A System for Three-Dimensional Robotic Mapping of Underground Mines

Michael Montemerlo, Dirk Hähnel, David Ferguson,
Rudolph Triebel, Wolfram Burgard, Scott Thayer
William Whittaker, and Sebastian Thrun

Oct 27, 2002

CMU-CS-02-185

School of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213

This research is sponsored by by DARPA's MARS Program (Contract number N66001-01-C-6018) and the National Science Foundation (CAREER grant number IIS-9876136 and regular grant number IIS-9877033), all of which is gratefully acknowledged. The views and conclusions contained in this document are those of the author and should not be interpreted as necessarily representing official policies or endorsements, either expressed or implied, of the United States Government or any of the sponsoring institutions.

Abstract

We describe two robotic systems [6] for acquiring high-resolution volumetric maps of underground mines. Our systems have been deployed in an operational coal mine in Bruceton, Pennsylvania, where they have been used to generate interactive 3-D maps. Our approach includes a novel sensor head, assembled from multiple SICK laser range finders, and a real-time algorithm for scan matching that generates accurate volumetric maps. The scan matching algorithm performs horizontal and vertical simultaneous localization and mapping (SLAM). Data from the horizontal scans is used to remove artifacts in the vertical scans, and vice versa. The system can construct full 3-D volumetric maps hundreds of meters in diameter, even when no odometry information is available.

Keywords: Robot mapping, mine mapping, mobile robotics, probabilistic robotics

1 Introduction

Throughout the industrialized world, the lack of accurate maps of inactive, underground mines poses a serious threat to public safety. According to a recent article [1], "Tens of thousands, perhaps even hundreds of thousands, of abandoned mines exist today in the United States. Not even the U.S. Bureau of Mines knows the exact number, because federal recording of mining claims was not required until 1976."[1] In July of 2002, nine miners were nearly killed in the QueCreek Mine in Somerset, Pennsylvania when they accidentally drilled into the abondoned Saxmon Mine, releasing millions of gallons of water in the QueCreek mine. This accident highlights the pressing need for accurate maps of abandoned mines.

Hazardous operating conditions and difficult access routes suggest that robotic exploration and mapping of abondoned mines may be necessary. Robotic mine mapping has been pursued by various research groups around the world. Corke and colleagues [3] have built vehicles that can acquire and utilize accurate 2-D maps of flat mines. Similarly, Baily [14] reports 2-D mapping results of an underground area using advanced mapping techniques. The mine mapping problem is made challenging by the lack of global position information underground. As a result, mine mapping must be approached as a *simultaneous localization and mapping*, or SLAM, problem [4, 8, 13, 16]. The robot must construct a map of the mine, while estimating its own position at the same time. The SLAM problem is known to be particularly difficult when the environment possesses loops [5, 15]. Unfortunately, mines typically contain a large number of cycles, and we know of no robotic system that could handle such maps. Moreover, none of the existing robotic mine mapping systems produce accurate volumetric 3-D maps.

The systems described in this paper are capable of generating volumetric 3-D models of mines. Our first system makes the common (but unrealistic) assumption of a flat floor inside the mine. This system has been used to generate accurate volumetric maps of relatively flat mines. The second, more elaborate system does not rely on a flat world assumption. It uses multiple range finders to generate

[1] See the course page http://www.cs.cmu.edu/afs/cs.cmu.edu/academic/ class/16861-f02/www/ for more information.

accurate volumetric maps for mines that change elevation. At the core of both systems are 2-D laser range finders, which are used for position referencing and for the recovery of the volumetric structure of a mine. Our initial system used two such sensors, one for each of the two functions described above. To accommodate uneven terrain, our second system uses four sensors, effectively extending the mine mapping capabilities of our first system into a vertical dimension. We present in this paper a new scan matching algorithm that exploits the overlapping laser range scans, to correct for noise and alignment errors in the data. The resulting mine maps are highly accurate 3-D models that can be visualized interactively by mining staff.

2 The Systems

Figures 1 and 2 show our two volumetric mine mapping systems. Our first prototype, shown in Figure 1, consists of a modified Pioneer AT robot. It is equipped with two SICK laser range finders, one pointing forward parallel to the floor, and one pointing upward perpendicular to the robot's heading direction. In addition, the robot is equipped with two wheel encoders to measure approximate robot motion. The forward-pointing laser scanner is used for simultaneous localization and mapping (SLAM) in 2-D. Using this data, the robot acquires an accurate 2-D map of the environment. The upward-pointing laser is used to reconstruct the 3-D shape of the walls and the ceiling of the mine, registered in space according to position estimates gathered from the 2-D map.

The limitations of the robotic system are immediately apparent. First and foremost, the system is confined to flat surfaces, due to its inability to sense or incorporate variations in elevation while performing SLAM. In this way, the system bears close resemblance to existing work on volumetric mapping of indoor environments [11, 7, 9], which principally lacks an extension into the third, vertical dimension when performing SLAM. Additionally, the robot platform was not rugged enough to handle the uneven, frequently wet terrain common in mines. Most notably, the robot was not able to cross rail-road tracks used to transport ore inside the mine.

To overcome these limitations, we developed the sensor cart assembly shown

Figure 1: Mine mapping robot with two laser range finders.

in Figure 2. This system is equipped with four SICK laser range finders. Two of these sensors point forward, but with a ninety degree offset in orientation. With this configuration, SLAM can be performed horizontally and vertically, capturing the missing dimension in the SLAM process. The other two lasers are mounted perpendicular to the motion direction of the cart, one pointing up (as on our robot), and one pointing down to map the texture of the floor and the lower portions of the wall. The four lasers together can acquire a full 3-D map of the mine, even under uneven terrain. Unfortunately, our cart is not equipped with odometry sensors. Hence highly accurate scan matching is essential in order to acquire large mine maps. The cart is pulled manually through a mine during mapping.

3 Software

Our approach is based on previous work on building large-scale 2-D maps of cyclic environments [17]. As such, it builds on a large body of literature on scan matching [5, 10] and probabilistic SLAM [4, 16]. However, the use of our new sensor assembly makes it possible to combine two processes of scan matching—one in the vertical dimension and one in the horizontal dimension—which is a key capability necessary to build maps of the scale and accuracy presented in this paper. All software described in this section (with the exception of the off-

Figure 2: Mine mapping cart with four laser range finders, for our new 4x2-D volumetric mapping approach.

the-shelf VRML viewer) is incremental and is executed in real-time, on laptop computers.

3.1 2-D SLAM

The robotic system shown in Figure 1 uses an improved version of the scan matching algorithm described in [17] for performing simultaneous localization and mapping (SLAM) in two dimensions. In essence, the problem is one of determining the shape of the environment from local sensors and (in the case of our robot) odometry data, while the same some maintaining an estimate of the robot's relative location and orientation in its ever-growing map. Our approach relies on scan matching as the basic mechanism for aligning scans. In doing so, it can eliminate the odometric error between subsequent scans almost entirely. Error that remains is due to multiple factors, such as the effect of uneven flooring and the noise in the sensor measurements.

Our approach deviates from our previous work in [17] in the way we perform the scan matching. Instead of matching scan points directly, our approach generates a local map out of a set of recent scans. Such maps are usually quite accurate, because they are constructed using scans that were nearly aligned already. To cope with errors and discontinuities in the maps stemming from residual errors in the

Figure 3: Two-dimensional scan alignment. Map created out of the most recent scans (left image), range scan measurement (center image), and resulting alignment (right image).

range registration process and discretization errors, we convolve this map with a Gaussian kernel. One advantage of this technique is that the every incoming scan is compared to a consistent local map, which reduces potential errors introduced by occluded portions of the environment or areas which have not been scanned due to the limited angular resolution of the scanner. The standard point-matching approach can diverge in cases where there is no odometry present (as is the case for our cart): in such situations, the match successively increases the distance between these scans, as an artifact of the exact spacing of the points in the scan. Second, the results of our approach are an order of magnitude more accurate. The scan alignment makes it possible to traverse hundreds of meters while maintaining an overall error in the centimeter range. Such accuracies were impossible to achieve using our previous software, and they are a direct result of our improved scan matching representation. Figure 3 shows a typical application of the range registration. The left image depicts the reference map constructed from 50 scans. The center image contains the scan that is aligned with this map. The right image shows the final position of the scan after applying the range registration procedure.

3.2 2x2-D SLAM

The key innovation of this paper is the use of two forward-pointed laser scanners, for performing SLAM simultaneously in both the horizontal and the vertical direction. At first glance, one might be tempted to simply run two such processes in

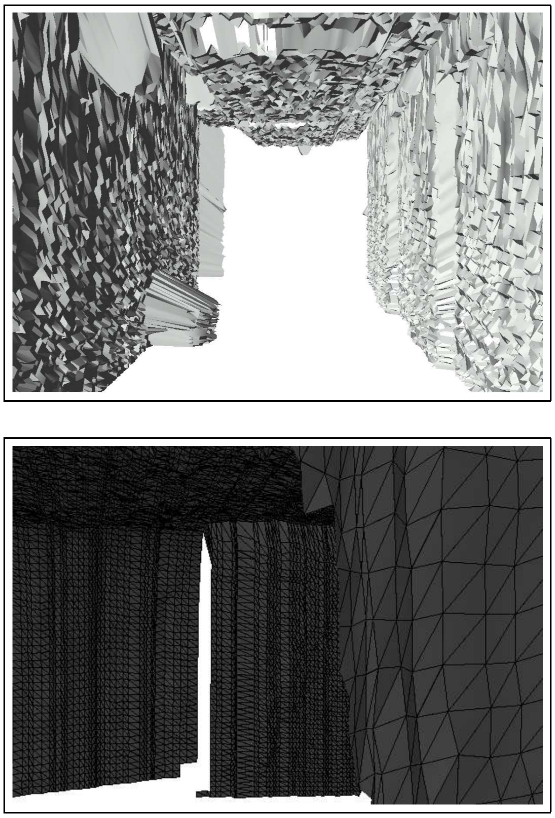

Figure 4: Surface visualization using past scan matching techniques (top), compared to a mine view using our present scan matching process (bottom). In 3-D, the difference of improved scan matching is much more apparent than in 2-D.

parallel, resulting in accurate 2-D cross-sectional maps of the mine that together allow for a recovery of the 3-D structure (under the obvious assumption that overall, the floor of the mine is not slanted sidewards). However, such a methodology is prone to fail in real mines.

The reason for such failure lies in the effect that variations in one dimension have on the measurements in the other. Consider, for example, a dip in the floor of the mine. This is clearly a vertical feature, and the vertical SLAM process can easily measure and map such a relief feature. However, as the cart is being moved through the mine, its horizontal sensor may see the ground, creating a 'phantom' obstacle in front of the robot. Phantom obstacles are usually fatal to

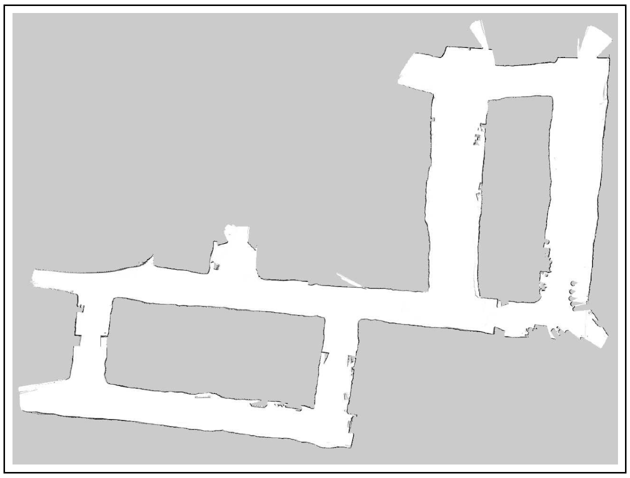

Figure 5: A 2-D map of the mine, acquired by our first system under a flat surface assumption. This 2D map is used for scan registration to localize the robot during mapping.

scan matching-based SLAM algorithms, a detection of the ground plane with the horizontal sensor appears as a massive obstacle in front of the robot. If these readings are used in the scan matching process, the robot will suffer unrecoverable errors in its longitudinal motion estimates. By symmetry, the same effect will corrupt the measurements of the vertical sensor. When the robot turns a corner, the vertical sensor will measure phantom objects that cannot be explained by a vertical view of the world alone; rather, these readings correspond to side walls that are being mapped by the horizontal sensor.

To accommodate for this interplay of horizontal and vertical structure, our system uses the vertical sensor to filter out phantom measurements in the horizontal sensor, and vice versa. In particular, we rely on the vertical sensor to detect when the horizontal sensor is close to detecting the ground plane. Our system uses the horizontal sensor to detect phantom objects in the vertical scans, which occur when the system turns or the mine is not straight, and the laser hits a side wall. This procedure automatically removes artifacts from the sensor measurements that result from the fact that the system is operating in an environment with non-trivial horizontal and vertical structure. Empirically, we found this approach to be necessary for the success of both SLAM components in mines with uneven surface properties. Since at the core, SLAM is still performed at a 2-D level (and

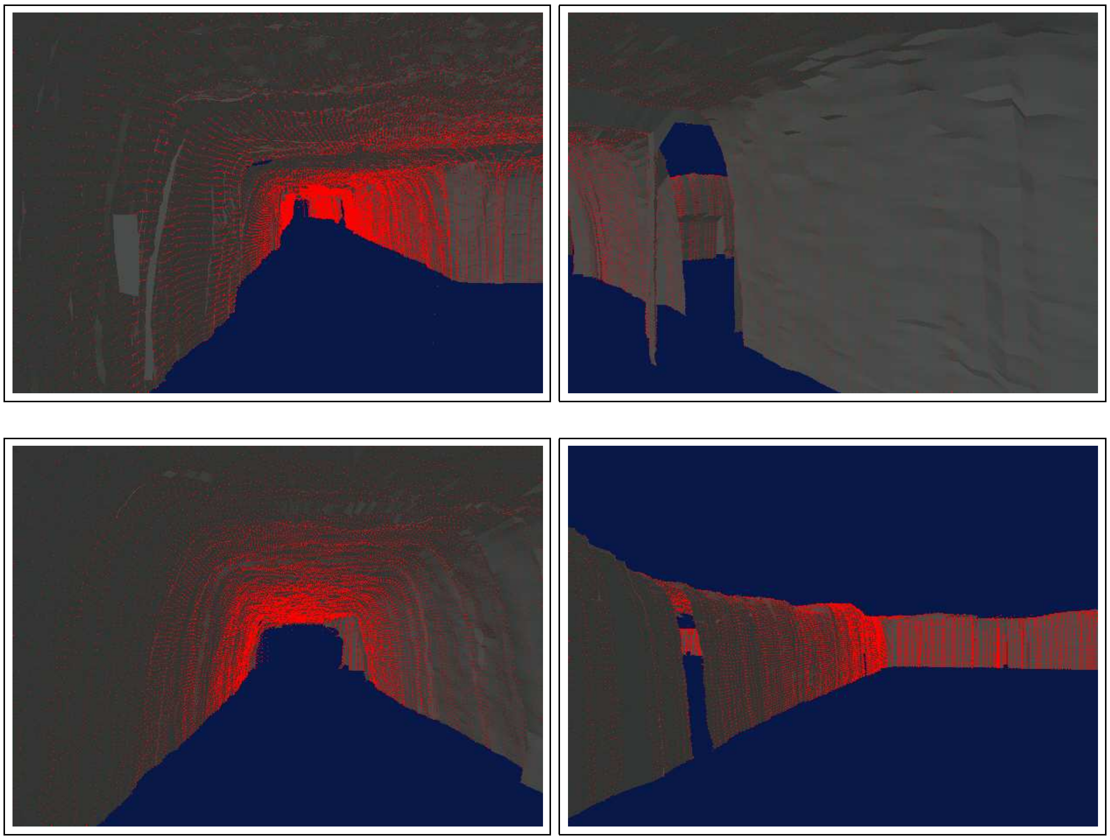

Figure 6: Sequence of 3-D visualizations of the planar surface volumetric mine map, acquired with the mobile robot. Shown in red are the sensor measurements used for generating the mine map.

not the full 3-D level due to the lack of full range cameras), we call the resulting approach 2x2-D SLAM.

3.3 3D Reconstruction

The 3-D volumetric reconstruction is achieved by using the remaining sensors, pointed upwards and (in the case of the cart) downwards in a direction perpendicular to the robot's heading direction. The reconstruction relies completely on the accuracy of localization during SLAM: Here our improved scan matching algorithm has a tremendous effect on the visual accuracy and integrity of the resulting maps when compared to our previous work. In particular, Figure 4 shows a cross-section of the raw 3-D data obtained by straightforward interpolation between adjacent sensor scans, and compares it with previous results obtained with our older scan matching approach [17]. This improved visual accuracy is partially a function of our improved scan matching.

In addition to that, we employ a local smoothing operator that further smooths the surface. Similar smoothing techniques were applied in [17].

4 Results

All results have been obtained in two different setions of an experimental coal mine in Bruceton, Pennsylvania. This mine is operated as a research mine by the U.S. Bureau of Mines, enabling us to operate robotic equipment without the need for explosion-proof certification. A partial map of the mine is shown in Figure 7 (bottom panel).

Figure 5 shows the result of 2-D mapping using our robotic system, of a small fraction of the mine with a sufficiently flat floor. This section of the mine had a concrete floor, facilitating its use as a research mine. However, concrete flooring is clearly unrepresentative of existing, and abandoned mines. As argued above, the flatness of the floor is essential for the success of our initial robotic system, which only performs SLAM in the horizontal direction.

3-D volumetric maps obtained using this system are shown in Figure 6. This visualization shows only the upper fraction of the mine. The map is incomplete due to the use of a single sensor for volumetric mapping on our robot. Nevertheless, these results illustrate that under idealized conditions, our initial system is indeed capable of acquiring accurate mine maps. However, our system failed in more realistic setting, where uneven floors and other artifacts (tracks, mud, water) made it impossible to acquire accurate maps.

These limitations were overcome with our mapping cart. Figure 7 shows a 2-D projection of the horizontal SLAM process, using 2x2-D SLAM as described above. Also shown in this figure is a blue-print of the mine map for comparison. It is important to notice that the 2-D SLAM map has been constructed without the use of odometry: The position estimate is solely the result of our 2x2-D scan

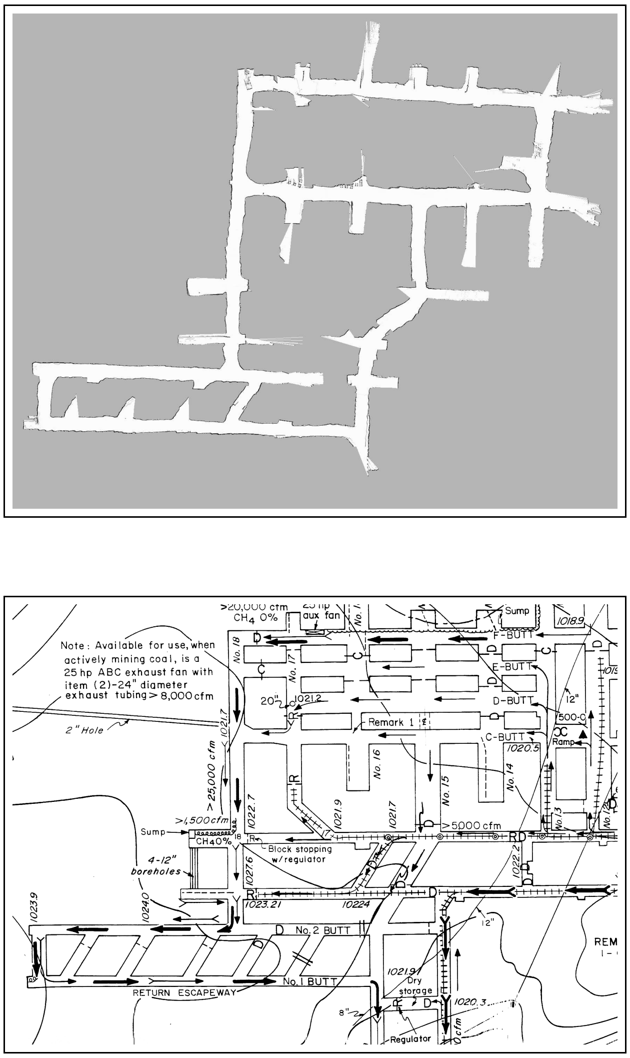

Figure 7: 2-D projection of the 3-D volumetric mine map, acquired by our mine mapping cart without odometry. A comparison with the manually constructed 2-D map illustrate the accuracy of our automatically acquired mine map.

matching approach. The largest loop in this map is several hundred meters in circumference, making this one of the largest loops in a confined environment ever mapped by probabilistic scan matching techniques. The lack of completeness of the map is due to closed doors and other massive obstacles that rendered many of the corridors in the mine inaccessible.

The resulting 3-D map is visualized in Figure 8. This map is represented in VRML format, allowing for an interactive visualization and exploration of the mapped mine. As these screen shots illustrate, the resulting map captures the full 3-D structure of the parts of the mines accessible to the cart. The resulting map captures both the horizontal and the vertical structure of the mine. The visualization tool enables mine personnel to inspect the mine from views that cannot be physically attained, such as the outside visualization shown in the bottom panel of that figure.

5 Conclusion

We have presented two implemented systems for acquiring volumetric maps of mines. Our first system relied on a robotic platform, equipped with two laser scanners. Our second, more versatile system used four range sensors, and was mounted on a cart. To achieve accurate mine mapping, we have developed a new scan matching algorithm that fuses information from a horizontal and a vertical sensor while performing SLAM in 2D. The volumetric map is then reconstructed from measurements acquired by additional laser sensors. As the results in this paper illustrate, our new scan matching approach enables us to obtain consistent volumetric maps of mines with significant vertical and horizontal structure. The fact that our final results were obtained in the absence of any odometry data illustrates the robustness of our approach.

We believe that the volumetric mine maps are unprecedented in the robotics literature in their scale, resolution, and by virtue of the fact that they are volumetric, and not just two-dimensional. The 2x2-D system is presently been extended to a rugged ATRV platform capable of traversing the type terrain found in mines, in a self-propelled mode. We anticipate that this will provide us with an automated robotic system for acquiring large maps of mines. We also believe that existing techniques for mobile robot exploration [2, 12, 18] can be adapted for the purpose of autonomously exploring mines.

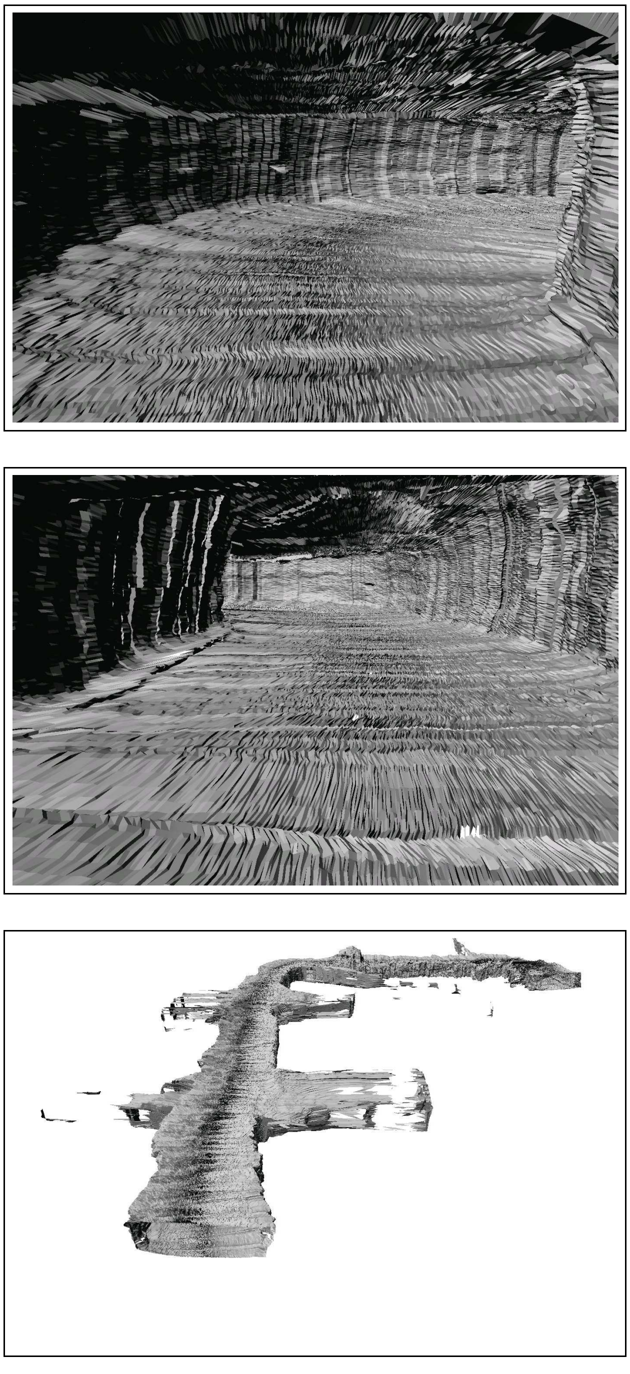

Figure 8: Sequence of 3-D visualizations of the full 3-D volumetric mine map. This map has been built using our new sensor cart and using our 2x2-D scan matching algorithm, without any odometry information.

Acknowledgment

This research was motivated and supported by an ongoing mine mapping course offered by Red Whittaker and Scott Thayer, whose encouragement is greatly acknowledged. The research has been sponsored by DARPA's MARS Program (contract N66001-01-C-6018 and contract NBCH1020014), which is gratefully acknowledged. Special thanks to the Bruceton Research Center for access to the Bruceton experimental mine.

References

[1] J.J. Belwood and R.J. Waugh. Bats and mines: Abandoned does not always mean empty. *Bats*, 9(3), 1991.

[2] W. Burgard, D. Fox, M. Moors, R. Simmons, and S. Thrun. Collaborative multi-robot exploration. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, San Francisco, CA, 2000. IEEE.

[3] P. Corke, J. Cunningham, D. Dekker, , and H. Durrant-Whyte. Autonomous underground vehicles. In *Proceedings of the CMTE Mining Technology Conference*, pages 16–22, Perth, Australia, September 1996.

[4] G. Dissanayake, H. Durrant-Whyte, and T. Bailey. A computationally efficient solution to the simultaneous localisation and map building (SLAM) problem. Working notes of ICRA'2000 Workshop W4: Mobile Robot Navigation and Mapping, April 2000.

[5] J.-S. Gutmann and K. Konolige. Incremental mapping of large cyclic environments. In *Proceedings of the IEEE International Symposium on Computational Intelligence in Robotics and Automation (CIRA)*, 2000.

[6] D. Hähnel, M. Montemerlo, D. Ferguson, R. Triebel, W. Burgard, and S. Thrun. A system for volumetric robotic mapping of underground mines. Submitted for publication, 2002.

[7] L. Iocchi, K. Konolige, and M. Bajracharya. Visually realistic mapping of a planar environment with stereo. In *Proceesings of the 2000 International Symposium on Experimental Robotics*, Waikiki, Hawaii, 2000.

[8] J. Leonard, J.D. Tardós, S. Thrun, and H. Choset, editors. *Workshop Notes of the ICRA Workshop on Concurrent Mapping and Localization for Autonomous Mobile Robots (W4)*. ICRA Conference, Washington, DC, 2002.

[9] Y. Liu, R. Emery, D. Chakrabarti, W. Burgard, and S. Thrun. Using EM to learn 3D models with mobile robots. In *Proceedings of the International Conference on Machine Learning (ICML)*, 2001.

[10] F. Lu and E. Milios. Globally consistent range scan alignment for environment mapping. *Autonomous Robots*, 4:333–349, 1997.

[11] H.P. Moravec and M.C. Martin. Robot navigation by 3D spatial evidence grids. Mobile Robot Laboratory, Robotics Institute, Carnegie Mellon University, 1994.

[12] R. Simmons, D. Apfelbaum, W. Burgard, M. Fox, D. an Moors, S. Thrun, and H. Younes. Coordination for multi-robot exploration and mapping. In *Proceedings of the AAAI National Conference on Artificial Intelligence*, Austin, TX, 2000. AAAI.

[13] R. Smith, M. Self, and P. Cheeseman. Estimating uncertain spatial relationships in robotics. In I.J. Cox and G.T. Wilfong, editors, *Autonomous Robot Vehicles*, pages 167–193. Springer-Verlag, 1990.

[14] Bailey T. *Mobile Robot Localisation and Mapping in Extensive Outdoor Environments*. PhD thesis, University of Sydney, Sydney, NSW, Australia, 2002.

[15] S. Thrun. A probabilistic online mapping algorithm for teams of mobile robots. *International Journal of Robotics Research*, 20(5):335–363, 2001.

[16] S. Thrun. Robotic mapping: A survey. In G. Lakemeyer and B. Nebel, editors, *Exploring Artificial Intelligence in the New Millenium*. Morgan Kaufmann, 2002. to appear.

[17] S. Thrun, W. Burgard, and D. Fox. A real-time algorithm for mobile robot mapping with applications to multi-robot and 3D mapping. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, San Francisco, CA, 2000. IEEE.

[18] B. Yamauchi, P. Langley, A.C. Schultz, J. Grefenstette, and W. Adams. Magellan: An integrated adaptive architecture for mobile robots. Technical Report 98-2, Institute for the Study of Learning and Expertise (ISLE), Palo Alto, CA, May 1998.