6D SLAM with an Application in Autonomous Mine Mapping

Andreas Nüchter, Hartmut Surmann,
Kai Lingemann, Joachim Hertzberg
Fraunhofer Institute for Autonomous Intelligent Systems (AIS)
Schloss Birlinghoven
D-53754 Sankt Augustin, Germany
{firstname.surname}@ais.fraunhofer.de

Sebastian Thrun
Computer Science Department
Stanford University
Stanford, CA, USA
thrun@stanford.edu

Abstract— To create with an autonomous mobile robot a 3D volumetric map of a scene it is necessary to gage several 3D scans and to merge them into one consistent 3D model. This paper provides a new solution to the simultaneous localization and mapping (SLAM) problem with six degrees of freedom. Robot motion on natural surfaces has to cope with yaw, pitch and roll angles, turning pose estimation into a problem in six mathematical dimensions. A fast variant of the Iterative Closest Points algorithm registers the 3D scans in a common coordinate system and relocalizes the robot. Finally, consistent 3D maps are generated using a global relaxation. The algorithms have been tested with 3D scans taken in the Mathies mine, Pittsburgh, PA. Abandoned mines pose significant problems to society, yet a large fraction of them lack accurate 3D maps.

I. INTRODUCTION

The increasing need for rapid characterization and quantification of complex environments has created challenges for data analysis. This critical need comes from many important areas, including industrial automation, architecture, agriculture, and the construction or maintenance of tunnels and mines. Especially mobile systems with 3D laser scanners that automatically perform multiple steps such as scanning, gaging and autonomous driving have the potential to greatly advance the field of environment sensing. On the other hand, having 3D information available in real-time enables autonomous robots to navigate in unknown environments, e.g., in the field of autonomous mine inspection.

The problem of automatic environment sensing and modeling is complex, because a number of fundamental scientific issues are involved in this research. One issue is the control of an autonomous mobile robot and scanning the environment with a 3D sensor. Another question is how to create a volumetric consistent scene in a common coordinate system from multiple views. The latter problem is addressed here: The proposed algorithms allows to digitize large environments fast and reliably without any intervention and solve the simultaneous localization and mapping (SLAM) problem. Finally robot motion on natural outdoor surfaces has to cope with yaw, pitch and roll angles, turning pose estimation as well as scan matching or registration into a problem in six mathematical dimensions. This paper presents a new solution to the SLAM problem with six degrees of freedom. A fast variant of the iterative closest points (ICP) algorithm registers the 3D scans in a common coordinate system and relocalizes the robot. The computational requirements are reduced by two new methods:

First we reduce the 3D data, i.e., we compute depth images that approximate the scanned 3D surface and contain only a small fraction of the 3D point cloud. Second a fast approximation of the closest point for the ICP algorithm is given. These extenstions of the ICP result in an algorithm for generating overall consistent 3D maps using global error minimization.

This paper describes an algorithm for acquiring volumetric maps of underground mines. Mapping underground mines is of enormous societal importance [16], as a lack of accurate maps of inactive, underground mines poses a serious threat to public safety. According to a recent article [3], "tens of thousands, perhaps even hundreds of thousands, of abandoned mines exist today in the United States and worldwide. Not even the U.S. Bureau of Mines knows the exact number, because federal recording of mining claims was not required until 1976." The lack of accurate mine maps frequently causes accidents, such as a recent near-fatal accident in Quecreek, PA [18]. Even when accurate maps exist, they provide information only in 2D, which is usually insufficient to assess the structural soundness of abandoned mines. Accurate 3D models of such abandoned mines would be of great relevance to a number of problems that directly affect the people who live or work near them. One is subsidence: structural shifts can cause collapse on the surface above. Ground water contamination is another problem of great importance, and knowing the location, volume, and condition of an abandoned mine can be highly informative in planning and performing interventions. Accurate volumetric maps are also of great commercial interest. Knowing the volume of the material already removed from a mine is of critical interest when assessing the profitability of re-mining a previously mined mine.

Hazardous operating conditions and difficult access routes suggest that robotic mapping of abandoned mines may be a viable option to traditional manual mapping techniques. The idea of mapping mines with robots is not new. Past research has predominantly focused on acquiring maps for autonomous robot navigation in active mines. For example, Corke and colleagues [8] have built vehicles that acquire and utilize accurate 2D maps of mines. Similarly, Baily [2] reports 2D mapping results of an underground area using advanced mapping techniques. The experiments reported in this paper utilize data by CMU's mine mapping robot Groundhog [10], [24], which relies on 2D mapping for exploring and mapping of abandoned mines. While Groundhog and a related bore-hole

deployable device "Ferret" [16] utilizes local 3D scans for navigation and terrain assessment, none of these techniques integrates multiple 3D scans and generates full volumetric maps of abandoned mines.

The paper is organized as follows: The next two sections describe the state of the art in automatic 3D mapping and present the autonomous mobile robot. Section IV presents the registration algorithms and the solution to the SLAM problem. Furthermore we show how data reduction and access methods speed up computation and the methods become real-time capable. Section V describes the Mathies Mine experiment. Finally, section VI summerizes the results and concludes the paper. The paper is accompanied with a 3D map, given as video. The video is available for download at www.ais.fraunhofer.de/ARC/3D/mine/.

II. STATE OF THE ART

Some groups have attempted to build 3D volumetric representations of environments with 2D laser range finders. Thrun et al. [23] use two 2D laser range finder for acquiring 3D data. One laser scanner is mounted horizontally and one is mounted vertically. The latter one grabs a vertical scan line which is transformed into 3D points using the current robot pose. The horizontal scanner is used to compute the robot pose. The precision of 3D data points depends on that pose and on the precision of the scanner. All these approaches have difficulties to navigate around 3D obstacles with jutting out edges. They are only detected while passing them.

The published 2D probabilistic localization approaches, e.g., Markov models or Monte Carlo algorithms work well in flat and structured 2D environments but an extension in the third dimension is still missing since the algorithm do not scale with additional dimensions. Common approaches to SLAM use global relaxation after incremental 2D scan matching in order to create a globally consistent map [11].

A few other groups use 3D laser scanners [1], [14], [22]. A 3D laser scanner generates consistent 3D data points within a single 3D scan. The AVENUE project develops a robot for modeling urban environments using a CYRAX laser scanner and a feature-based scan matching approach for registration of the 3D scans in a common coordinate system [1]. The research group of M. Hebert has reconstructed environments using the Zoller+Fröhlich laser scanner and aims at building 3D models without initial position estimates, i.e., without odometry information [14].

III. THE GROUNDHOG ROBOT

The robot has been deployed and built by the CMU Mine Mapping Team [10], [24]. Groundhog's chassis unites the front halves of two all terrain vehicles, allowing all four of Groundhog's wheels to be both driven and steered. The two Ackerman steering columns are linked in opposition, reducing Groundhog's outside turning radius to approximately 2.44m. A hydraulic cylinder drives the steering linkage, with potentiometer feedback providing closed-loop control of wheel angle. Two hydraulic motors coupled into the front and rear

Fig. 1.   The Groundhog robot

stock ATV differentials via 3:1 chain drives result in a constant 0.145 m/sec velocity. When in motion, Ground hog consumes upwards of 1kW, where processing and sensing only draw 25W and 75W respectively. Therefore, time spent sensing and processing has minimal impact on the operational range of the robot. The high power throughput combined with the low speed of the robot means that Groundhog has the torque necessary to overcome the railway tracks, fallen timbers, and other rubble commonly found in abandoned mines. Equipped with six deep-cycle lead-acid batteries, and in later experiments with eight such battereis, Groundhog has a locomotive range greater than 3km.

For acquiring 3D scans in a stop and go fashion, Groundhog is equipped with tiltable SICK laser range finders on either end. The area of $180°(h) \times 60°(v)$ is scanned with the horizontal resolution of 361 pts. and vertical of 341 pts. A slice with 361 data points is scanned in 26ms by the 2D laser range finder (rotating mirror device). Thus a scan with 361×341 data points needs 8.9 seconds. Fig. 4 shows an example scan of the mine.

IV. RANGE IMAGE REGISTRATION AND ROBOT RELOCALIZATION

Multiple 3D scans are necessary to digitalize environments without occlusions. To create a correct and consistent model, the scans have to be merged into one coordinate system. This process is called registration. If the localization of the robot with the 3D scanner were precise, the registration could be done directly by the robot pose. However, due to the unprecise robot sensors, the self localization is erroneous, so the geometric structure of overlapping 3D scans has to be considered for registration.

The matching of 3D scans can either operate on the whole 3D scan point set or can be reduced to the problem of scan matching in 2D by extracting, e.g., a horizontal plane of fixed height from both scans, merging these 2D scans and applying the resulting translation and rotation matrix to all points of the corresponding 3D scan.

Matching of complete 3D scans has the advantage of having a larger set of attributes (either pure data points or extracted features) to compare the scans. This results in higher precision and lowers the possibility of running into a local minimum of

the cost function. Furthermore, using three dimensions enables the robot control software to recognize and take into account changes of height and roll, yaw and pitch angles of the robot. This 6D robot relocalization is essential for robots driving cross country or in mines.

6D matching approaches of 3D surfaces can be classified into two categories: First, scan matching as optimization problem uses a cost function for the quality of the alignment of the scans. The range images are registered by determining the rigid transformation (rotation and translation) which minimizes the cost function. Second, feature based scan matching extracts distinguishing features of the range images and uses corresponding features for calculating the alignment the scans. Even though through this approach is more intuitive, it cannot be applied to scan matching in mines, since the surface structure of the mine is too simple. In consequence there are not many features and an algorithm based on feature matching will fail [22], [17].

A. Matching as Optimization

The following method for registration of point sets is part of many publications, so only a short summary is given here. The complete algorithm was invented in 1992 and can be found, e.g., in [5]. The method is called *Iterative Closest Points (ICP) algorithm*.

Given two independently acquired sets of 3D points, M (model set, $|M| = N_m$) and D (data set, $|D| = N_d$) which correspond to a single shape, we want to find the transformation consisting of a rotation $\mathbf{R}$ and a translation $\mathbf{t}$ which minimizes the following cost function:

$$E(\mathbf{R}, \mathbf{t}) = \sum_{i=1}^{N_m} \sum_{j=1}^{N_d} w_{i,j} \left\| \mathbf{m}_i - (\mathbf{R}\mathbf{d}_j + \mathbf{t}) \right\|^2. \quad (1)$$

$w_{i,j}$ is assigned 1 if the i-th point of M describes the same point in space as the j-th point of D. Otherwise $w_{i,j}$ is 0. Two things have to be calculated: First, the corresponding points, and second, the transformation $(\mathbf{R}, \mathbf{t})$ that minimize $E(\mathbf{R}, \mathbf{t})$ on the base of the corresponding points.

The ICP algorithm calculates iteratively a local minimum of equation (1). In each iteration step, the algorithm selects the closest points as correspondences $w_{i,j}$ and calculates the transformation $(\mathbf{R}, \mathbf{t})$ for minimizing equation (1). Fig. 2 shows three steps of the ICP algorithm. Besl and McKay prove that the method terminates in a minimum [5]. The assumption is that in the last iteration step the point correspondences are correct.

In each ICP iteration, the transformation is calculated by the quaternion based method of Horn [15]: A unit quaternion is a 4 vector $\dot{q} = (q_0, q_x, q_y, q_z)^T$, where $q_0^2 + q_x^2 + q_y^2 + q_z^2 = 1, q_0 \geq 0$. It describes a rotation axis and an angle to rotate around that axis. A 3×3 rotation matrix $\mathbf{R}$ is calculated from the unit quaternion according the the following scheme: $\mathbf{R} =$

$$\begin{pmatrix} (q_0^2 + q_x^2 - q_y^2 - q_z^2) & 2(q_x q_y + q_z q_0) & 2(q_x q_z + q_y q_0) \\ 2(q_x q_y + q_z q_0) & (q_0^2 - q_x^2 + q_y^2 - q_z^2) & 2(q_y q_z - q_x q_0) \\ 2(q_z q_x - q_y q_0) & 2(q_z q_y + q_x q_0) & (q_0^2 - q_x^2 - q_y^2 + q_z^2) \end{pmatrix}.$$


Fig. 2. Left: Initial odometry based pose of two 3D scans. Middle: Pose after five ICP iterations. Right: final alignment.

To determine the transformation, the mean values of the paired points (centroid vectors) $\mathbf{c}_m$ and $\mathbf{c}_d$ are subtracted from all points in M and D, respectively, resulting in the sets M' and D'. The rotation expressed as quaternion that minimizes equation (1) is the largest eigenvalue of the cross-covariance matrix

$$\mathbf{N} = \begin{pmatrix} (S_{xx} + S_{yy} + S_{zz}) & (S_{yz} + S_{zy}) \\ (S_{yz} + S_{zy}) & (S_{xx} - S_{yy} - S_{zz}) \\ (S_{zx} + S_{xz}) & (S_{xy} + S_{yx}) \\ (S_{xy} + S_{yx}) & (S_{yz} + S_{zy}) \\ & \\ (S_{zx} + S_{xz}) & (S_{xy} + S_{yx}) \\ (S_{xy} + S_{yx}) & (S_{zx} + S_{xz}) \\ (-S_{xx} + S_{yy} - S_{zz}) & (S_{yz} + S_{zy}) \\ (S_{zx} + S_{xz}) & (-S_{xx} - S_{yy} + S_{zz}) \end{pmatrix},$$

with $S_{xx} = \sum_{i=1}^{N_m} \sum_{j=1}^{N_d} w_{i,j} \ m'_{ix} d'_{jx}$, $S_{xy} = \sum_{i=1}^{N_m} \sum_{j=1}^{N_d} w_{i,j} \ m'_{ix} d'_{jy}$, ... After calculation the rotation $\mathbf{R}$, the translation is determined by $\mathbf{t} = \mathbf{c}_m - \mathbf{R}\mathbf{c}_d$ [15]. Fig. 2 shows two 3D scans in their initial, i.e., odometry-based pose, after 5 iterations, and the final pose. 40 iterations are needed to align these two 3D scans correctly.

B. Matching Multiple 3D Scans

To digitalize environments, multiple 3D scans have to be registered. After registration, the scene has to be globally consistent. A straightforward method for aligning several 3D scans is *pairwise matching*, i.e., the new scan is registered against the scan with the largest overlapping areas. The latter one is determined in a preprocessing step. Alternatively, Chen and Medioni [7] introduced an *incremental matching* method, i.e., the new scan is registered against a so-called *metascan*, which is the union of the previously acquired and registered scans. Each scan matching has a limited precision. Both methods accumulate the registration errors such that the registration of many scans leads to inconsistent scenes and to problems with the robot localization.

Pulli presents a registration method that minimizes the global error and avoids inconsistent scenes [19]. This method distributes the global error while the registration of one scan is followed by registration of all neighboring scans. Other matching approaches with global error minimization have been published, e.g., by Benjemaa et. al. [4] and Eggert et. al. [9].

Based on the idea of Pulli we have designed a method called *simultaneous matching* [17], [22]. Thereby, the first scan is the master scan and determines the coordinate system. This scan is fixed. The following steps register all scans and minimize the global error:

1) Based on the robot odometry, pairwise matching is used to find a start registration for a new scan. This step speeds up computation.
2) A queue is initialized with the new scan.
3) Three steps are repeated until the queue is empty:
   a) The current scan is the first scan of the queue. This scan is removed from the queue.
   b) If the current scan is not the master scan, then a set of neighbors (set of all scans that overlap with the current scan) is calculated. This set of neighbors forms one point set M. The current scan forms the data point set D and is aligned with the ICP algorithms.
   c) If the current scan changes its location by applying the transformation (translation or rotation), then each single scan of the set of neighbors that is not in the queue is added to the end of the queue.

Note: One scan overlaps with another iff more than 250 corresponding point pairs exist.

In contrast to Pulli's approach, the proposed method is totally automatic and no interactive pairwise alignment has to be done. Furthermore the point pairs are not fixed [19]. The accumulated alignment error is spread over the whole set of acquired 3D scans. An explicit detection of closed loops for the proposed solution to the SLAM problem is not necessary, multiple overlapping 3D scans are sufficient to diffuse the alignment error equally over the set of 3D scans.

C. Data Reduction

The computational expense of the ICP algorithm depends mainly on the number of points. In a brute force implementation the point pairing is in $O(n^2)$. Data reduction reduces the time required for matching. Several approaches have been presented for subsampling the data, including randomized sampling, uniform sampling, normal-space sampling and covariance sampling [20], [12]. Randomized sampling selects points at random, uniform sampling draws samples equally distributed samples from the input point cloud. Normal space sampling, as proposed by Rusinkiewicz and Levoy, aims at constraining translational sliding of input meshes, generated from the point cloud [20]. Their algorithm tries to ensure that the normals of the selected points uniformly populate the sphere of directions. Covariance sampling as proposed by Levoy et al. and extends the nomal space approach. They identify whether a pair of meshes will be unstable in the ICP algorithms by estimating a covariance matrix from a sparse uniform sampling of the input [12].

The data reduction proposed here considers the procedure of the scanning process, i.e., the spherical and continuous measurment of the laser. Scanning is noisy and small errors


Fig. 4. A point cloud of a 3D laser scan taken inside the Mathies Mine (perspective projection). Top left: Viewing pose 2.5 meter behind the scan pose. Top right: Top view. Bottom: Reduced and filtered point cloud. The reduced points have been enlarged. The number of points was reduced from 123101 to 10535.

may occur. Two kinds of errors mainly occur: Gaussian noise and so called salt and pepper noise. The latter one occurs for example at edges, where the laser beam of the scanner hits two surfaces, resulting in a mean and erroneous data value. Furthermore reflections lead to suspicious data. Without filtering, only a few outliers lead to multiple wrong point pairs during the matching phase and results in an incorrect 3D scan alignment.

We propose a fast filtering method to reduce and smooth the data for the ICP algorithm. The filter is applied to each 2D scan slice, containing 361 data points. It is a combination of a median and a reduction filter. The median filter removes the outliers (Fig. 3) by replacing a data point with the median value of the n surrounding points (here: $n = 7$). The neighbor points are determined according to their number within the 2D scan, since the laser scanner provides the data sorted in a counter-clockwise direction. The median value is calculated with regards to the Euclidian distance of the data points to the point of origin. In order to remove salt and pepper noise but leave the remaining data untouched, the filtering algorithm replaces a data point with the corresponding median value if and only if the difference (Euclidian distance) between both is larger than a fixed threshold (here: threshold = 200 cm). The data reduction works as follows: The scanner emits the laser beams in a spherical way, such that the data points close to the source are more dense. Multiple data points located close together are joined into one point. This reduction lowers the Gaussian noise. The number of these so called *reduced points* is in the mine application one order of magnitude smaller than the original one (Fig. 3). Finally the data points of a slice have a minimal distance of 10 cm and approximate the surface. The clue of the algorithm is that it is nearly impossible to detect differences between the median filtered and the reduced data (Fig. 3). The reduction fulfills the sampling criterions stated by Boulanger et al. [6], i.e., sampling the range images, such that the surface curvature is maintained.

The data for the scan matching is collected from every third scan slice. This fast vertical reduction yields a good surface description. Data reduction and filtering are online algorithms and run in parallel to the 3D scanning.

Fig. 3. Filtering the data for removing noise and reducing the computational expense for the ICP algorithm. Left: Original data. Middle: median filter. Right: median filtering combined with reduction. Notice: The data points have been connected with lines for better demonstration of effects from the median filter. The reduced and filtered data fits the scanned surface perfectly.

D. Data Access

The ICP algorithms spends most of its time in creating the point pairs. kD-trees (here $k = 3$) have been suggested to speed up the data access [21]. The kd-trees are a binary tree data structure with terminal buckets. The data is stored in the buckets and the keys are selected, such that a data space is divided into two equal parts. This ensures that a data point can be selected in $O(\log n)$. The proposed SLAM algorithm benefits from fast data access. In addition, the time spent on creating the tree is important.

For a given data set, larger bucket size results in smaller number of terminal buckets and hence less computational time to build the tree. The implemented algorithm uses a bucket size of 10 points and cuts recursively the scanned volume in two equal-sized halves. Once the tree is built, the nearest neighbor search for a given data point $\mathbf{p}$ starts at the root of the tree. Each tree node contains the cut dimension, i.e., orientation of the cut plane and the cut value. By comparing the coordinate of the given 3D point at the cut dimension with the cut value of this tree node, the search knows which branch to go for next level of tree node, it will compare the cut value of this tree node and go down further. This process is repeated until the terminal bucket that contains the closest data points $\mathbf{p}_b$ is reached. It may occur, that the true neighbor lies in a different bucket, e.g., if the distance between $\mathbf{p}$ and a boundary of its bucket region is less than the distance $\mathbf{p}$ and $\mathbf{p}_b$. In this case kd-tree algorithms have to backtrack, until all buckets that lie within the radius $||\mathbf{p} - \mathbf{p}_b||$ are explored. This is known as the Ball-Within-Bounds tests [13]. The number of distance computations is minimal for the smallest bucket size, i.e., one point per bucket. Nevertheless the running time of a kd-tree decreases for a slightly larger bucket size. This is due to the greater cost of backtracking as compared to a simple linear traversal of a small list within the bucket.

Recently, Greenspan and Yurick introduced Approximate kd-trees (Apr-kd-tree) [13]. The idea behind this is to return as an approximate nearest neighbor $\mathbf{p}_a$ the closest point $\mathbf{p}_b$ in the bucket region where $\mathbf{p}$ lies. This value is determined from the depth-first search, thus expensive Ball-Within-Bounds tests and backtracking are not used [13]. In addition to these ideas we avoid the linear search within the bucket. During the computation of the Apr-kd-tree, the mean values of the points within a bucket are computed and stored. Then the mean value of the bucket is used as approximate nearest neighbor,


Fig. 5. Running time of scan registration using kd-tree search and approximate kd-tree search with different bucket sizes.

replacing the linear search.

The search using the Apr-kd-tree is applied until the error function $E(\mathbf{R}, \mathbf{t})$ (1) stops decreasing. To avoid misalignments due to the approximation, the registration algorithm is restarted using the normal kd-tree search. A few iterations (usually 1-3) are needed for this final corrections. Fig. 5 shows the registration time of two 3D scans in dependence to the bucket size using kd-tree or Apr-kd-tree search. Both search methods have their best performance with a bucket size of 10 points.

V. The Mathies Mine Experiment

Groundhog's development began in the Fall of 2002 by the CMU Mine Mapping Team [10], [24]. The robot was extensively tested in a well-maintained inactive coal mine accessible to people: the Bruceton Research Mine located near Pittsburgh, PA. However, this mine is technically not abandoned and therefore not subject to collapse and deterioration. On May 30, 2003 Groundhog finally entered an inaccessible abandoned mine in fully autonomous mode. The mine is known as the Mathies mine and is located in the same geographic area as the other mines. The core of this surface-accessible mine consists of two 1.5-kilometer long corridors which branches into numerous side corridors, and which are accessible at both ends. This was an important feature of this mine, as it provided natural ventilation and thereby reduced the chances of encountering combustible gases inside the mine.

To acquire an accurate 3D map of one of the main corridors, the robot was programmed to autonomously navigate through the corridor. 250 meters into the mine, the robot encountered a broken ceiling bar draping diagonally across its path. The robot made the correct decision to retract. The data acquired on these runs has provided us with models of unprecedented detail and accuracy, of subterranean spaces that may forever remain off limits for people.

points used	time	# ICP iterations
all points & brute force search	4 h 25 min	45
all points & kD–tree	6.8 sec	45
all points & Apx-kD–tree	5.9 sec	45
reduced points & Apx-kD–tree	<0.62 sec	42
3D SLAM with		42 (step 1)
reduced points & Apx-kD–tree	52 sec	497 (step 3)

VI. RESULTS AND CONCLUSIONS

The algorithms have been applied to data collected in the Mathies Mine after the robot returned. Table I summarizes the results for the 3D scan matching and 6D SLAM. It is shown that using approximate kd-tree search decreases the running time of the proposed scan matching algorithms about 15%. Nevertheless, the main speedup is reached by the data reduction, resulting in a real-time capable ICP algorithm. The 6D SLAM algorithm can be used on an inspection robot for mines, the time needed for global consistent registration roughly corresponds to the time, needed to drive to the next scanning pose.

Fig. 6 shows the result of the Mathies Mine mapping. The top plot shows the 2D map, i.e., xz-map, where z is the depth axis. The bottom part shows the elevation, i.e., the xy-map. The Groundhog robot had to overbear a height of 4 meters during its 250 meter long autonomous drive.

To visualize the scanned 3D data, a viewer program based on OPENGL has been implemented. The task of this program is to project the 3D scene to the image plane, i.e., the monitor, such that the data can be drawn and inspected from every perspective. Fig. 2 and 4 show rendered 3D scans. A video of all matched 3D scans is available for download at www.ais.fraunhofer.de/ARC/3D/mine/. C This paper has presented a new solution to the simultaneous localization and mapping (SLAM) problem with six degrees of freedom. Based on the ICP algorithm the registration error is globally spread over all 3D scans and thus minimized. The presented algorithms are significant speeded up with data reduction that maintains the surface structure and with approximate kd-tree for closest point search.

REFERENCES

[1] P. Allen, I. Stamos, A. Gueorguiev, E. Gold, and P. Blaer. AVENUE: Automated Site Modeling in Urban Environments. In *Proc. IEEE 3DIM*, Canada, 2001.
[2] T. Bailey. *Mobile Robot Localisation and Mapping in Extensive Outdoor Environments*. PhD thesis, University of Sydney, Australia, 2002.
[3] J.J. Belwood and R.J. Waugh. Bats and mines: Abandoned does not always mean empty. *Bats*, 9(3), 1991.
[4] R. Benjemaa and F. Schmitt. Fast Global Registration of 3D Sampled Surfaces Using a Multi-Z-Buffer Technique. In *Proc. IEEE 3DIM*, Canada, 1997.
[5] P. Besl and N. McKay. A method for Registration of 3–D Shapes. *IEEE Transactions on PAMI*, 14(2):239 – 256, February 1992.
[6] P. Boulanger, O. Jokinen, and A. Beraldin. Intrinsic Filtering of Range Images Using a Physically Based Noise Model. In *Proc. 15th Int. Conf. on Vision Interface*, Canada, 2002.
[7] Y. Chen and G. Medoni. Object Modelling by Registration of Multiple Range Images. In *Proc. IEEE ICRA*, USA, April 1991.

Fig. 6. Complete 3D map of the mine. Top: zx-map (view from top). Bottom: yx-map (view from side). z is the depth axis and and y the elevation. units: cm.

[8] P. Corke, J. Cunningham, D. Dekker, , and H. Durrant-Whyte. Autonomous underground vehicles. In *Proc. of the CMTE Mining Technology Conference*, pages 16–22, Australia, September 1996.
[9] D. Eggert, A. Fitzgibbon, and R. Fisher. Simultaneous Registration of Multiple Range Views Satisfying Global Consistency Constraints for Use In Reverse Engineering. *Computer Vision and Image Understanding*, 69:253 – 272, March 1998.
[10] D. Ferguson, A. Morris, D. Hähnel, C. Baker, Z. Omohundro, C. Reverte, S. Thayer, W. Whittaker, W. Whittaker, W. Burgard, and S. Thrun. An autonomous robotic system for mapping abandoned mines. In S. Thrun, L. Saul, and B. Schölkopf, editors, *Proc. of NIPS*. MIT Press, 2003.
[11] J. Folkesson and H. Christensen. Outdoor Exploration and SLAM using a Compressed Filter. In *Proc. of the IEEE ICRA*, Taiwan, 2003.
[12] N. Gelfand, S. Rusinkiewicz, and M. Levoy. Geometrically Stable Sampling for the ICP Algorithm. In *Proc. IEEE 3DIM*, Canada, 2003.
[13] M. Greenspan and M. Yurick. Approximate K-D Tree Search for Efficient ICP. In *Proc. IEEE 3DIM*, Canada, 2003.
[14] M. Hebert, M. Deans, D. Huber, B. Nabbe, and N. Vandapel. Progress in 3–D Mapping and Localization. In *Proc. SIRS*, France, 2001.
[15] B. Horn. Closed–form solution of absolute orientation using unit quaternions. *Journal of the Optical Society of America A*, 4(4):629 – 642, April 1987.
[16] A. Morris, D. Kurth, W. Whittaker, and S. Thayer. Case studies of a borehole deployable robot for limestone mine profiling and mapping. In *Proc. FSR*, Japan, 2003.
[17] A. Nüchter, H. Surmann, K. Lingemann, and J. Hertzberg. Consistent 3D Model Construction with Autonomous Mobile Robots. In *Proc. 26th German Conference on AI KI 2003*, Germany, 2003.
[18] E. Pauley, T. Shumaker, and B. Cole. Preliminary report of investigation: Underground bituminous coal mine, non-injury mine inundation accident (entrapment), July 24, 2002, Quecreek, Pennsylvania, 2002. Black Wolf Coal Company, Inc. for the PA Bureau of Deep Mine Safety.
[19] K. Pulli. Multiview Registration for Large Data Sets. In *3DIM*, 1999.
[20] S. Rusinkiewicz and M. Levoy. Efficient variants of the ICP algorithm. In *Proc. IEEE 3DIM*, pages 145 – 152, Canada, 2001.
[21] D. Simon, M. Hebert, and T. Kanade. Real–time 3–D pose estimation using a high–speed range sensor. In *Proc. IEEE ICRA*, USA, 1994.
[22] H. Surmann, A. Nüchter, and J. Hertzberg. An autonomous mobile robot with a 3D laser range finder for 3D exploration and digitalization of indoor environments. *Robotics and Autonomous Systems*, 12 2003.
[23] S. Thrun, D. Fox, and W. Burgard. A real-time algorithm for mobile robot mapping with application to multi robot and 3D mapping. In *Proc. of the IEEE ICRA*, USA, 2000.
[24] S. Thrun, D. Hähnel, D. Ferguson, M. Montemerlo, R. Triebel, W. Burgard, C. Baker, Z. Omohundro, S. Thayer, and W. Whittaker. A system for volumetric robotic mapping of abandoned mines. In *Proc. of the IEEE ICRA*, 2003.

ACKNOWLEDGMENT