integrated deployment system than used in prior configurations. Hand management and storage of electrical and communications tethers is feasible if the tether is robust. For CoreHoleFerret the delicate but necessary fiber optic would likely be damaged if not shielded from the environment and clumsy device operators. CoreHoleFerret employs a hollow strength member as a protective jacket for the fiber and for lifting and lowering the downhole portion of the device. The fiber optic, along with a multi-conductor electrical cable for signal and power are housed inside a hydraulic hose. A tether reel assists in storing and managing the hydraulic hose. The tether reel also houses the laser range finder electronics, control computer, power conversion, and wireless communications inside of the rotor. This eliminates the need for a fiber optic slip ring and provides environmental protection of these delicate components.

The tether reel is designed to be modular and easily broken down for transport. In addition to the scanner head and tether rotor, the complete system also contains a tether

stator, a generator and a laptop, as shown in figure 7.21. The tether stator provides bearings and a mounting interface for supporting and attaching the tether reel rotor. The reel also eases the burden of hauling the entire system to remote field locations by acting as a cart for the heaviest components. A servo motor equipped with a planetary gearhead spools and unspools the tether through a non-backdrivable worm gear transmission. The worm shaft also has an external drive head which can be actuated by a hand or drill should the servo motor fail. An industrial amplifier with integrated servo controller is also integrated into the stator. The rotor computer controls tether payout through the stator servo controller and with the aid of an external friction wheel based payout encoder.

The entire CoreHoleFerret system is powered from a single 120VAC source, coming either from a portable gasoline generator or from a car battery via an inverter. Internal wiring inside the tether stator and rotor handle power distribution and voltage level conversions. The operator laptop is also powered from this 120VAC source and runs a graphical user interface which controls CoreHoleFerret and data visualization software. CoreHoleFerret's software performs many of the same functions found in previous generations, but with a more robust underlying software architecture. Additionally and most importantly, 3"Ferret's software provides model data updates in real-time in addition to reusing post-processing and data manipulation tools developed for analyzing complete scans developed for prior Ferret generations.

CoreHoleFerret is a design, not a working device. Therefore, CoreHoleFerret does not have data results from either the lab or the field. CoreHoleFerret will be the smallest device of its class capable of collecting subterranean modeling data, although this assertion has yet to be substantiated. The revolutionary nature of its design enables more efficient data collection through tighter coupling between drilling and scanning. CoreHoleFerret can be deployed down more holes and into shorter voids than prior, larger diameter, greater swept volume systems. Therefore, CoreHoleFerret will supersede prior systems as the constrained entry, zero mobility configuration of choice.

Chapter 8

Constrained Entry, Limited Mobility

Constrained entry limited mobility configurations represent the best of mobility subject to the sever constraints of borehole deployment. These configurations constitute unprecedented modeling capabilities in the design space of subterranean modeling robots. While these configurations require larger boreholes then zero mobility configurations and have inferior range and mobility to unconstrained entry configurations, they can model underground voids not viewable by either of these other configuration classes. Despite the challenges inherent in developing configuration of this class, the middle ground of the configuration space fills a critical gap in the capabilities required to successfully and efficiently model subterranean voids, and therefore should be pursued.

Configurations which can deploy down boreholes yet travel hundreds of feet along subterranean corridors require robots that are self reconfigurable. While some configurations that look and behave the same during and after deployment are possible, the limitations of these classes of devices, such as limited sensor payload capacity, poor sensor vantage points, and insufficient mobility (see section 2.4.2) make them unsuitable for subterranean modeling. Therefore configurations that have distinct deployment and operational configurations are examined.

After discussing concepts for configurations in this class, this section explores a specific configuration that was taken from concept to prototype to system. Finally, some reflections on the capabilities, limitations, and reasons for exploring this challenging region of the design space are recorded.

8.1 Possible configurations

While the problem of developing robot configurations capable of meaningful locomotion after deployment down small diameter boreholes is challenging, the span of potential configurations is large. The one requirement that is common to all subterranean modeling robots is the need for 3D scanner coverage, which requires both the actuation to provide 3D coverage and scanner placement which keeps the range finder clean and unobscured. Mobile configurations must provide a means for driving over rough terrain and obstacles and the ability to turn and maneuver in constrained environments. Driving and steering capabilities on a constrained platform will never equal those possible on an unconstrained platform, but as the challenges of mines and caves don't vary with means of access, capabilities must at least be on the same order of magnitude for borehole and portal entry machines. Likewise, while a borehole-deployed mobile platform will never reach the level of miniaturization possible with a static sensor, the diameter constraints are of the same order of magnitude, thus requiring mobile borehole systems to compact into small diameter compatible configurations. The unique aspect of this class of device is the requirement for transformation between the two configurations. In order to postulate and evaluate mobile borehole robot configuration, these five elements, 3D scanner configuration, driving ability, steering performance, borehole compatibility, and self reconfiguration must be examined. Four concepts and how they satisfy these various requirements are outlined below.

8.1.1 Two Wheel Inflatable Differential Drive

Figure 8.1 shows a concept for borehole-deployed mobility that illustrates considerations for designing borehole explorers. The system concept, referred to as Cyclops due to its central, single laser scanner configuration, is built around gravity reaction locomotion using only two actuators. Large inflatable wheels provide ground clearance, compatibility and compliance with rough terrain, and a means of drastically but simply transforming the robots configuration after accessing a void. Driving is accomplished by rotating both wheels in the same direction. Some differentially driven robots rely on a reaction tail to provide the force closure necessary for locomotion. The mechanism required to deploy such a reaction tail would be challenging to design and actuate. Additionally, reaction tails result in significant system drag, reducing effective system range, and have the potential to snag or jam when executing turns in debris strewn environments. There, Cyclops relies instead on the offset mass of the central body to provide counteracting torques to those generated by the wheel motors. This results in the geometry of the robot's configuration limiting the maximum traversable incline or obstacle. Cyclops steers differentially, and can therefore execute an arc of any curvature, including a point turn.

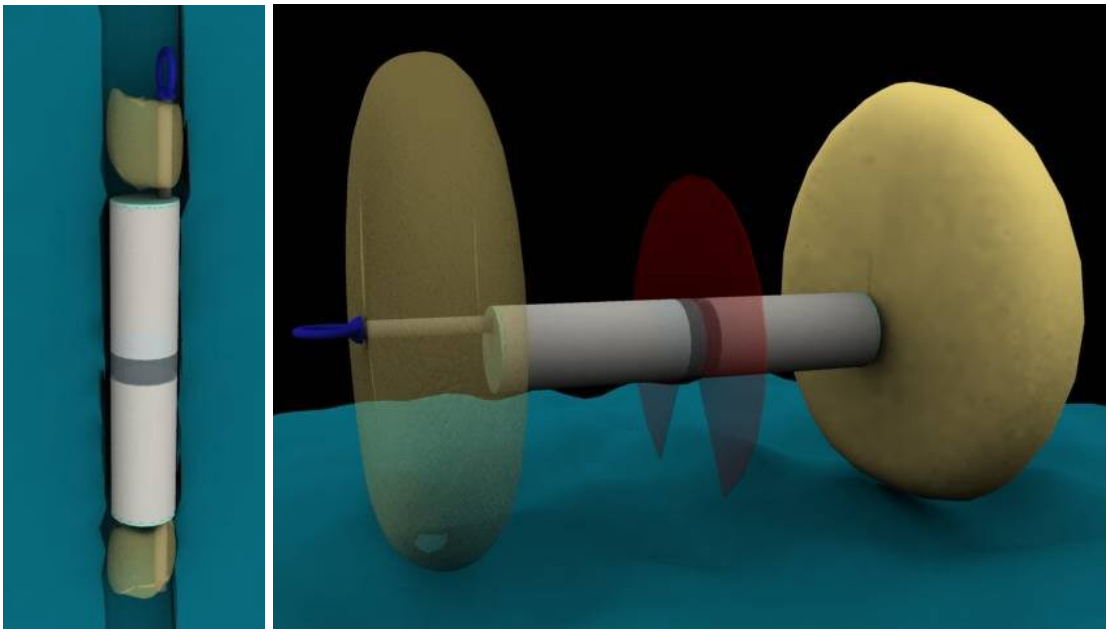

Figure 8.1: Cyclops constrained entry, limited mobility configuration concept during deployment (left) and during modeling (right)

Cyclops also has one of the simplest (in terms of actuators) 3D data acquisitions schemes possible. A centrally mounted laser scans a plan normal to the axis of the wheels with the aid of a rotating 45 degree mirror. This allows a standard commercially packaged single point laser range finder to be used, as such devices usually have their largest dimension along the axis of the laser. 3D coverage is achieved through integrated robot motions. For example a 'static' scan can be taken by executing a point turn, resulting in full spherical coverage. The envisioned nominal robot motion consists of a serpentine path composed of tangent arc segments. These motions result in oscillating modeling coverage of the environment in front of and behind the robot, thus providing the critical information necessary for automated path planning and obstacle avoidance. The caveat to this system is that exceptional inertial measurements are required in order to integrate the acquired data into consistent models. Rotations of the body result from varying drive torques and terrain and must be accounted for in scans. Additionally, large, soft, inflated wheels drastically reduce the accuracy of odometry and yield inaccurate arc following. Finally, the orientation of the scanner disallows position estimation through scan matching if the robot is moving in anything put a straight line.

A crude prototype of Cyclops' configuration was developed in order to test the feasibility of the concept. In order to increase available reaction torques an additional link was added to allow greater separation between the primary body and the wheel axis. Prototype testing demonstrated that insufficiency of obstacle climbing performance compared to the

expected subterranean environments. Only having two wheels providing force resulted in the prototype pivoting about a wheel after it encountered an obstacle, rather than climbing over it. Additionally, the failure mode of the gravity based reaction torque was found to be unacceptable. Maximum reaction torque is supplied when the wheel axis and robot center of gravity are horizontal. Increasing required wheel torque beyond this point puts the system into an unstable positive feedback loop that ultimately results in the device flipping over. While such a maneuver is not fatal due to large clearances around the main body, the implications (such as saturation of sensing due to frequent, large excursion motions) for inertial sensing and robot pose estimation are severe.

The Cyclops concept was terminated due to a variety of reasons including prototype performance (as outlined above), model integration, deployment, and inflatable technology. With scan matching not possible, the only avenue available to solving model integration for Cyclops was high end inertial systems, which likely would have dominated the overall system cost while still suffering from drift. Deployment was also a concern, primarily due to the estimated total length of the device. Many subterranean voids are only a few feet high. A robot with a single long rigid body would be limited to voids taller than the length of the robot. This also eliminates options for mechanically expanding wheels or increased range through added batteries, as both would add to the overall single segment length of the robot.

Inflatable wheels were troubling due to efficiency, reliability and extraction concerns. An inflated wheel loses energy as it deforms during rolling. The higher the inflation pressure, the higher the efficiency. However, low pressure tires will be compliant with terrain while puncture risks go up as tire pressure increases. Hence efficiency and reliability are mutually exclusive when dealing with inflated tires. Additionally, large balloon-like inflatable tires, while successfully demonstrated on some robots, are not commonly available and would therefore have to be custom fabricated. With a two-wheeled configuration, loss of integrity of one tire results in total device lose. Finally, while it is possible to hand pack and inflatable wheel, much like a parachute, prior to deployment, not such option exists for dealing with the balloon tire material when retrieving the robot at the end of a mission. Deflation is just as important as inflation for reliable device retrieval.

8.1.2 Four Wheel Inflatable Explicit Steered

An alternate configuration also employing inflatable wheels is shown in figure 8.2. Magellan's four powered wheels provide the torque needed to climb obstacles. An actuated center pivot, commonly seen on front end loaders, provides steering but only allows traversal of larger radius arcs. Turning in place is not possible with this configuration. The front and rear wheels are powered by independently-driven differentials, much like Ground-

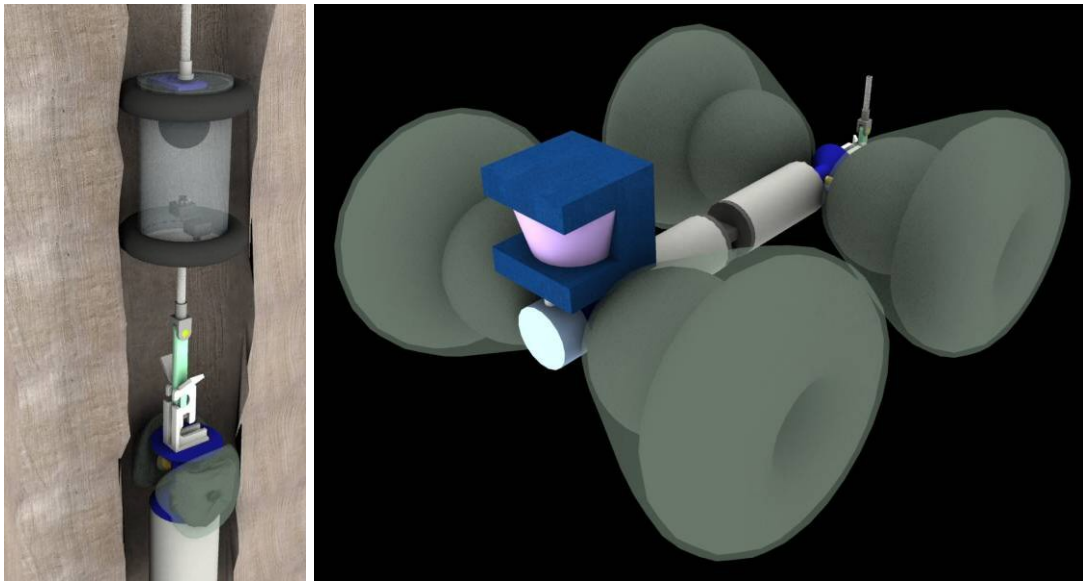

Figure 8.2: Magellan conceptual renderings: base station and robot rear wheels during deployment (left), robot during exploration (right)

hog's configuration. Also like Groundhog, a front mounted tilting laser scanner provides 3D modeling data when the robot is static and is locked into a horizontal orientation beneficial for 2D scan matching during driving. Magellan's operational approach is very similar to that demonstrated on Groundhog due to similarities in the locomotion and scanning configurations.

One of the key elements introduced by this design is the idea of a downhole base station as well as a means of disconnecting and reconnecting the mobile device to the base station for deployment and recovery. The base station shown in the rendering consists of two inflatable bladders used to brace the device against the end of the borehole. A panospheric camera provides complete visual coverage around the entry point into the void, and a downward facing camera coaxial to the borehole monitors the base station to robot coupling. One concept for coupling the base station and robot is shown in figure 8.3. borehole-deployment results in relationships between the base station and robot that can be exploited to create a passive latching mechanism that guarantees secure suspension of the robot during deployment and retrieval while allowing the robot to drive away from the latch without requiring any additional actuation. This elimination of exposed external actuators through innovative mechanical design is common across the various conceptual designs for borehole-deployed mobile systems due to the absence of precedence and due to explosion proofing and environmental sealing requirements.

Configuration concepts developed for mobile, borehole-deployed systems of necessity place greater emphasis on environmental sealing and explosion proofing than any other

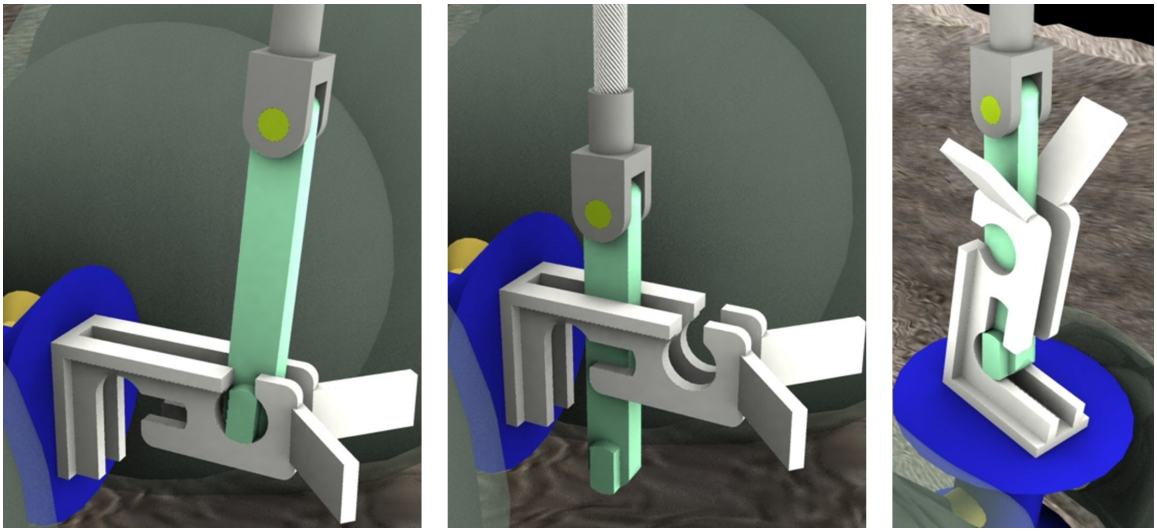

Figure 8.3: Coupling/Decoupling sequence showing geometric latching system: Decoupling after deployment (left), driving to latch the coupling (center), locking coupling during extraction (right)

class of device. Borehole-deployed static sensors avoid mud and water by remaining elevated above the mine floor. Mobile portal systems afford large ground clearances, keeping sensors and electronics enclosures away from hostile environmental elements. Mobile borehole systems are forced to operator closer to the ground due to their smaller size, and device reconfiguration often requires resting odd elements of the device against the ground. Subterranean devices benefit from being rated for operation in explosive environments, but during research and development experimental waivers of some safety regulations are often possible. Static sensors are exempted due to the circulation of fresh non-explosive air at the boreholes entrance into a void. Mobile portal systems can be exempted if they are operating in some naturally ventilated subterranean environments, such as the Mathies mine described in section 6.1.5. Inherent in the operations of mobile borehole devices is that they will operate in undisturbed and potentially explosive environments more often than other classes in the design space. Therefore sealing is required against both water and gases and actuation across sealed boundaries should be kept to a minimum.

The Magellan configuration was not selected for further development in part because of concerns about both sealing and fouling. Its modeling sensor, which is the first object to hit as the device is lowered onto the floor of a subterranean void, also requires wiring across the tilting mechanism exterior to the sealed and shielded central volumes of the device. The actuated center link reduces the required minimum void height for successful deployment, an improvement relative to Cyclops. However, Magellan configuration also suffers from the entwined problems of inflatable wheel efficiency and reliability outlined for Cyclops in the

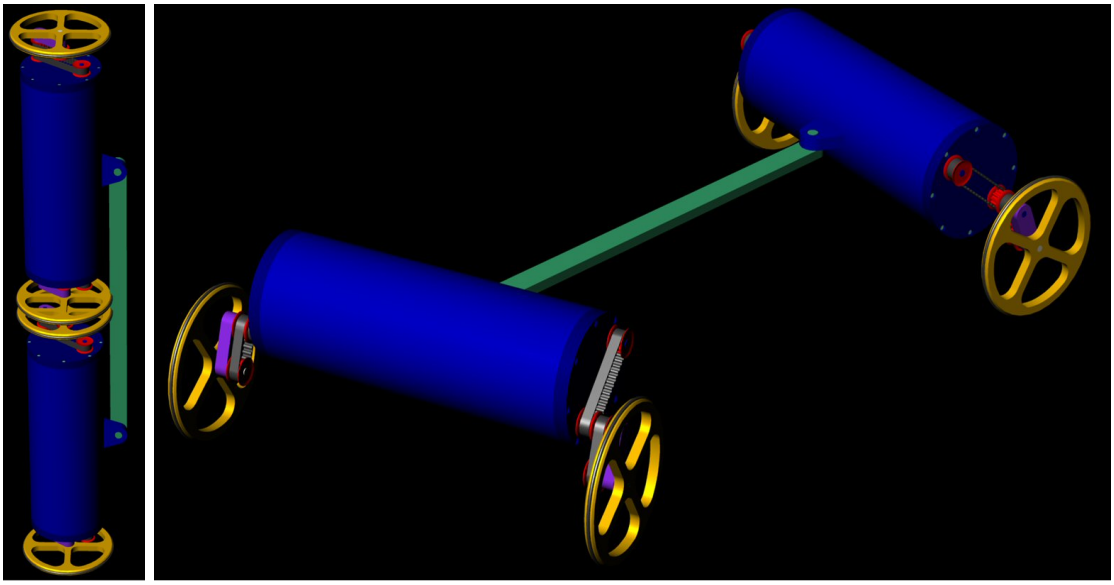

Figure 8.4: Four wheel, mechanically reconfigurable dual differentially driven concept shown stowed (left) and deployed (right).

previous section, 8.1.1. Additionally, the actuated center link would requires large swept volumes to keep the wheels from interfering, which results in the center of mass of the robot being well away from the center of the wheel base during tight turns. As wheel base is limited by the structural properties of the wheels, the solution common to front end loaders of widening the wheel base is not possible. Ultimately, these factors combined with the complexity of the design even in its earliest stage of development lead to abandonment of the concept.

8.1.3 Four Wheel Differential Drive

One of the larger drawbacks to the concepts discussed thus far is the need for inflatable wheels to increase wheel diameter and provide terrain compliance. Concepts which instead rely on mechanical deployment systems for increasing ground clearance and improving sensor vantage points eliminate the immature inflatable technology from the design while allowing more traditional robotic configuration concepts to be utilized. For example, figure 8.4 demonstrates a configuration loosely based on the Hyperion and Zoe mobile robots which utilizes a purely mechanical deployment system.

Four independently driven drive wheels allow sharp turns while guaranteeing power to all wheels by eliminating differentials from the system. In effect, the system acts as two differentially driven robots while the coupling between the sections acts to transfer reaction forces without the efficiency losses seen with drag tails. Ground clearance is provided after

deployment through an offset linkage. As rendered, ground clearance is slightly less than the minimum usable borehole diameter of the device, but multi-link systems with greater offsets could also be possible. Mechanical deployment has the advantage of providing a known configuration for extraction, versus the indeterminate, partially deflated state seen when using inflation to effect configuration change.

A three dimensional laser scanning system would be embedded in one of the two primary housings, with the other carrying computing and additional energy storage. The laser scanner requires two degrees of freedom, the drive wheels four, and if the deployment of the wheels can be mechanically linked across the primary housings, another two degrees of freedom are required for deployment. While this configuration eliminates the unknowns associated with inflation, it greatly increases the numbers of degrees of freedom and hence the complexity of the system. Additionally, synchronization between the four drive wheels across two semi-independent segments is required. The configuration and Magellan both suffer from insufficient simplicity at an early stage in the design.

8.1.4 Two Wheel Mechanical Helical Drive

Helix is a radically new configuration. Helix takes its name from the helical treads seen in figure 8.5 which theoretically provide both linear and rotary locomotion. In the course of development, helical vanes were abandoned for the advantages of paddlewheels and chevrons, but the Helix name remained. Utilizing only two primary actuators to provide both driving and steering ability, but configured to deploy mechanically and provide exceptional sensor positioning, this concept attempted to utilize an as yet undemonstrated means of locomotion.

Locomotion using two oppositely handed threaded cylinders is an established concept that has been deployed on amphibian and arctic platforms. This concept spins the two parallel drive cylinders in opposite directions, resulting in the thrust from both cylinders going in the same direction. This method of locomotion relies on sliding contact between the screw threads and the terrain, resulting in low efficiency but high force motions, especially suited for operations in ice and mud. Turning is accomplished by differentially driving the two cylinders, resulting in a net moment about the center of the vehicle.

Helix, unlike prior helically driven devices, employs two right handed threaded cylinders to drive like a steam roller. Driving both cylinders forward results in the robot driving forward. Driving the cylinders in opposition results in the robot turning in place as the screw threads provide equal and opposite thrusts forces to the robot. Deployment is achieved with the aid of gravity and the use of solenoids to lock-in deployed and stowed configurations. Helix's modeling sensor and computing are largely independent from the drive mechanisms and reside in a center segment that is elevated after deployment. This provides

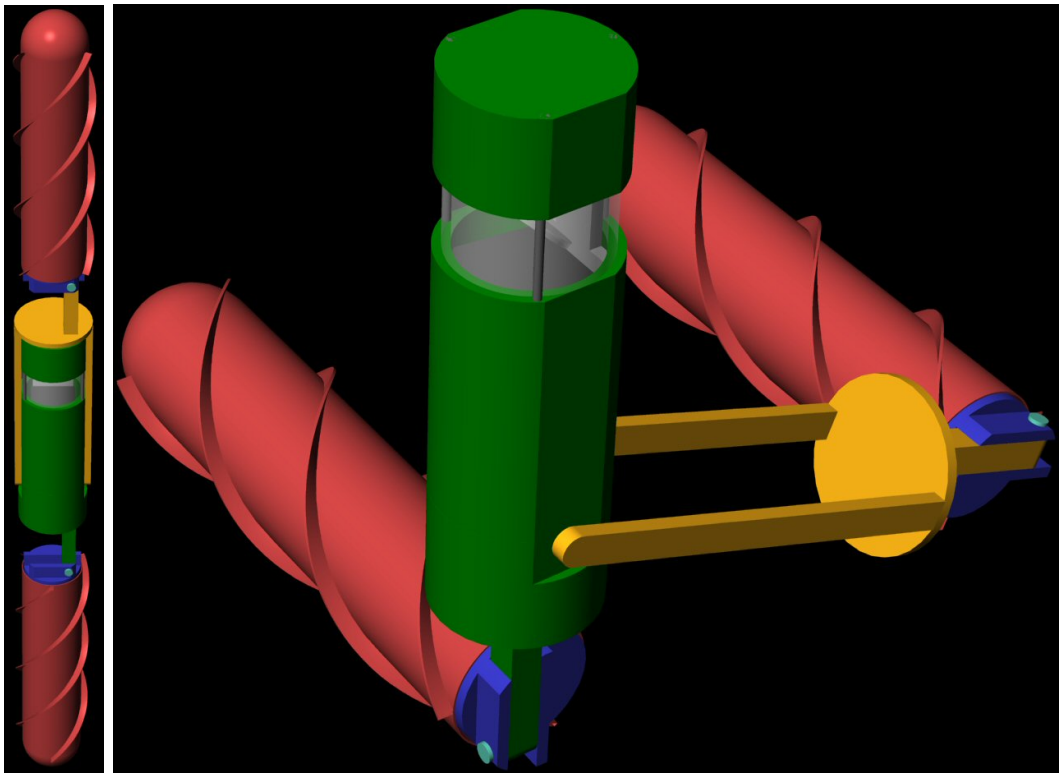

Figure 8.5: Helically driven and steered, mechanically deployed concept renderings: stowed configuration (left), deployed configuration (right).

an improved vantage point for scanning. Actuating the solenoid pin maintaining the angle of the laser housing and driving the two threaded cylinders apart or together results in 3D laser coverage. Sensor vantage (high) and nominal configuration (horizontal plane) are both favorable aspects of the Helix concept. Three roughly equal length segments allow deployment into shorter subterranean voids than any of the other concepts discussed thus far. Finally, storing the batteries and drive motors in the drive cylinders lowers the overall robot center of mass, providing greater stability.

A crude prototype of this concept was developed and tested. Testing demonstrated that the first implementation of the Helix locomotion concept was insufficient for capable operations. During driving over hard packed ground the device moved consistent with the direction of rotation. However, in softer terrains where the screw threads on the cylinders often bit into the ground, resulting in the device translating along the screw axis rather than perpendicular to it. This problem was especially apparent when a rigid obstacle was encountered. Rather than climbing the obstacle, the prototype would instead circle around it. While this mechanical encirclement of obstacles is interesting from a mobility stand point, it doesn't achieve the application objective.

Figure 8.6: Early Helix prototype used to test mobility with helical treads (left) and straight treads (right).

An alternate but similar configuration was developed to address these mobility concerns. Shown in figure 8.6, this concept more directly mirrors steam roller locomotion, with steering accomplished by angling the two drive cylinders. In order to avoid lateral translation when an obstacle is encountered, the drive cylinder tread was changed to have no screw components. Testing results from this modified prototype were improved versus the prior configuration, with obstacles becoming traversable and motions becoming deterministic. Of the four configurations presented, Helix was deemed to be the most likely to succeed for the class of borehole-deployed mobile robot. The Helix design was refined and iterated and a full system prototype was developed.

8.2 Representative Implementation: Helix

Despite losing the helical treads from which its name was derived, the Helix concept retained its name and received a major redesign. The resulting concept, shown in figure 8.7, shows the reconfigured system which incorporates all the necessary elements of a subterranean mobile modeling robot. The two drive cylinders contain all the motors, control electronics, and energy storage to provide complete locomotion. The center section provide computation and houses the laser range finder and its associated scanning hardware. Details of each of these aspects of the system design are detailed below.

8.2.1 Helix Mechanical Design

Helix is designed to be deployed down lined six-inch boreholes. When stowed for deployment, it is five and a half inches in diameter and six feet in length. After deployment, Helix's eighteen inch wide drive cylinders are nominally two feet apart and its laser scanner

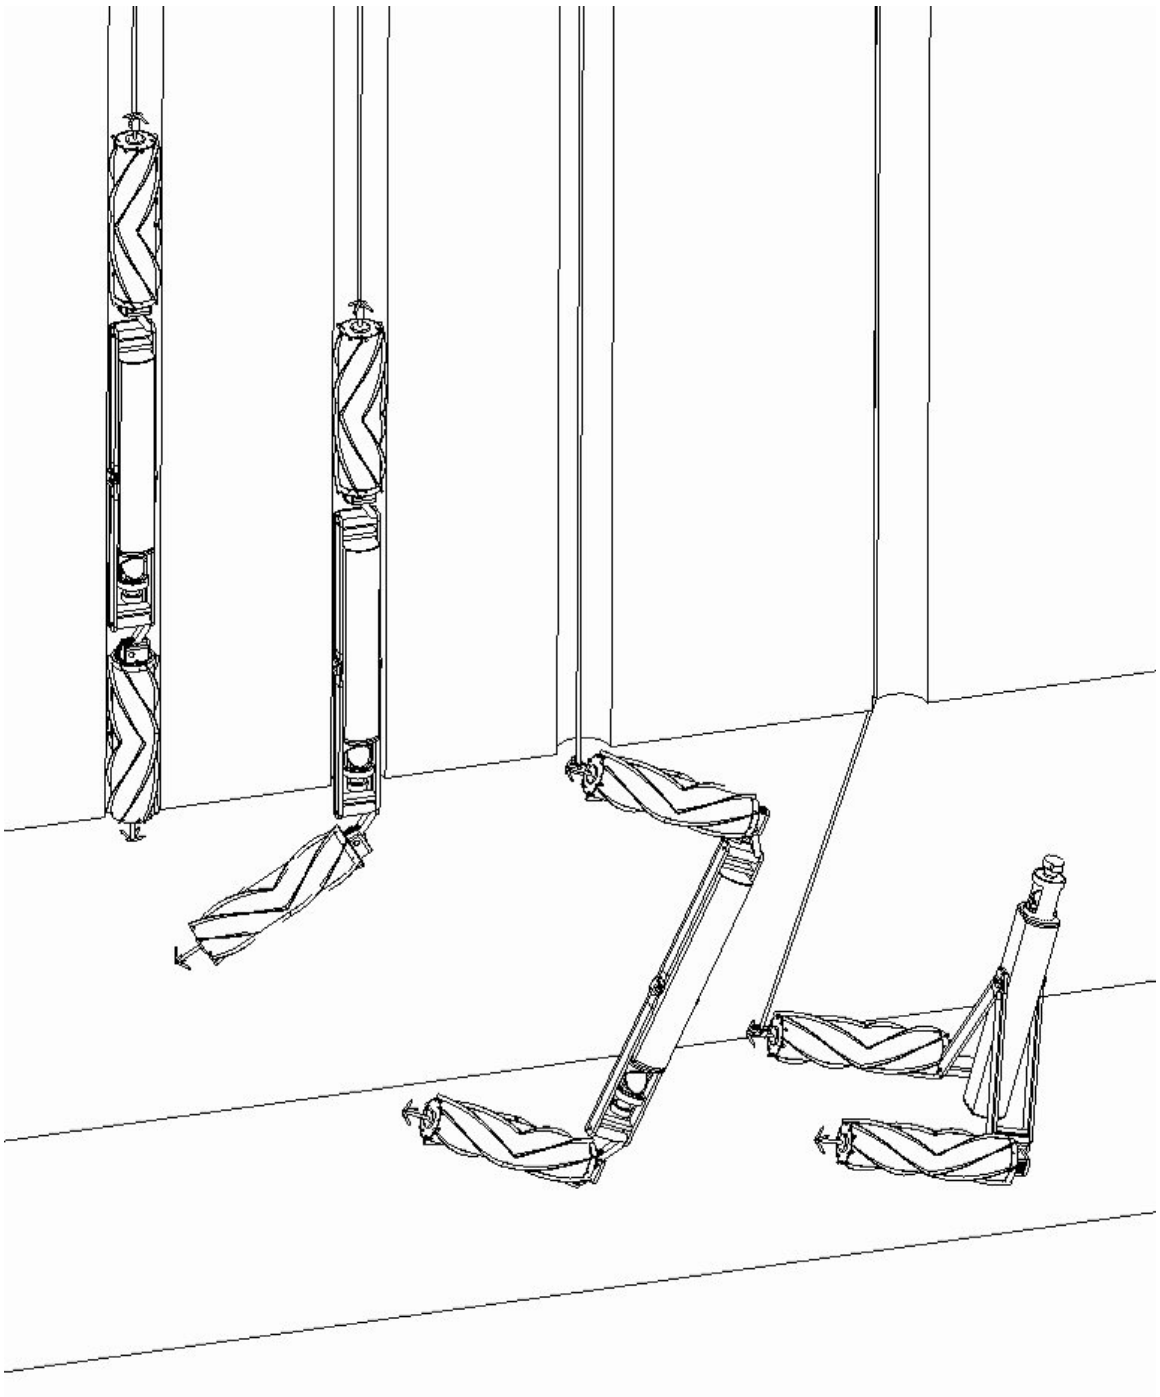

Figure 8.7: Helix deployment scenario.

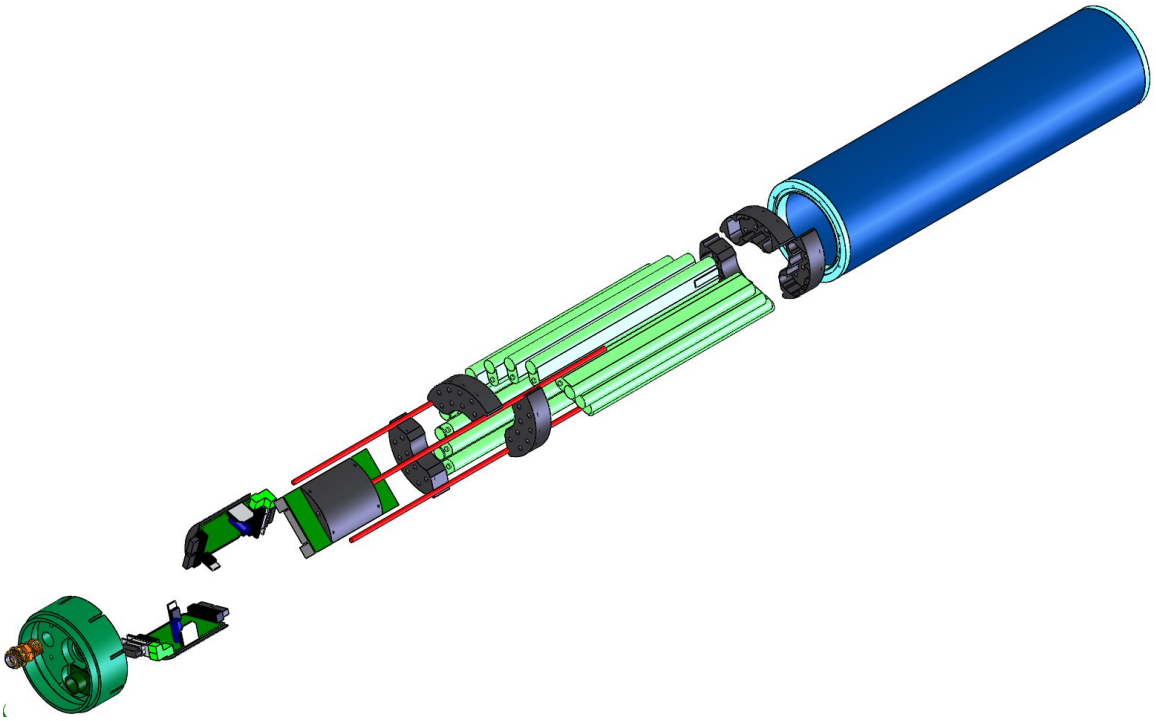

Figure 8.8: Exploded view of one of Helix's drive cylinders shown battery and electronic packaging.

sweeps a plane thirty inches above the ground. It weighs in at seventy pounds and has an estimated round trip range of four hundred feet. Helix can fit down smaller diameter holes than the current field deployed FerretIII.

Helix achieves locomotion and scanning through six independent motions. Two motions in the center section provide three dimensional laser range finder coverage, while two motions in each of the drive cylinders provide driving, steering and deployment capabilities. Mechanically, parallelogram based arms connect the drive cylinders to the center section. Electrically and environmentally, wiring and the sealed pressure volume of the center section are connected to the two drive cylinders via flexible tubing.

An exploded view (figure 8.8) shows the components packaged into each drive cylinder. Each drive cylinder consists of a tightly packed and environmentally sealed inner cylinder, surrounded by a rotating outer cylinder which converts drive torque into linear robot motion. The inner cylinder is primarily filled with batteries, twelve stacks of ten AA batteries, packaged in a two wide by five high configuration. The end of each inner cylinder that is closest to the center section houses three printed circuit boards. Two commercial motor controllers for controlling the drive rotation and angle of the cylinder, and one custom board designed to handle power distribution and battery management. Both ends of the

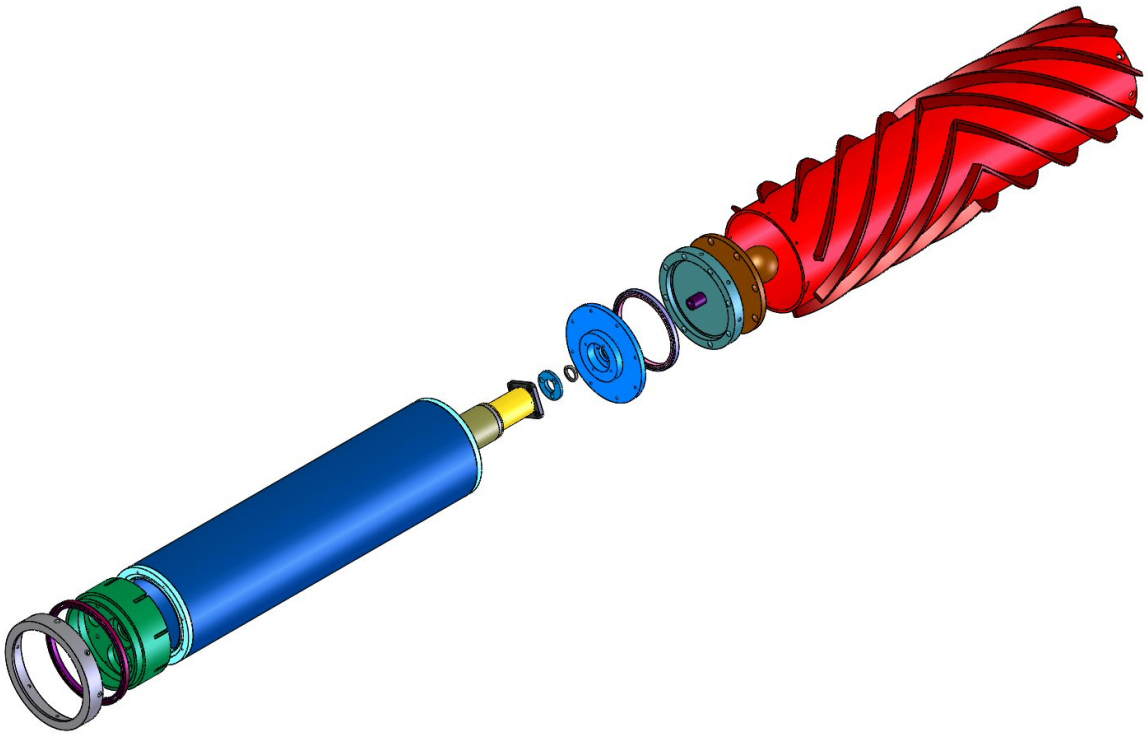

Figure 8.9: Helix's drive system, exploded view.

inner cylinder are capped welded flanges which feature an O-ring groove and radial bolt hole pattern for sealing and mechanical connection to the rest of the system.

Rotation of the outer cylinder is accomplished through a single brushed motor acting through a 5000:1 planetary gearhead, as shown in figure 8.9. The motor bolts to a static end cap for the inner cylinder. A rotatory end cap is supported by a thin section ball bearing and interfaces to the motor with a broached shaft. An internal dynamic canted coil lip seal between the rotary and static end caps maintains the environmentally seal of the device. Another thin section bearing interfacing between components on the opposite end of the drive cylinder provides a second bearing for passive support of the outer cylinder. The rotary end caps are designed to allow rapid change out of outer cylinder tubes for testing of new tread configurations, repair of damaged treads, or access to the interior of the drive cylinder. The majority of Helix's cross section as viewed from the ground is composed of actuated drive cylinders. This, combined with slow speeds, allows Helix to operate in the quasi-static region. Whereas CaveCrawler deals with rough terrain through superior compliance, Helix's configuration solves the problem by actuating nearly all possible contact surfaces with the ground. Hence Helix operates without compliant suspension. The low center of gravity due to battery placement stabilizes Helix for scaling large obstacles through improved stability, even in the absence of an engineered terrain compliance system.

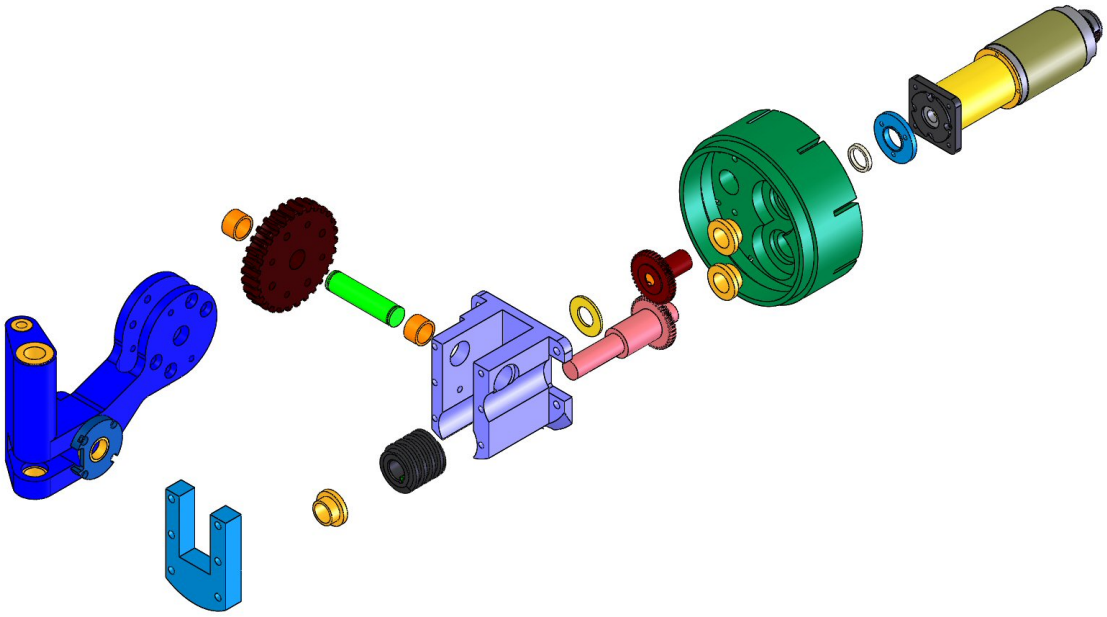

Figure 8.10: Helix's steering system: exploded (left) and cutaway (right)

Drive cylinder angle relative to the center section is controlled by a motor mounted to the other end of the drive cylinder, as shown in figure 8.10. The steering motor, geared by a factor of 1000, connects through a dynamic seal identical to that around the drive motor to a spur gear at the start of a transmission. A mating spur gear drive a worm which in turn rotates a worm gear rigidly mounted to the termination of the deployment mechanism coming from the center section. Each cylinder steers thirty degrees forward or back from the nominal (straight ahead) configuration. This results in two useful, but non-symmetric steering configurations. The robot will follow the arc centered at the intersection of the two drive cylinder axes and going through the centers of the two drive cylinders. However, since the steering pivot point is actually outside the rotating portion of the drive cylinder the minimum possible arc occurs when the far ends of the drive cylinders are almost interfering. Even though the range of motion of the steering actuator is symmetric, the resulting traversable arcs are not, in effect favoring right turns (toed-in steering configuration) over left turns (toed-out steering configuration), as shown in figure 8.11.

Designing and implementing an effective deployment mechanism was one of the most challenging aspects of Helix's configuration. The original concept of utilizing the weight of the device to actuate its deployment provide infeasible due to frictional forces involved, even in a pristine lab environment. Numerous alternatives were explored, such as spiral cam paths and additional actuators, before the system shown in figure 8.12 was finally devised. A pair of struts in a parallelogram configuration keep the two steering joints parallel with

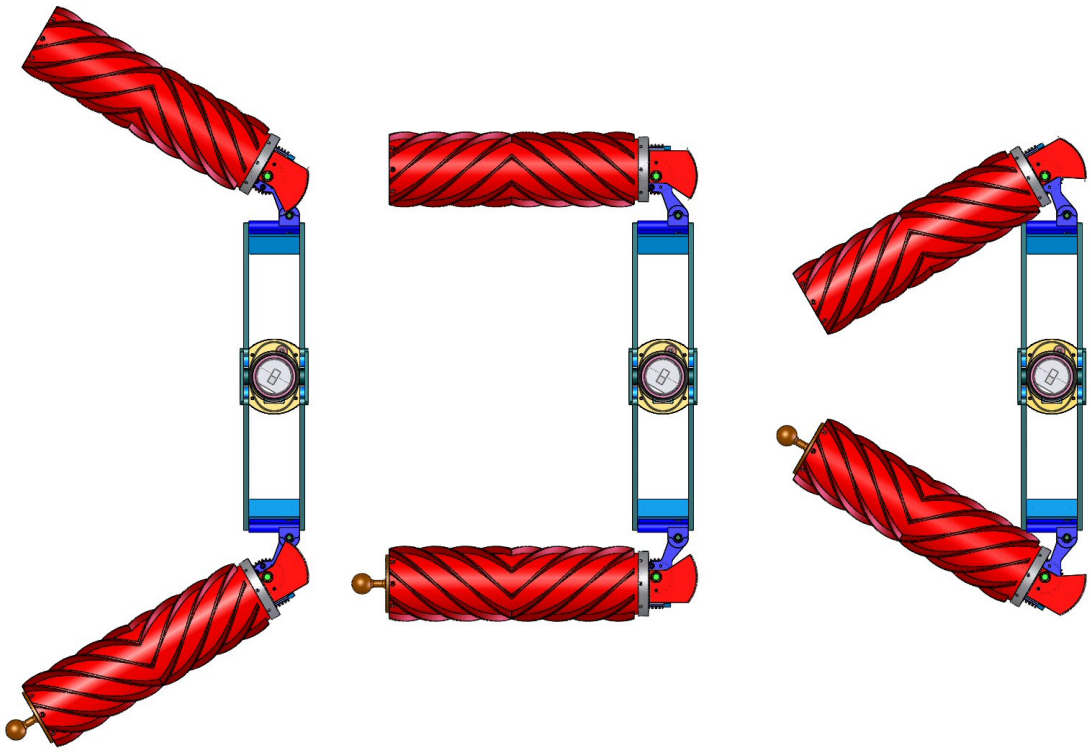

Figure 8.11: Possible steering configurations: toed-out (left), straight ahead (center) and toed-in (right)

the center joint of the center section in all configurations, thus preventing out-of-plane motions when steering. This design also provides hard stops at both ends of the range of motion for the deployment mechanism, corresponding to a linear and an equilateral triangle configuration at the two ends of the sixty degrees of travel. Actuation of the deployment mechanism comes from sixty degrees of otherwise unused steering angle motion. A sector gear interfaces with a small spur gear which in turn drives a worm. The corresponding worm gear is rigidly attached to the larger of the two parallel deployment struts. This mechanism enables deployment to a known configuration without requiring any additional actuation and without compromising any system seals.

The deployment struts terminate in two plates which are attached to the shaft about which the center section of the robot rotates. The interior of this shaft provides passthroughs for wiring, while a large diameter worm gear is mounted to its exterior. The center section tilt angle is controlled by a worm connected to a motor housed inside the central body, as shown in figure 8.13. This tilting action is necessary for shifting from stowed to deployed configurations and also provide the second motion necessary to provide three dimensional range data coverage. Additionally, in rough terrain where the ground clearance beneath the center section becomes an issue this motion allows the central body to be stored parallel

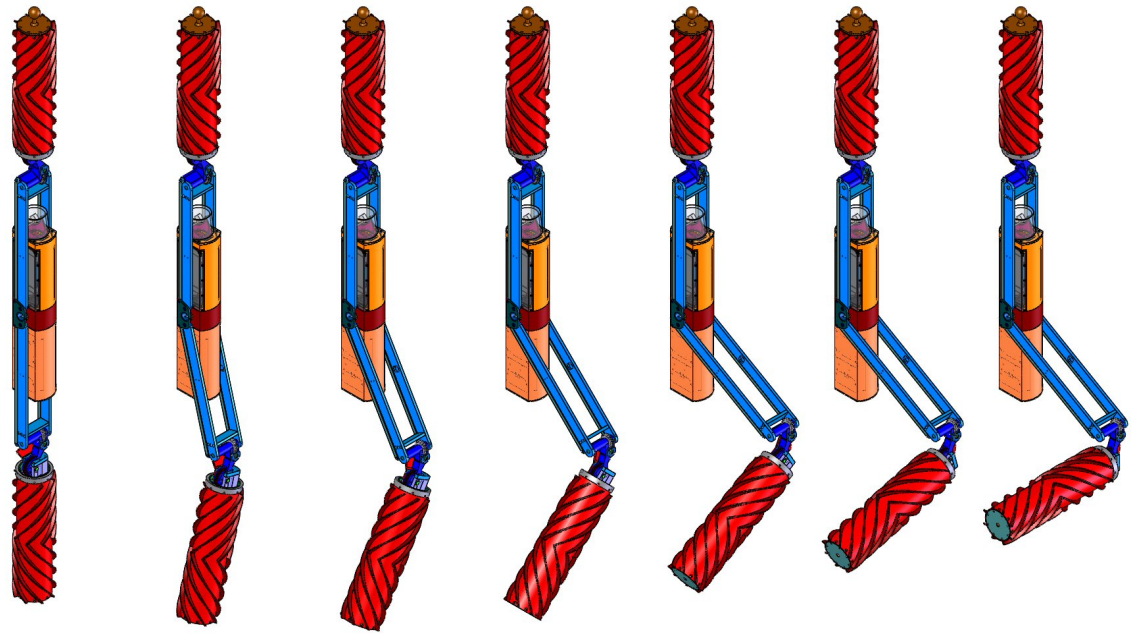

Figure 8.12: Helix's deployment system: coupled with steering angle for sixty degrees (first five images), decoupled and held in place by a worm gear for steering angles between 60 and 120.

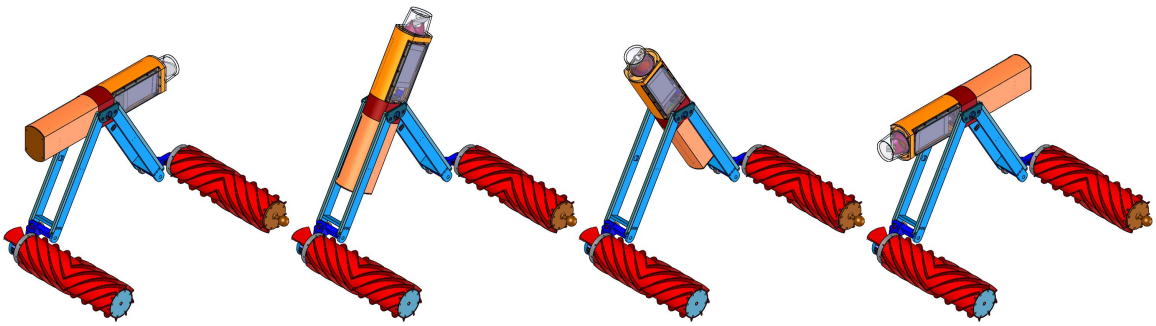

Figure 8.13: Tilting of the central body of Helix to provide 3D laser range finder coverage.

with the deployment struts, thus presenting the minimum possible cross section capable of bottoming out the robot. During nominal operations, the joint only passes through 180 degrees of rotation, limiting the required service loops on the wires from the drive cylinders while still providing maximum possible sensor coverage.

The center section is roughly balanced about the tilt shaft both geometrically and in terms of mass. This reduces the torque required to rotate the center section while also providing two large bays for electronics and instrumentation. The lower center electronics bay houses Helix's computer, some power electronics, and the motor controllers necessary for

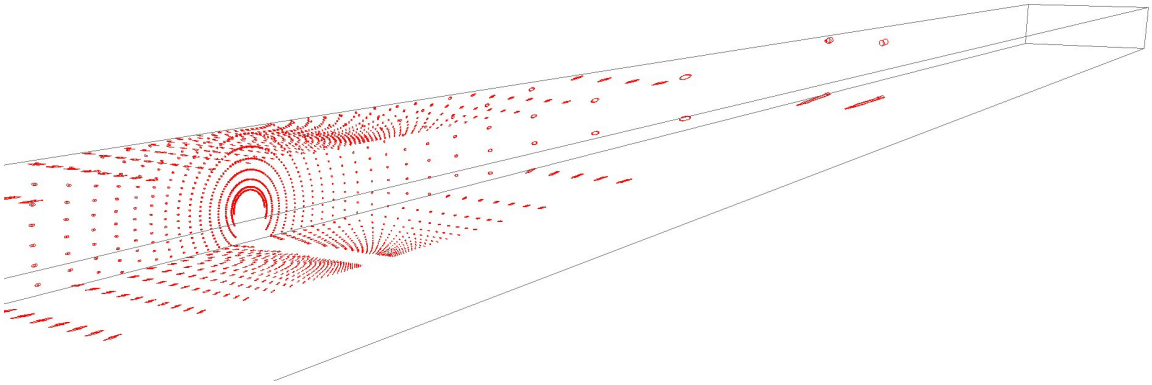

Figure 8.14: Helix's sensor coverage pattern and data density looking down a corridor.

providing 3D scanning capability. At the bottom of the computer bay is an interface panel behind a sealed cover which allows rapid interfacing to the computer, power, and motor control systems for testing, field repairs and development. The upper center electronics bay houses the laser range finder and an inertial measurement unit. Also in this bay are the motors and mechanisms for controlling laser scans. The single point laser is coaxial with the center section and is aimed at a 45 degree mirror. The mirror is rotated to produce a plane of laser data by Helix's sixth and final motor.

8.2.2 Helix Electrical and Sensing Design

The combination of rotating mirror and tilting center section results in nearly full spherical coverage minus an infinite length cylinder coaxial with the center section tilt axis. The cylinder has a radius equal to the distance between the tilt axis and the point where the laser beam intersects the mirror, as shown in figure 8.14. This cylinder without coverage is not aligned with the direction of robot motion, so as soon as the robot has moved roughly one center sections length the missing data will be filled in by a new scan. Helix subscribes to a concept of exploration similar to that used by Groundhog, where terrain is scanned in three dimensions while scan matching for path tracking occurs in only two dimensions, hence the nominal position of the laser is to scan a horizontal plane while the center section if vertical.

Helix's laser range finder has the smallest cross section of those which were commercially available at the time of its purchase. However, this small size comes at the cost of laser range and support electronics. Similarly, Helix is equipped with a six axis IMU, but its miniature size and MEMS components result in high drift rates and noisier signals, even when the robot is static. Helix sacrifices some modeling capability in order to meet entry constraints, but mobility, especially in small subterranean spaces, has great advantage and

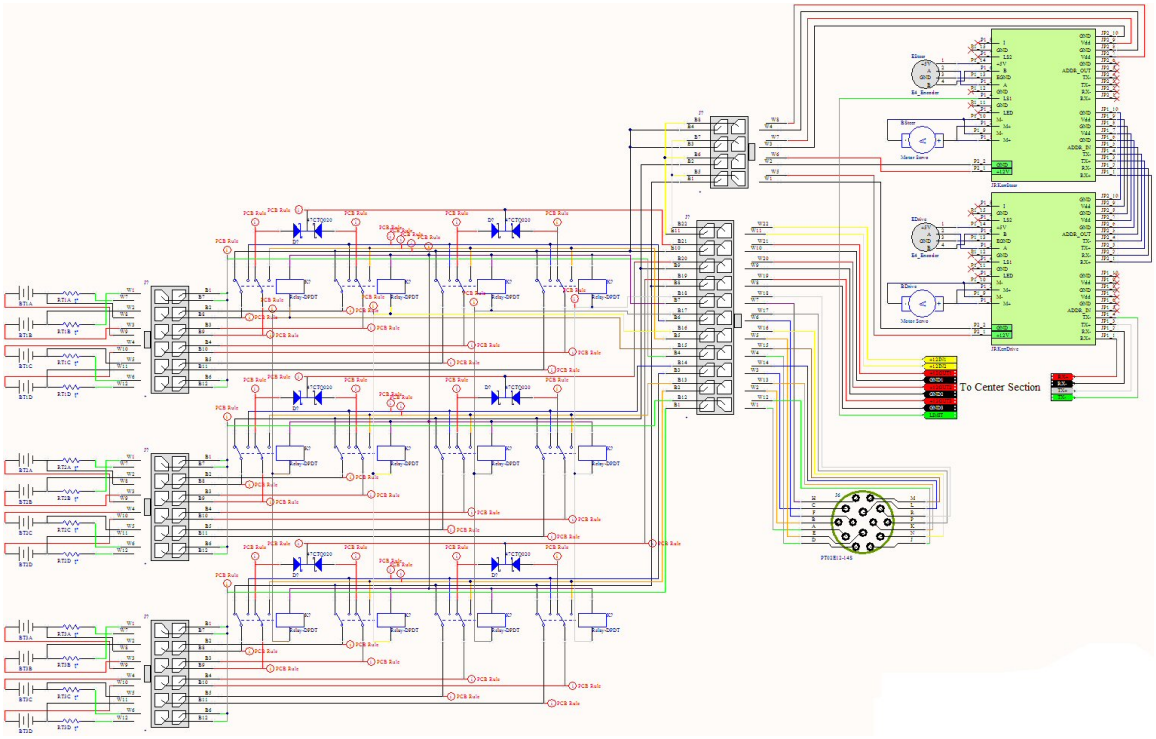

Figure 8.15: Schematic of Helix's drive cylinders. From left to right, three groupings of four battery packs each, relays and protection diodes for charging, charging connector, drive and steer motor control boards, and finally connections to the central electronics housing.

capability relative to the loss of static scan capacity.

Mobility requires energy storage. Helix's power system is built around 240 AA nickel metal hydride cells, wired as 24 parallel stacks of 10 cells each to provide a nominal 12VDC system bus. Each drive cylinder contains an identical package of 120 batteries. The risk of wiring battery stacks in parallel is that a weak or failed string will discharge the remaining cells, thus drastically reducing overall energy storage and ultimately robot range. To protect against this each string of 10 cells is wired into a reverse polarity protection diode, as shown in figure 8.15. Charging so many batteries is also challenging, especially considering the restrictions on wire gauge and numbers due to mechanical limitations. The schematic reference above shows the relay board inside each drive cylinder that allows all 120 batteries to be charge through a single eighteen pin connector, the maximum allowable pin count given size and connector pin gauge constraints. Only three battery stacks per drive cylinder can be charged at one time, requiring four one hour charging cycles for each robot recharge. Each cell has a 1600mAh capacity at 1.2V, resulting in a total robot energy storage capacity of 460Wh.

Power lines from the two drive cylinders are bound together in the center section. A

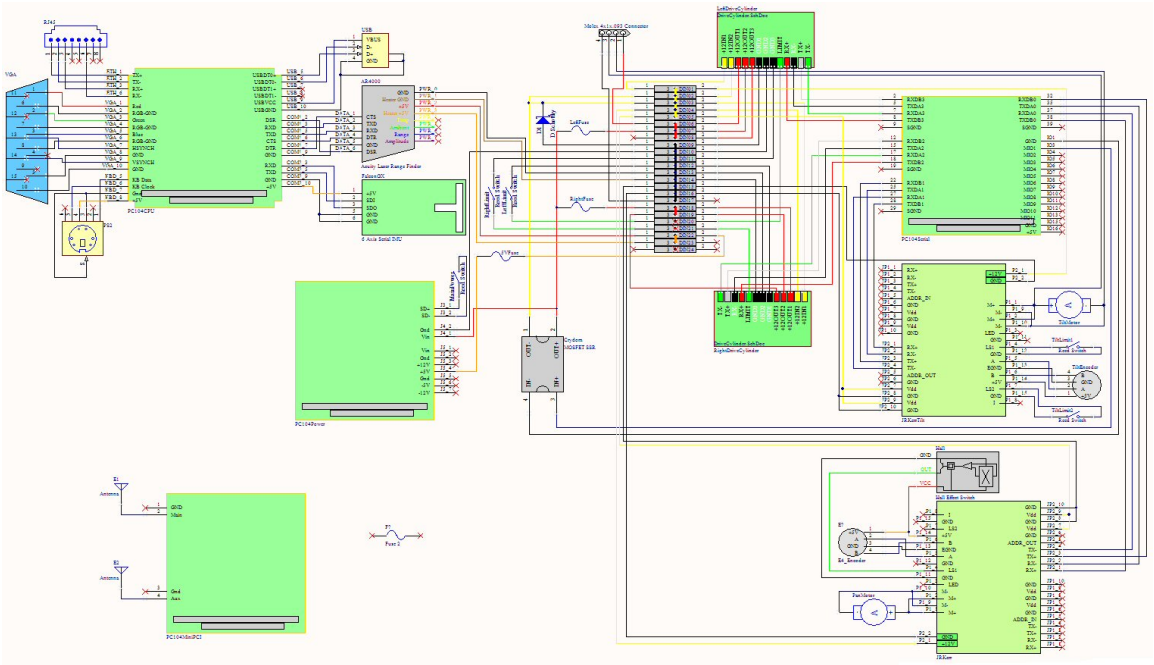

Figure 8.16: Schematic of Helix's central electronics enclosure. From left to right, PC104 motherboard with interface connectors, laser range finder and inertial sensing, power control and wiring to drive cylinders, and motor controllers for the two axes of laser scanning.

single solid state relay controls power flow to the motor power bus, as shown in figure 8.16. Helix has only two power buses, the primary 12VDC bus which powers the laser, all six motors, and a PC104 form factor DC/DC conversion card, which generates a 5VDC bus for the computing stack and the inertial measurement unit. The computing stack incorporates a 600Mhz motherboard, wireless Ethernet adapter, and quad RS485 serial port card. Both motor controllers in each drive cylinder are controlled through a single RS485 bus. Interfacing to the IMU and laser is accomplished through two RS232 ports integrated into the motherboard. Overall, Helix has a minimalist electrical system due to tight volumetric constraints disallowing many of the additional components found in configurations such as CaveCrawler or FerretIII.

Helix's knowledge of its mechanical state is minimal but sufficient. Magnetic reed switches in the center section respond to a magnet mounted on the central worm gear attached to the tilt shaft. This enables homing of the encoder monitoring center section tilt angle. Reed switches mounted within the center section sense magnets on the deployment struts through the aluminum of the center section housing. The geometrical configuration requires the tilt axis angle to be known, but also means that the switches only trigger when the robot is in its stowed position. Therefore the homing procedure is to rotate the tilt axis until a limit switch is detected, then move the tilt axis to its stowed configuration. Next

the deployment mechanism is actuated to stow the device until the deployment strut limit switches trigger, thus indicating the robot is in its home or stowed configuration. Once absolute position is determined through this homing procedure, incremental optical encoders mounted to all six motors maintain angular position information.

Helix's software system directly controls each motor velocity while an independent data logger records range scan information. While Helix's software is limited but well designed from a testing and debugging standpoint. Nearly full access to primary computer functions is provided in a single control panel. This panel also houses the primary power switch for the robot, and e-stop which cuts power to the motors, and over ride lines for the tilt access, as only certain tilt angles allow unobstructed access to the laser and computing bays in the center section. While these features are useful during debugging and development, during operations a wireless Ethernet connection and mechanical steel cable are the only external interfaces provided to the robot.

8.2.3 Helix Field Experiments

Much of the testing relative to demonstrating Helix's configuration has occurred in the lab. Figure 8.17 shows a sequence of images of Helix being deployed from its stowed configuration. Steering was also tested and both toeing in and toeing out were shown to be effective means of steering. Laser scans of the highbay area in which testing occurred were also taken, and are shown in figure 8.18. Early testing demonstrated the sufficiency of Helix's configuration for entering, exploring, modeling, and exiting subterranean voids via boreholes. However, some deficiencies in the design were also discovered.

Problems in the deployment sequence and the center body tilt axis were discovered during initial testing. Ideally, effective deployment could occur with the robot suspended in the air inside a void, or with the deployment line slack and the robot laid out on the void floor in its stowed configuration. In practice, careful coordination between the tether spool and Helix's deployment mechanisms are necessary in order to reconfigure the robot. The combination of required gearing ratios, available volume, and large moment loads incured during the transition from stowed to deployed disallows Helix from actuating its deployment mechanism while suspended. Additionally, even should fully reconfiguration be achieved, the tether termination at the end of one drive cylinder combined with the center of mass of Helix results in Helix being placed on the void floor upside down. Therefore deployment must occur while Helix is in contact with the floor both to avoid damage to the mechanism and to ensure the robot is in its upright configuration following deployment.

Actuation of the tilt axis results in the center section sweeping through the two sets of parallelogram deployment struts. While the struts are very rigid in one plane, they are weak in torsion. When Helix is traversing uneven terrain, the weight of the drive cylinders

Figure 8.17: Helix transforming from its stowed (left) to deployed (bottom right) configurations.

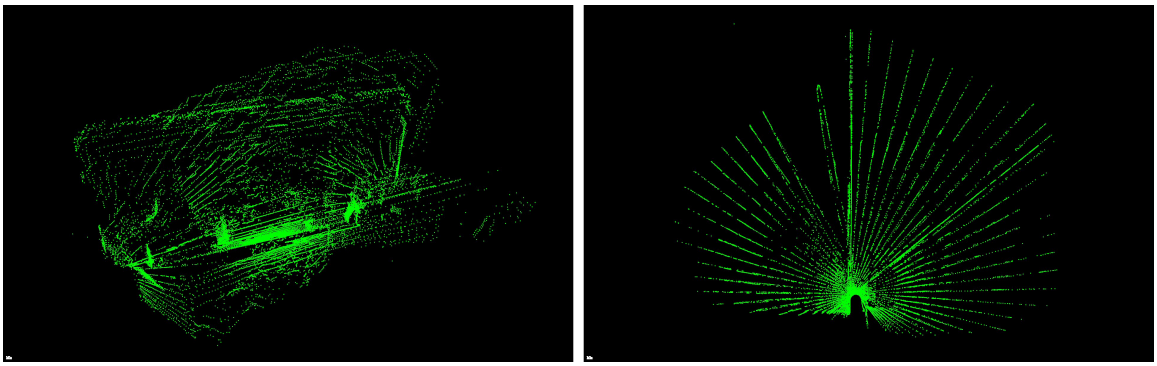

Figure 8.18: Laser scans of an industrial highbay: perspective (left) and with the tunnel of missing data emphasized (right).

often induces torsional stresses, resulting in visible displacements of the deployment struts. These displacements in turn infringe upon the swept volume of the center section, which when combined with the aluminum on aluminum contact between center section and strut results in binding, galling and ultimately stalling of the tilt axis motor. A redesign of the deployment struts and their attachment to the center section will solve this problem.

8.3 Summary

There are no precedents for constrained entry, limited mobility subterranean modeling configurations, so these new designs are seminal. Helix is the first-of-class concrete, realized configuration which demonstrates all the elements necessary to successfully access, explore and model subterranean voids. Helix's configuration combines deployment mechanisms, locomotion elements, and three dimensional scanning into a steam roller like platform that is simple given the varied tasks it must perform and effective underground with its deployment down six inch diameter holes and five hundred foot range.

Borehole-deployed mobile robot configurations combine some of the miniaturization of static sensors with some of the mobility of portal deployed mobile systems. This middle ground acquires more data from each borehole, enables modeling of some subterranean areas that can't be drilled into due to conditions on the surface, and allows inspection around the corner underground. Therefore constrained entry, limited mobility configurations are demonstrated to be possible, useful, and necessary when complete modeling of difficult to access subterranean voids is required.

Chapter 9

Conclusions

This thesis conceives, classifies, develops and deploys robot systems to model subterranean voids. These robots are distinguished for their relevance in two types of subterranean void entry and three categories of locomotion, resulting in six classes of robot designs (figure 9.1). Unlimited mobility can't be achieved when an entry to a void is constrained. Two classes of robot design, unconstrained entry, zero mobility and unconstrained entry, limited mobility, are achievable but do not fully realize subterranean modeling robot potential. The remaining three regions in the design space span the vast majority of technical issues and application needs. These robotic classes compel new robot configurations that are tailored for subterranean exploration and modeling. This thesis explores multiple design configurations for unconstrained entry, unlimited mobility; constrained entry, zero mobility; and constrained entry, limited mobility.

Unconstrained entry, unlimited mobility devices explore vast subterranean expanses using configurations descendant from conventional mobile robots. However, these configurations are unique due to scanning range sensors which improve three dimensional modeling coverage, lowered profiles which enable operations in environments with obstacles rising up from the ground and protruding down from the ceiling, and navigational sensors and systems which can not rely on GPS to provide localization information. Two robots, Groundhog and CaveCrawler, demonstrate the potential of this class of device and have produced subterranean models of unmatched extent, density and accuracy.

Constrained entry, zero mobility devices can be deployed into voids through the smallest of boreholes drilled in the remotest and most difficult to reach locations. While limited to a single vantage point, these device provide close to full spherical modeling coverage and are the cheapest, most reliable means of acquiring model data when borehole locations are favorable for modeling. This class of device also introduces new operational approaches, where subterranean void data acquired from existing holes is used to place new holes. The four Ferret devices embody this configuration and have demonstrated their worth through

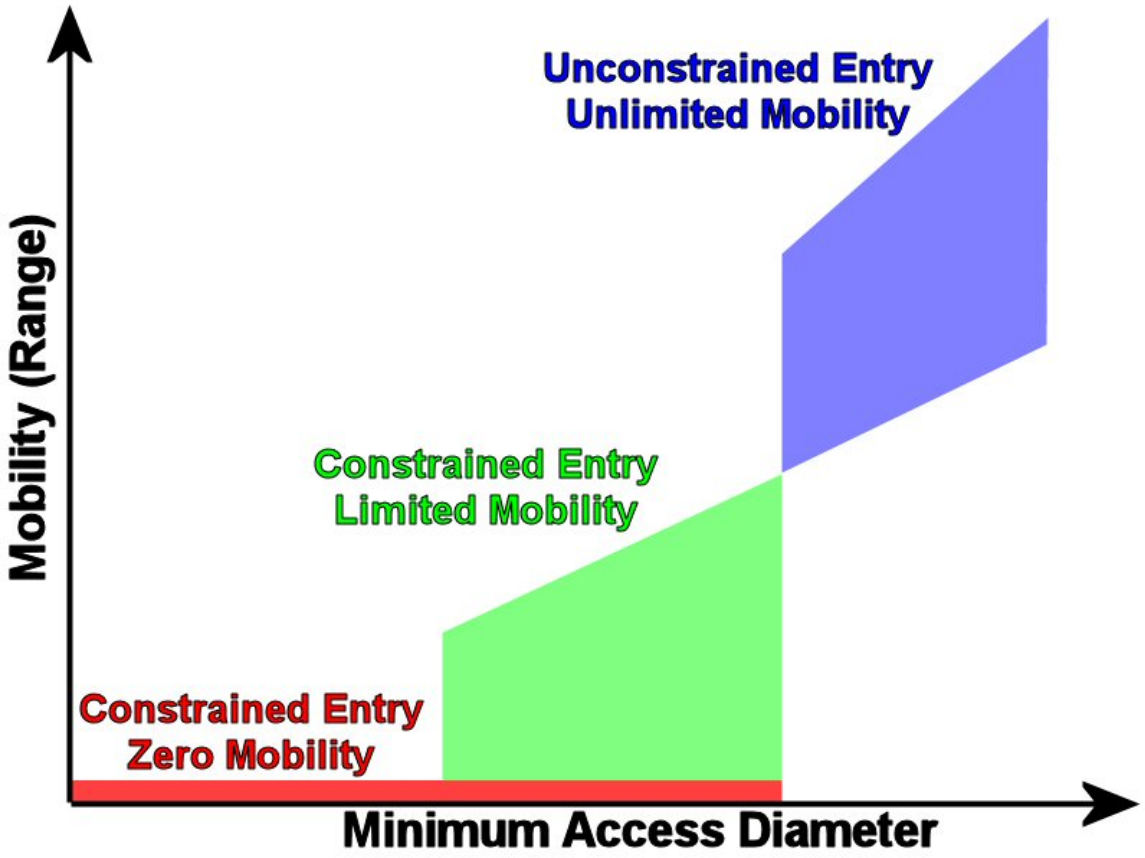

Figure 9.1: Regions in the design space of subterranean modeling robots.

a series of successfully completed modeling tasks. Constrained entry, limited mobility configurations achieve some of the mobility of unconstrained mobile robots and exhibit some of the miniaturization of static sensors. Robots of this class cover a range of modeling tasks that would not be possible otherwise by improving the extent of modeling coverage possible from a single borehole. Helix embodies one design within this class and demonstrates borehole-deployment, self reconfiguration, terrain traversal, and three dimensional modeling, thus demonstrating the feasibility of robot configuration within this challenging region of the design space.

Subterranean modeling is a difficult task, and not solvable with a single configuration, concept, or method of operation. Research conducted for this thesis envisioned concepts, developed robots, and produced hitherto impossible models which address subterranean challenges. Ultimately, this thesis distinguishes the core challenges and early solutions to address the problem of robotic modeling of subterranean voids.

9.1 Contributions

1. This research investigated the unexplored area of robot configurations for modeling subterranean voids.

2. This work distinguished three classes of underground void modeling robot. Static sensors can be deployed down the smallest boreholes, providing spherical range data from a fixed vantage. Mobile portal robots can use their superior mobility, sensing and computing to explore vast underground voids from a single large portal. Mobile borehole robots combine some mobility with borehole-deployment capabilities to explore beyond what is observable from a fixed vantage. Understanding when to use each type of robot and the transitions between these classes of robot in the configuration design space of subterranean robots is a core contribution of this thesis.

3. This work acquired robotically generated three dimensional subterranean models of unprecedented resolution, scale, and accuracy. Robotically produced data, maps, and models of voids are changing the way the world views, understands, and interacts with the subterranean world.

4. This work developed numerous robots which are contributions in their own right. Through the design, development and testing of a half dozen robots this thesis has demonstrated the effectiveness of robots in exploring and modeling subterranean voids. The capabilities and limitations of various configuration archetypes has also been determined in the real world. The robots developed are now ready to serve as software development platforms for the next generation of subterranean robotics researchers.

5. Beyond the intellectual and creative aspects of this research, the quality, quantity and significance of field data is unparalleled. On order of 1 billion subterranean void model data points were collected during the course of this research. This data was gathered from a half dozen robots that were envisioned, prototyped, and deployed into underground voids. One large, geo-referenced model, acquired from multiple modeling experiments conducted over the course of months, contains roughly a half billion data points, making it the largest subterranean void model yet acquired. Void data acquired during the course of this research provided definitative answers to otherwise unanswerable questions about underground voids. Data acquired has been used to determine underground pipeline route viability, backfill effectiveness, and foundation competence. Void data acquisition was made possible by the innovative robot configurations developed in this thesis.

6. This initiative patented important results of the research. Operational concepts and the core ideas related to robotic modeling of voids have been patented under US patent #7069124.

9.2 Future Work

Underwater This thesis investigated operations in dry voids. Water-filled voids are equally important. Many voids are completely submerged, and a full range of conditions, from fully dry, to lake like, to fully submerged are possible. An additional axis representing the conditions within voids could be added to the design space graph shown in figure 9.1, resulting in a 3D graph with a third access measuring water depth relative to void floor. Underwater robots will evolve to map submerged voids. Submergence might make some problems easier. Underwater robots could effortlessly swim or float through situations impassable by robots in air-filled voids. Some initial conceptual work in this areas has been performed, but much remains to be done.

Optimization This thesis has demonstrated first configurations and iterated improvements for the identified subterranean robot classes. It has not demonstrated optimal configurations. Applying configuration optimization techniques[Leger, 1999] to the domain of subterranean robots has great potential, but has yet to be explored.

Autonomy Mechanisms, sensing and operational concepts for subterranean robots have been developed in this research. The innovations are in full collaboration with co-development for autonomy, mapping, exploration and intelligence software. The synergy has elsewhere been referred to collectively with configuration design as automobility[Apostolopoulos, 2001]. Continued collaborations will transform subterranean robots from electro-mechanisms to automatons.

References

[Altafini, 1999] Altafini, C. (1999). A path-tracking criterion for an LHD articulated vehicle. *The International Journal of Robotics Research*, 18:435–441.

[Apostolopoulos, 2001] Apostolopoulos, D. (2001). *Analytical Configuration of Wheeled Robotic Locomotion*. PhD thesis, Robotics Institute, Carnegie Mellon University, Pittsburgh, PA.

[Automated Mining Systems, 2006] Automated Mining Systems (2006). see http://www.robominer.com/guidance.html.

[Bailey, 2003] Bailey, D. (2003). Mystery and joy as miners emerge. BBC News Online.

[Baker et al., 2004] Baker, C., Morris, A. C., Ferguson, D., Thayer, S., Whittaker, C., Omohundro, Z., Reverte, C., Whittaker, W. R. L., Haehnel, D., and Thrun, S. (2004). A campaign in autonomous mine mapping. In *Proceedings of the IEEE Conference on Robotics and Automation (ICRA)*.

[Banta et al., 1995] Banta, J., Y.Zhein, X.Z.Wang, G.Zhang, M.T.Smith, and M.A.Abidi (1995). A "next-best-view" algorithm for three-dimensional scene reconstruction using range images. In *SPIE*, volume 2588, pages 418–29.

[Champeny-Bares et al., 1991] Champeny-Bares, Coppersmith, L. S., and Dowling, K. (1991). The terregator mobile robot. Technical Report CMU-RI-TR-93-03, Robotics Institute, Carnegie Mellon University, Pittsburgh, PA.

[C.I.Conolly, 1985] C.I.Conolly (1985). The determination of next best views. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 432–435.

[CNN,] CNN. Chine mine blast kills 203.

[Corke et al., 1998] Corke, P. I., Roberts, J. M., and Winstanley, G. J. (1998). Robotics for the mining industry. *Autonomous Robotics Systems*, 236:163–181.

[Council, 2002] Council, N. R. (2002). *Coal Waste Impoundments: Risks, Responses, and Alternatives*, chapter 5. National Academy Press, Washington, D.C.

[DeKok, 1986] DeKok, D. (1986). *Unseen Danger: A Tragedy of People, Government, and the Centralia Mine Fire*. University of Pennsylvania Press.

[Directions Magazine , 2003] Directions Magazine (2003). see http://www.directionsmag.com/article.php?article id=407.

[Drenner et al., 2002] Drenner, A., Burt, I., Kratochvil, B., McMillen, C. P., Nelson, B., Papanikolopoulos, N., Rybski, P., Stubbs, K., Waletzko, D., and Yesin, K. B. (2002). Mobility enhancements to the scout robot platform. In *Proceedings of the 2002 IEEE International Conference on Robotics and Automation*, pages 1069–1074, Washington, DC.

[Duff and Roberts, 2003] Duff, E. and Roberts, J. (2003). Wall-Following with Constrained Active Contours. In *Proceedings of the Fourth International Conference on Field and Service Robotics*.

[Duff et al., 2003] Duff, E., Roberts, J., and Corke, P. (2003). Automation of an Underground Mining Vehicle using Reactive Navigation and Opportunistic Localization. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 3775–3780.

[Eisen et al., 1997] Eisen, H., C.W.Buck, G.R.Gillis-Smith, and J.W.Umland (1997). Mechanical design of the mars pathfinder mission. Technical report, NASA Jet Propulsion Laboratory.

[Elphinstone, 2006] Elphinstone, C. (2006). Minegem loader automation system. Product Brochure.

[Fotta et al.,] Fotta, B., Peters, R., and Mallett, L. Safety challenges at thin seam mines. NIOSH Online Publication.

[Gonzalez-Bantos and J.C.Latombe, 2002] Gonzalez-Bantos, H. and J.C.Latombe (2002). Navigation strategies for exploring indoor environments. *International Hournal of Robotics Research*, 21:829–848.

[Hainsworth et al., 1990] Hainsworth, D. W., Mallett, C. W., and Stacey, M. R. (1990). Project numbat - an emergency mine survey vehicle. In *Proce. Third National Conference on Robotics*, pages 91–98, Melbourne, Australia.

[H.H.Gonzalez-Banos and J.C.Latombe, 2001] H.H.Gonzalez-Banos and J.C.Latombe (2001). A randomized art-gallery algorithm for sensor placement. In *Proceedings 17th ACM Symposium on Computational Geometry (SoCG'01)*, pages 232–240.

[J.Maver and R.Bajcsy, 1993] J.Maver and R.Bajcsy (1993). Occlusions as a guide for planning the next view. In *IEEE Trans. Pattern Analysis and Machine Intelligence*, volume 15(5), pages 417–433.

[Jones, 2001] Jones, J. (2001). Inflatable robotics for planetary applications. In *6th International Symposium on Artifical Intelligence, Robotics, and Automation in Space, I-SAIRAS*, Montreal, Canada.

[Kuipers and Byan, 1991] Kuipers, B. and Byan, Y. (1991). A robot exploration and mapping strategy based on a semantic hierarchy of spatial representations. *J. Robot. Auton. Syst.*, 8:47–63.

[Leger, 1999] Leger, P. C. (1999). *Automated Synthesis and Optimization of Robot Configurations: An Evolutionary Approach*. PhD thesis, Robotics Institute, Carnegie Mellon University, Pittsburgh, PA.

[McAteer et al., 2006] McAteer, J. D., Bethell, T. N., Monforton, C., Pavlovich, J. W., Roberts, D., and Spence, B. (2006). The sago mine disaster: A preliminary report to governor joe manchin iii. Technical report, Wheeling Jesuit University, Buckhannon, West Virginia.

[PADEP, 2002] PADEP (2002). Report of commission on abandonded mine voids and mine safety. Technical report, Pennsylvania Department of Environmental Protection.

[RedZone Robotics, 2006] RedZone Robotics (2006). see http://www.redzone.com.

[Roberts et al., 2002] Roberts, J., Duff, E., and Corke, P. (2002). Reactive navigation and opportunistic localization for autonomous underground mining vehicles. *The International Journal of Information Sciences*.

[Roberts et al., 2000] Roberts, J., Duff, E., Corke, P., Sikka, P., Winstanley, G., and Cunningham, J. (2000). Autonomous Control of Underground Mining Vehicles using Reactive Navigation. In *Proceedings of the IEEE International Conference on Robotics and Automation*, pages 3790–3795.

[Roberts et al., 1999] Roberts, J. M., Corke, P. I., and Winstanley, G. J. (1999). Development of a 3500-tonne field robot. *The International Journal of Robotics Research*, 18(7):739–754.

[Ruegsegger and Lefchik, 1999] Ruegsegger, L. R. and Lefchik, T. E. (1999). Managing car-crunching sinkholes. *Public Roads*, 63(1).

[Scheding et al., 1997] Scheding, S., Dissanayake, G., Nebot, E., and Durrant-Whyte, H. (1997). Slip modelling and aided inertial navigation of an LHD. In *Proceedings of the IEEE International Conference on Robotics and Automation*, pages 1904–1909.

[Schempf, 1995] Schempf, H. (1995). Houdini: Site and locomotion analysis-driven design of an in-tank mobile cleanup robot. In *American Nuclear Society Winter Meeting Transactions*.

[Schempf et al., 1995] Schempf, H., Chemel, B., and Everett, N. (1995). Neptune: Above ground storage-tank inspection robot system. *IEEE Robotics and Automation Society Magazine*.

[Shaffer, 1995] Shaffer, G. (1995). *Two-Dimensional Mapping of Expansive Unknown Areas*. PhD thesis, Robotics Institute, Carnegie Mellon University, Pittsburgh,PA.

[Shammas et al., 2003] Shammas, E., Wolf, A., Brown, H. B., and Choset, H. (2003). New joint design for three-dimensional hyper redundant robots. In *Proceedings of the 2003 IEEE/RSJ International Conference on Intelligent Robots and Systems. (IROS '03)*, volume 4, pages 3594 – 3599.

[Stentz et al., 1999] Stentz, A., Ollis, M., Scheding, S., Herman, H., Fromme, C., Pedersen, J., Hegadorn, T., McCall, R., Bares, J., and Moore, R. (1999). Position Measurement for Automated Mining Machinery. In *Proceedings of the International Conference on Field and Service Robotics*, pages 299–304.

[The RATLER Project, 2006] The RATLER Project (2006). see http://www.ri.cmu.edu/projects/project_361.html.