2012 International Conference on Indoor Positioning and Indoor Navigation, 13–15th November 2012

Mobile 3D Indoor Mapping Using the Continuous Normal Distributions Transform

Dylan Campbell
School of Mechanical and Manufacturing Engineering
University of New South Wales
Sydney, Australia
dylan.campbell@student.unsw.edu.au

Mark Whitty
School of Mechanical and Manufacturing Engineering
University of New South Wales
Sydney, Australia
m.whitty@unsw.edu.au

Samsung Lim
School of Surveying and Geospatial Engineering
University of New South Wales
Sydney, Australia
s.lim@unsw.edu.au

***Abstract*—Existing approaches for indoor mapping are often either time-consuming or inaccurate. This paper presents the Continuous Normal Distributions Transform (C-NDT), an efficient approach to 3D indoor mapping that balances acquisition time, completeness and accuracy by registering scans acquired from a rotating LiDAR sensor mounted on a moving vehicle. C-NDT uses the robust Normal Distributions Transform (NDT) algorithm for scan registration, ensuring that the mapping is independent of the long-term quality of the odometry. We demonstrate that C-NDT produces more accurate maps than stand-alone dead-reckoning, achieves better map completeness than static scanning and is at least an order of magnitude faster than existing static scanning methods.**

I. Introduction

The mapping of indoor environments is a time-consuming and expensive process. Contemporary static scanning systems can create maps with high accuracies and resolutions, but require a considerable investment of time and expertise [1]. For applications where mapping accuracy is not the dominant consideration, a mobile mapping system would be ideal. In addition to being operable by non-experts, mobile systems can map at a significantly faster rate and produce more complete maps.

However, mobile mapping systems are subject to dynamic effects, such as vibrations, which introduce errors relative to the platform. Furthermore, the accuracy of the platform's position estimate deteriorates over time, since dead-reckoning produces small errors that accumulate. The combination of relative errors and localisation errors significantly reduces the quality of mapping.

To compensate for these incremental errors, exteroceptive sensors, such as laser rangefinders and cameras, can be used to align the new viewpoint with previously mapped areas. The Iterative Closest Point (ICP) algorithm [2], [3] is frequently used to register scans, thereby localising the platform without cumulative errors from the odometry.

The main contribution of this work is the integration of the 3D Normal Distributions Transform (3D-NDT) scan registration algorithm [4], [5] with mobile LiDAR to create a system for real-time indoor surveying, which is called Continuous NDT (C-NDT). The algorithm's robustness to significant rotations between scans solves the primary problem associated with mobile scanning, that is, the problem that vehicle rotations are less controllable in mobile scanning than in static scanning.

There are numerous applications for 3D maps of indoor environments. For example, navigation systems require a map to plan the optimal path to a destination, such as for an autonomous forklift. In addition, the provision of a 3D map would assist a device equipped with exteroceptive sensors (such as cameras or laser scanners) to localise itself precisely. Other applications include engineering analysis, restoration or renovation planning and documenting the structural design of heritage buildings [6], [7]. Mobile mapping systems are of particular interest to the robotics community, since the ability to map the environment while simultaneously localising itself allows the robot to perform useful tasks autonomously, see Thrun et al. [8] for an overview.

II. Literature Review

A. 3D Indoor Mapping Systems

Indoor mapping systems have varied widely throughout history. Originally, mapping was carried out by taking physical measurements with tape measures and theodolites. Since the invention of the laser rangefinder, this process has been greatly simplified. Current 3D indoor mapping systems can be categorised based on their method of data acquisition and registration.

Pomerleau et al. [9] used a total station to localise a set of static laser scanner positions with millimetre precision. To measure the orientation of the scanner as well as the position, they mounted three reflective prisms around the platform, allowing the total station to make three separate distance measurements. However, this system became inefficient when the position and mounting angles of the scanner could not be measured from the position of the total station. When this occurred, the total station had to be moved and relocalised, greatly increasing the mapping time.

978-1-4673-1954-6/12/$31.00

A faster and more flexible technique is to distribute reflective targets about the scene and use them to register the scans in post-processing. Zimmermann and Eßer [7] used this technique to survey a catacomb, since it allowed some flexibility in positioning the scanner. They reported an average error of 5–6mm when registering scan positions, given a minimum of five reflective targets viewable in common. However, Brenner et al. [1] noted that this technique has several drawbacks. They reported that target distribution, collection and manual intervention during registration takes five times longer than the scanning time. In addition, they reported an uneven distribution of alignment errors due to the ad hoc positioning of targets.

A hybrid technique is often applied for larger sites, where the cumulative errors of target-based registration can become too large [7]. In these cases, a set of ground control points are established by conducting a traverse with a total station and these are used as a baseline for target registration.

Another category of registration technique eschews the use of external targets altogether. The Iterative Point Algorithm (ICP) [2], [3] repeatedly refines the relative position and orientation between two overlapping scans by minimising the sum of squared distances between corresponding points. For each point in one point cloud, the corresponding point is chosen as its nearest neighbour in Euclidean space. Nüchter et al. [10] used ICP to register scans taken from a mobile robot using a "stop-and-scan" methodology, that is, frequent short-term static scanning. They achieved a coarsely estimated 0.5–3.8% deviation from ground-truth.

While ICP is a simple and reasonably accurate registration algorithm, it has several limitations. Firstly, the algorithm may converge to an incorrect local minimum, particularly when the initial position and orientation estimate is poor. Another disadvantage is that points from non-overlapping sections of the scans may cause a systematic bias. The main disadvantage, however, is that ICP requires a reasonable position and orientation estimate to converge correctly and rapidly.

In comparison, the three-dimensional Normal Distributions Transform (3D-NDT) [4], [11] is more robust to poor initial alignments [12]. Magnusson et al. [5] used NDT to register statically-acquired scans from a mine, reporting faster and more reliable performance than ICP. A more detailed description of 3D-NDT is given in Section III.

A thorough comparison of ICP and 3D-NDT, undertaken by Magnusson et al. [12], concluded that 3D-NDT converged from a greater range of initial pose estimates. For the tested datasets, 3D-NDT converged in 77% of cases, compared to 30% for ICP. The tested valley of convergence was ± 2 m in x and y and $\pm 90^\circ$ in yaw. In addition, the study found that 3D-NDT was a faster process, with a median processing time of 2.2s as compared to 5s for ICP. However, ICP was shown to behave more predictably than 3D-NDT. That is, 3D-NDT may converge from a pose estimate with large initial error, but fail from a pose estimate with less error.

Another scan registration technique that improves ICP's robustness to rotational deviations is Metric-Based ICP (MBICP), as introduced by Minguez et al. [13]. MBICP replaces the Euclidean distance used to determine the nearest neighbour in ICP with a distance measure that simultaneously accounts for translational and rotational displacements. This takes into account the fact that a small rotational displacement will cause points far from the sensor to be displaced significantly from their correspondents.

A modification of this algorithm has been used successfully by Milstein et al. [14] to align point clouds to an occupancy grid of previously acquired scan data. However, this occupancy grid application highlights one of the main limitations of MBICP. Techniques used to expedite the search for nearest neighbours in traditional ICP, such as cached *k*d-tree searches [10], cannot be used with MBICP. Instead, a slower search in polar space is required, which would not scale to full 3D point cloud matching.

A more efficient way to map an indoor environment is to use a mobile scanner. Talaya et al. [15] fused data acquired from a moving laser scanner with a position estimate from GPS and an orientation estimate from an inertial measurement unit (IMU). Indoor applications, however, preclude the use of absolute positioning systems like GPS. Since errors in the position and orientation estimates obtained from dead-reckoning accumulate, the registered scans will deviate from ground-truth. As a result, a registration method independent of the odometry is required.

Before scan registration techniques can be applied, the continuously acquired laser data must be segmented into 3D point clouds. The standard approach is to segment the data at the extents of motion for a pitching or yawing laser and use the odometry to form a coherent point cloud [16]. Harrison and Newman [17] applied odometry correction to further ensure the point cloud quality by inferring that near-vertical planes were in fact vertical.

While researchers have applied ICP to register point clouds segmented from a mobile scanner [16], [18], we introduce a system for mobile mapping using NDT. This registration algorithm is particularly suitable for mobile mapping, because it has a wider valley of convergence than ICP and is more robust to rotations [12]. This is advantageous, because a mobile scanner that is turning as it scans will have a considerable angular offset between the segmented point clouds.

B. Benchmarking

The accuracy of scan-matching algorithms is often non-trivial to quantify, due to the difficulty associated with obtaining accurate ground-truth measurements. If accurate 3D environment maps were already available, robotic mapping would be redundant.

Henry et al. [19] obtained ground-truth maps using 2D Simultaneous Localisation and Mapping (SLAM) algorithms, on which the 3D map is overlaid. Similarly, Wulf et al. [20] used architectural floor plans from surveyors as a benchmark and compared them to the 3D maps generated by SLAM algorithms. They also suggested using calibrated aerial photographs or satellite images if existing maps are unavailable.

The 3D map is converted into a 2D map by removing all points that do not lie on vertical surfaces (for example, walls) and then projecting the remaining points onto the horizontal plane. A qualitative, visual inspection of map quality is then possible by checking whether the walls align as expected. Wulf et al. [20] provided a quantitative measure by running a Monte Carlo Localisation algorithm with the reference and projected 2D maps as inputs. However, these methods are not capable of determining the yaw, pitch or elevation accuracies, due to the inherent limitations of the 2D reference map.

C. Overview

The objective of this work is to demonstrate that, for certain applications, C-NDT achieves a more satisfactory balance between acquisition time, completeness and accuracy than static scanning and existing odometry-based mobile mapping systems.

C-NDT uses the robust 3D-NDT algorithm to register scans that have been acquired while the scanner is moving. As a result, it is independent of the long-term quality of the attitude and position estimated from the gyroscopes and wheel encoders. The concept and implementation of C-NDT has been subjected to a number of tests for the purposes of comparison with static scanning and mobile dead-reckoning.

The scope of this work does not extend to a rigorous comparison of the 3D-NDT and ICP algorithms, since this has already been undertaken by Magnusson et al. [12]. However, the ICP algorithm was used in the experiments to confirm the findings of [12] and to benchmark the accuracy of the C-NDT system.

Section III introduces the 3D-NDT scan registration algorithm. The proposed C-NDT mapping system is described in Section IV in detail. Section V presents the experiments undertaken and the results attained, categorised into investigations of accuracy, completeness and efficiency. Finally, the conclusions of this work are summarised in Section VI.

III. Fundamentals

The simplest representation of laser data is the point cloud, a collection of points in 3D space sampled from a surface or surfaces, as shown in Figure 1. While the point cloud is useful for visualisation, it does not describe the characteristics of the underlying surface explicitly. In contrast, the three-dimensional Normal Distributions Transform (3D-NDT) incorporates these characteristics in a representation that can be used for efficient scan registration [5].

The 3D-NDT algorithm subdivides the point cloud into a 3D grid of cells (known as voxels) and computes a Probability Density Function (PDF) for each cell. The algorithm then finds the transformation that maximises the likelihood that the points of another point cloud lie on this reference surface. This transformation represents the change in position and orientation of the laser rangefinder and so can be used to track the robot's pose.

The PDF used in the NDT algorithm is a mixture of a normal distribution and a uniform distribution, as given in Equation 1, where $\mathbf{x}$ is a point in 3D space, μ is the centroid of the voxelised point cloud, Σ is the covariance matrix of the point cloud, c_1 and c_2 are constants and p_o is the expected ratio of outliers. A mixed PDF is used since a pure normal distribution is not robust to outliers [21].

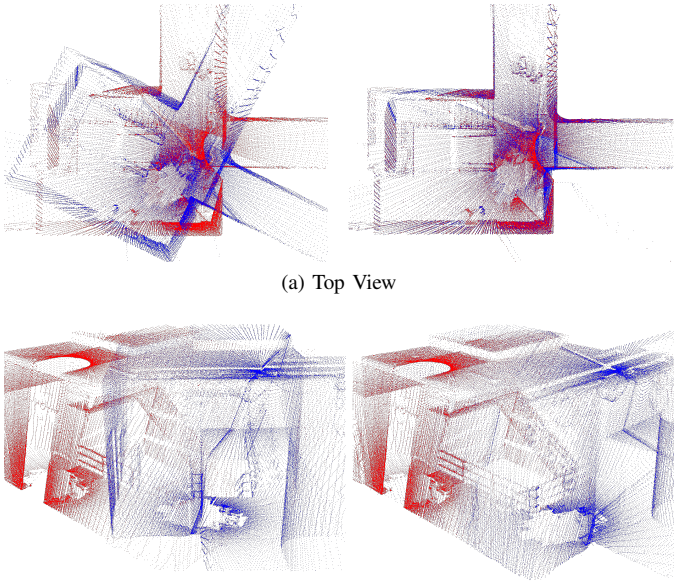

(a) Top View

(b) Oblique View

Fig. 1. The top and oblique views of two non-aligned 3D point clouds are shown in red and blue on the left. On the right, they have been aligned using 3D-NDT and merged to form an accurate and more complete representation of the environment.

$$p(\mathbf{x}) = c_1 \exp\left(-\frac{1}{2}(\mathbf{x}-\mu)^T\Sigma^{-1}(\mathbf{x}-\mu)\right) + c_2 p_o \qquad (1)$$

To find the transformation that maximises the likelihood that another point cloud lies on the NDT surface, the negative log-likelihood of the function Ψ is minimised, as given in Equation 2. $T(\mathbf{p}, \mathbf{x}_k)$ is the transformation that translates and rotates a point $\mathbf{x}_k$ in the new point cloud by the pose $\mathbf{p}$.

$$\Psi = \prod_{k=1}^{n} p\left(T(\mathbf{p}, \mathbf{x}_k)\right) \qquad (2)$$

The log-likelihood of Ψ is approximated by a Gaussian distribution, from which the derivatives can easily be calculated. Since the NDT surface is piecewise smooth with continuous first- and second-order derivatives, standard numerical optimisation techniques, such as Newton's method, can be used. The details of the algorithm and the relevant equations can be found in Magnusson [11].

The reasons for using NDT in a continuous mapping system, rather than the more common ICP algorithm, are twofold. Firstly, NDT obviates the need for the computationally expensive nearest neighbour search present in the ICP algorithm, resulting in faster execution times. Secondly, NDT is more robust to poor initial alignments than ICP [12]. This is crucial for a mobile system, since, unlike a static system, rotations between scans cannot be avoided.

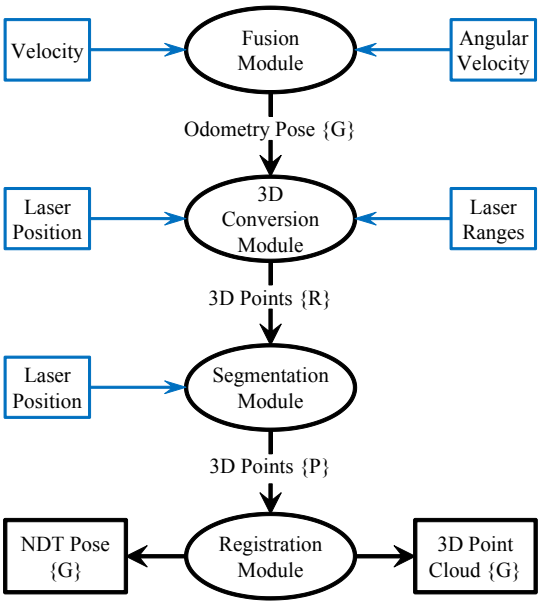

Fig. 2. Flowchart of the C-NDT system. Input and output data are shown in boxes with blue and black outlines respectively and processing modules are shown as black ellipses.

IV. Method

To construct a coherent 3D map without relying on the long-term accuracy of onboard odometry or independent sensors such as GPS, C-NDT segments the continuously acquired range data into 3D point clouds and registers them into the map frame using NDT. In the following treatment, $\{G\}$, $\{R\}$ and $\{P\}$ are the global, robot and point cloud reference frames respectively. In addition, 'pose' refers to the position and orientation of the body in six degrees of freedom (x, y, z, roll, pitch and yaw). The laser rangefinder is mounted on a robot and is rotated by $\pm 90°$ about a near-vertical axis by a DC motor. A flowchart of the process is shown in Figure 2.

An outline of the fusion module is given in Guivant and Nebot [23]. In brief, the fusion module estimates the robot's pose with six degrees of freedom using a dead-reckoning approach. This allows inclines and unstructured natural environments to be navigated and mapped. The attitude is estimated by integrating the angular velocities from the gyroscopes in the Inertial Measurement Unit (IMU). By assuming that the velocity vector is parallel to the attitude vector (as in [24]), the 3D position of the robot can be calculated from the attitude and velocity, provided by the wheel encoders.

An outline of the 3D conversion module is given in Whitty et al. [22]. The module synchronises the data from the laser rangefinder and motor encoder. After synchronisation, the yaw angle of the laser is known accurately at the beginning and end of each laser scan. The polar coordinates from the rangefinder are transformed into the robot's Cartesian coordinate system $\{R\}$ and interpolated.

To make use of 3D scan-matching algorithms, the data stream from the rangefinder must first be segmented into 3D point clouds. Each $180°$ sweep of the rangefinder forms a full 3D point cloud and thus the extents of laser yaw become a natural boundary for segmentation. Consequently, the beginning of a new point cloud is triggered when the motor encoder angle nears the extents of rotation.

The robot's pose at the beginning of the sweep is taken as the point cloud coordinate frame $\{P\}$. Every point observed within the $180°$ sweep is transformed into this frame using the odometry pose. This is performed by first transforming the point into the global coordinate frame $\{G\}$ and then transforming back into the point cloud coordinate frame $\{P\}$. The first transformation uses interpolated poses that account for movement between the start and end of one laser scan.

Once the continuous laser data has been segmented into a distinct 3D point cloud, it is registered into the global or map frame $\{G\}$ using 3D-NDT. Scan-matching dissociates map accuracy from dead-reckoning accuracy. While one can be sufficiently confident in the accuracy of the odometry during the 2.7s taken to acquire the point cloud, the odometry produces small errors that accumulate over time. Instead of the odometry, C-NDT uses an independent exteroceptive source of information, the point clouds themselves, to localise.

The output of the C-NDT system is a 3D map and a sequence of poses that correspond to the path taken by the mobile scanner. The map can be viewed in real-time as the scanner progresses. This provides the operator with an immediate indication of mapping quality. Rather than finding occluded areas during post-processing and having to rescan, they can be identified by the operator during mapping and addressed immediately.

V. Experiments and Results

A. Equipment

The unmanned ground vehicle (UGV) used as a platform for this project has been described in Whitty et al. [22] and Guivant et al. [25] and is shown in Figure 3. It is equipped with a SICK LMS151 2D laser rangefinder, which has a full view angle of $270°$ and can measure distances up to a maximum of 50m with a 1σ statistical error of 12mm. In addition, a value for the reflection intensity is provided for each range measurement. A DC motor rotates the upward-facing laser scanner around a near-vertical axis by $90°$, capturing a wide field of view. The UGV is also fitted with a Microstrain 3DM-GX3 Inertial Measurement Unit (IMU) and motor encoders.

B. Mapping Details

The experiments took place on the second floor of an engineering laboratory building, containing cluttered rooms of various sizes and straight, sparsely-featured corridors. In the primary dataset, the robot travelled approximately 110m, entering four rooms and traversing long sections of corridor, as shown in Figure 5a. The area and volume mapped was 340m^2 and 1220m^3 respectively. Within some rooms, the robot turned a full revolution in order to exit from the same door, which would generally degrade the quality of mapping using dead-reckoning, since the gyroscope errors would accumulate quickly.

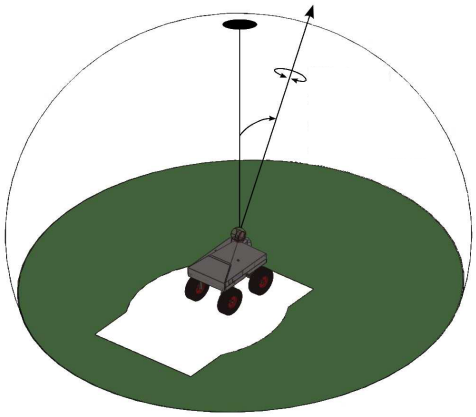

(a) The actuated laser arrangement gives a 360° field of view about the tilted axis of rotation.

(b) The actuated laser rangefinder can be seen on the top of the UGV.

Fig. 3. One of the UGVs developed by UNSW Mechatronics was used to collect data for the experiments in this paper.

C. Accuracy

To compare the accuracy of C-NDT to existing ICP-based and odometry-based mobile mapping systems, 3D point clouds were constructed from the primary dataset using each of these methods. Both qualitative and quantitative methods were used to determine model accuracy.

In order to establish a correct ground-truth for quantitative analysis, physical measurements were taken along the robot's path using a tape measure. The robot's poses, as generated by C-NDT, ICP and the odometry, were recorded at each of the measured locations. The Euclidean distance between a reference position towards the beginning of the traverse and any subsequent location was then compared to the corresponding ground-truth measurement.

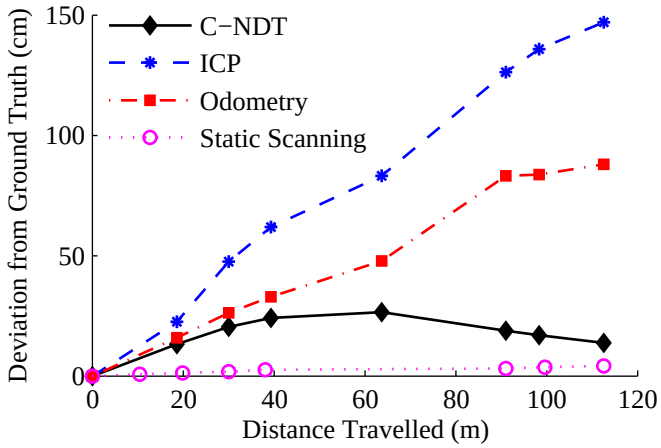

Fig. 4. Comparison of the accuracy of C-NDT, ICP, odometry and simulated static scanning using the traverse in Figure 5.

Figure 4 displays the deviation between the reference distances measured and the distances reported by C-NDT, ICP and the odometry. These deviations from ground-truth provide a metric that can be used to estimate accuracy and compare between mapping methods. The deviation is plotted against the cumulative distance travelled, since the mapping errors are expected to be proportional to distance. It is important to note that from 50m to 80m distance travelled, the robot mapped two confined rooms and was forced to turn sharply on many occasions. This would tend to reduce the accuracy of the odometry at a greater rate than the accuracy of C-NDT. The paths and reference positions are plotted in Figure 5a.

The results show that the map generated using C-NDT is significantly more accurate than that generated from the odometry or ICP. Moreover, the values for the odometry and ICP suggest that their mapping accuracies degrade at a significant rate with distance travelled, from the accumulation of errors. In contrast, the values for C-NDT show that its accuracy degrades much slower with distance travelled.

To estimate the comparative accuracy of static scanning, the error arising from the registration of discrete point clouds has been estimated. Zimmermann and Eßer [7] achieved a mean error of 5–6mm when registering scan positions, given a minimum of five reflective targets viewable in common. To register a scan from the end of the map to the beginning of the map, the geometry dictates that at least eight scanning positions would be required, every 10m. The cumulative registration error for static scanning is also shown in Figure 4.

The maximum registration error of 48mm compares favourably with the 136mm error obtained by C-NDT in the last scan of the experiment and the peak C-NDT error of 265mm. In addition, mobile scanning introduces errors relative to the moving platform due to the dynamic effects such as vibration. As a result, static scanning can achieve greater accuracies than mobile scanning.

(a) Experimental setup showing the path taken by the UGV overlaid on the existing floor plan. Reference points are detailed in the text.

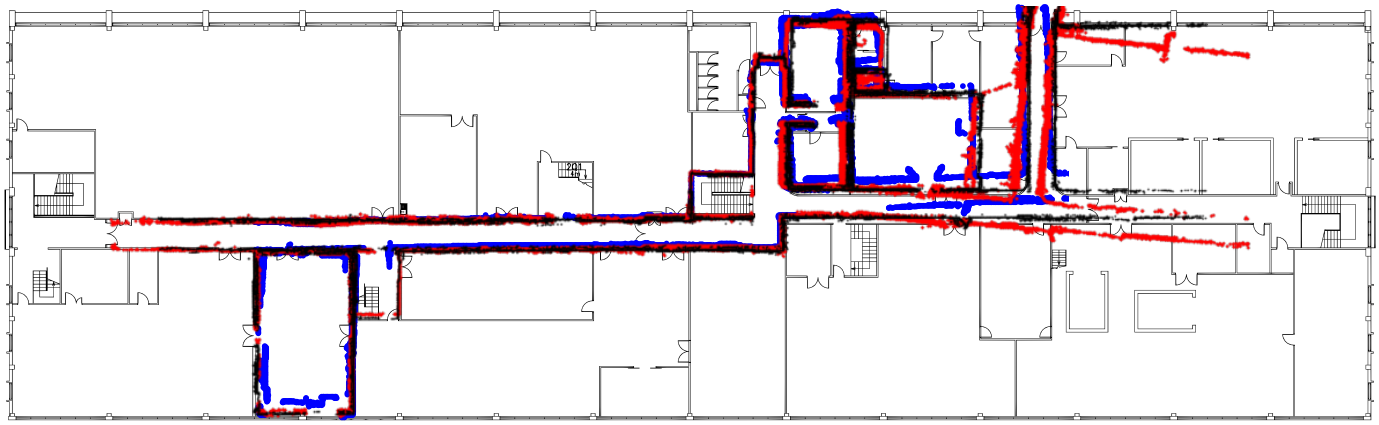

(b) A 2D projection of the wall points extracted from the C-NDT, ICP and odometry-based point clouds, shown in black, blue and red respectively.

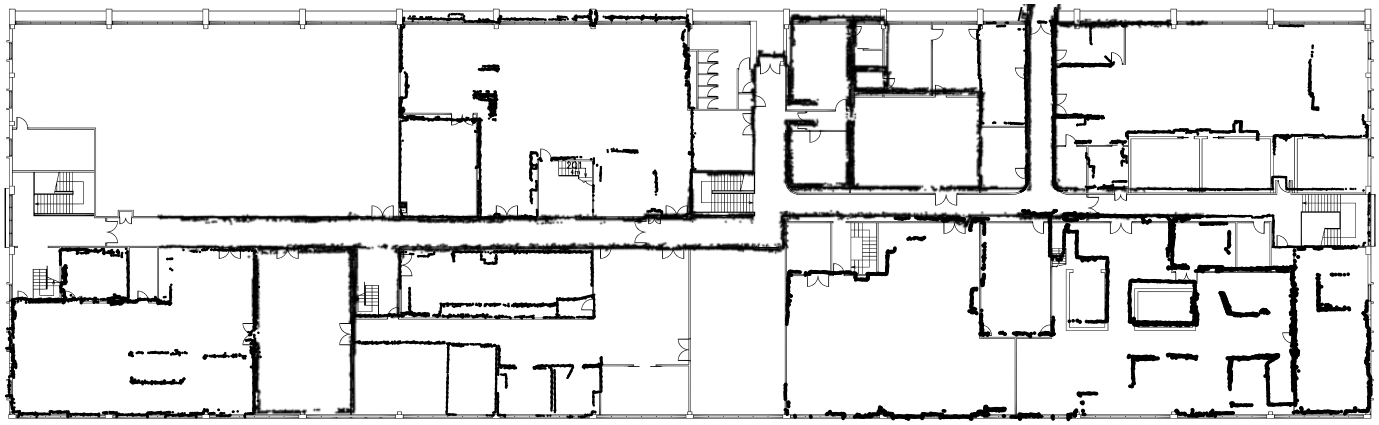

(c) A 2D projection of the wall points extracted from several C-NDT maps, acquired over multiple days. To aid visualisation, the full 3D point clouds, including floor and ceiling points, are not shown.

Fig. 5. A visual comparison of C-NDT, ICP and odometry-based pose estimates and mapping accuracy.

In Figure 5b, 2D projections of the point clouds constructed using C-NDT, ICP and pure odometry are shown, overlaying the floor plan. A qualitative visual inspection reveals that the point clouds constructed using odometry and ICP quickly accumulated significant errors in rotation and translation. In comparison, the C-NDT point cloud corresponded very closely to the ground truth and did not drift noticeably.

While there are a few apparent discrepancies, the field survey indicated that the "ground-truth" reference map was in fact less accurate than C-NDT. In particular, the two main corridors are offset incorrectly in the floor plan.

Figure 5c gives further qualitative evidence for the accuracy of C-NDT, showing a 2D projection of the map constructed from a series of traverses, conducted over several days.

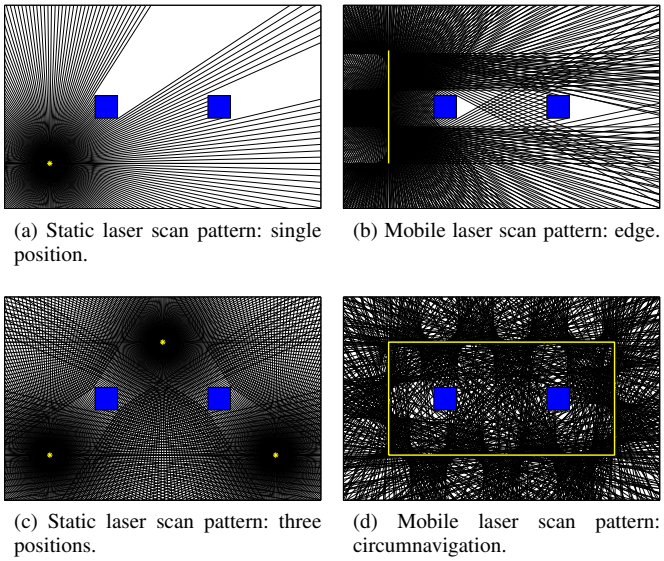

(a) Static laser scan pattern: single position.

(b) Mobile laser scan pattern: edge.

(c) Static laser scan pattern: three positions.

(d) Mobile laser scan pattern: circumnavigation.

Fig. 6. Comparing the completeness obtained by mobile and static scanning, using the example of a rectangular room containing two obstacles (in blue). The scanning points (at left) and path (at right) are displayed in yellow.

In this experiment, the traverse was open loop. For both static and mobile scanning, a closed loop traverse would have bounded the cumulative registration error, distributing it around the loop, such as in [26].

D. Completeness

A primary advantage of mobile mapping systems is that they are able to observe a scene from many different viewpoints. To demonstrate the consequent increase in map completeness, static and mobile scanning has been simulated in a simple 2D environment with two interior columns. Completeness is defined as the percentage of a scene that has been mapped to a given resolution.

Figure 6 shows the scan patterns formed by a fixed laser and a moving laser with the same angular resolution. It shows the shadowing caused by the columns, whereby sections of wall and column are not mapped. To ensure a complete map in the presence of occlusions, a number of static scanning positions are required. In comparison, a mobile scanner will obtain a complete map in the course of moving through the room.

The map completeness percentages for the static and mobile scanning simulations in Figures 6c and 6d were 30% and 87% respectively. These values refer to the percentage of fixed-size cells that contained at least one laser measurement along the surfaces of the simulated environment. Both simulations used the same angular resolution in the horizontal plane (1.34°) and a cell size of 1cm.

In this simulation, interior columns are used as a clear example of occlusion. However, this investigation applies equally to any other obstructions, such as furniture and partitions. A cluttered environment will require a large number of scan positions to achieve full completeness. Moreover, for real-world environments, choosing the scan positions is non-trivial and so heuristics are generally used. Poor choices will result in incomplete maps with greatly varying point densities.

TABLE I
ESTIMATED TIMES FOR STATIC AND MOBILE SCANNING.

Static	**Time (min)**	**Type**	**Description**
Setup	3	One-off	Arrange reflective targets
Setup	2	Per scan	Move, set-up and level tripod and scanner
Data Collection[1]	3	Per scan	Scanning to achieve the minimum resolution
Post-Processing	5	Per scan	Registration of each point cloud
Mobile	**Time (min)**	**Type**	**Description**
Setup	0.5	One-off	Start UGV and run programs
Data Collection	0.06	Per metre	Driving at 0.3 m/s
Post-Processing	0	Per metre	All processing is real-time

E. Efficiency

To demonstrate the efficiency of this mobile mapping system, the setup, data collection and processing times have been compared to the standard static approach. The variable that has been kept constant is map completeness.

Consider the mapping of a room with interior columns, such as shown in Figure 6. For n collinear columns, a minimum of n + 1 static scanning positions would be required to ensure full and even coverage of the room. A greater number of scanning positions would be necessary if line-of-sight was required between them, such as during a surveying traverse. In comparison, a mobile mapping solution would need only to circumnavigate the room.

For this experiment, the minimum resolution for both static and mobile mapping was chosen to be 2cm. Using the simulation, it was found that the mobile platform should drive at 0.3m/s to achieve this level of completeness. The values used to estimate setup, data collection and post-processing durations for mobile and static scanning are listed in Table I. The static scanning durations are low, conservative estimates chosen in consultation with a surveying professional.

Figure 7 demonstrates that for an increasing number of interior columns, the total time to survey the room statically is an order of magnitude greater, and increases at a greater rate, than for mobile scanning.

The comparison being made is between the mobile scanning system and static scanning using a typical surveying rangefinder. Surveying rangefinders often have fixed maximum rotation rates about the vertical axis, whereas the scanner used in the mobile system can rotate at an arbitrary speed. If the static scanner had the same scanning speed as the mobile scanner (2.7s per scan), the data collection component of

[1]Laser rangefinders used in surveying often have fixed maximum rotation rates about the vertical axis, which varies between products.

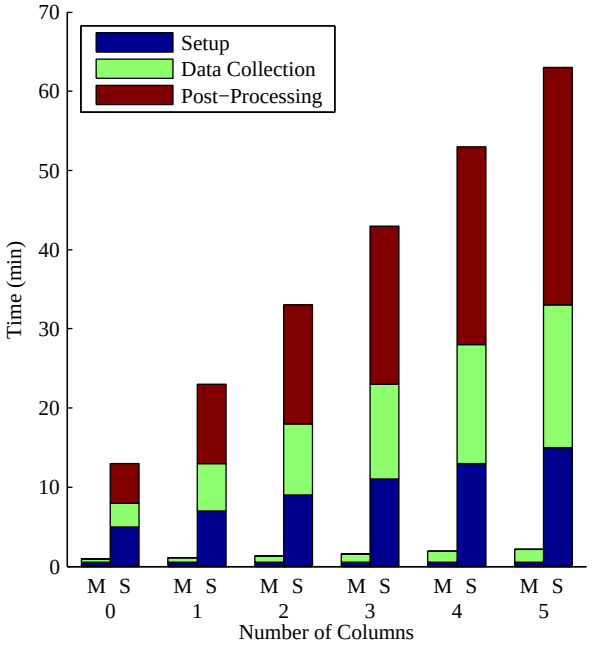

Fig. 7. Comparing the efficiency of mobile (M) and static (S) scanning using the estimated times from Table I.

the scanning time would be reduced by 98.5%. Even in this case, the total time to survey the room would be an order of magnitude greater than for mobile scanning.

One contribution to C-NDT's efficiency is its ability to perform all scan registration processing in real-time. The mean time taken to register a point cloud is 1.31s with a standard deviation of 0.73s, whereas the time taken to acquire the point cloud is 2.70s.

For the traverse shown in Figure 5, it took 6.3 minutes to map an area of 340m^2. In comparison, Dongzhen et al. [27] reported that it took 1.5 days to map an indoor area of 550-600m^2 using static scanning. Hence, per square metre, C-NDT is two orders of magnitude faster, taking only 1.4% of the time taken by static scanning.

In addition, the map shown in Figure 5c took 46.2 minutes to acquire with C-NDT. It covers an area of 2019m^2, including 24 rooms and sections of corridor. While it is non-trivial to determine the required static scanning locations for such a map, particularly since the rooms are cluttered and contain many obstacles, a very conservative estimate would be 50 scanning positions. Using the durations from Table I, this would take at least 8.4h, a tenfold increase in scanning time when compared to mobile scanning. This time estimate is itself conservative, because it does not include the time taken to move the reflective targets.

VI. Conclusion

In this paper, we have presented a system for indoor mapping using a mobile laser scanner: Continuous NDT. C-NDT segments continuously acquired laser data and processes it in real-time using the 3D-NDT scan registration algorithm.

We have demonstrated that C-NDT produces more accurate maps than stand-alone dead-reckoning, with a quantitative analysis showing 85% less deviation from ground-truth on the tested indoor traverse. While this is an indicative statistic only, the qualitative results presented in this paper, from many additional traverses, support this general conclusion. Nonetheless, a comprehensive study, examining many hundreds of datasets from different buildings, would be required to establish more generalisable results.

Furthermore, we have shown that a mobile scanner is able to achieve a higher map completeness than a static scanner, since it is able to observe a scene from many different viewpoints, reducing shadowing effects. Finally, we have demonstrated that C-NDT is at least an order of magnitude faster than static scanning for a constant map completeness.

We conclude that, for applications where the accuracy requirements are slightly lower, C-NDT achieves a more satisfactory balance between acquisition time, completeness and accuracy than static scanning and existing odometry-based mobile mapping systems.

Additional benchmarking methods could be used to more closely characterise the accuracies achievable with C-NDT. In particular, the deviation of the estimated pose could be determined more precisely by using a total station to measure the actual pose at a set of positions within the traverse.

Real-time extraction and matching of geometric features from each segmented point cloud opens the possibility of extending the C-NDT system. This would enable the automatic registration of maps produced on separate occasions or by separate vehicles and assist in closing the loop. While we have not studied the effect of closing the loop on C-NDT accuracy, results from the field of Simultaneous Localisation and Mapping (SLAM) suggest a large improvement in accuracy can be expected while maintaining real-time processing capability.

Acknowledgment

The authors would like to thank José Guivant, Steve Cossell and Jayantha Katupitiya for the use of the raw datasets. Modules for data acquisition ("Sensor.exe", "DMC_read.exe" and "LMS151_manager.exe"), data recording ("Possum.exe"), data replaying ("Simulator.exe") and vehicle 3D pose estimation ("Fusion.exe") using dead-reckoning were used with the permission of José Guivant and further details can be found here: www.possumrobot.com/Software/Possum/Applications/{SensorsExe.htm, DMC_read.htm, LMS151_manager.htm, Possum.htm, Simulator.htm, FusionExe.htm}.

References

[1] C. Brenner, C. Dold, and N. Ripperda, "Coarse orientation of terrestrial laser scans in urban environments," *ISPRS Journal of Photogrammetry and Remote Sensing*, vol. 63, no. 1, pp. 4–18, 2008.

[2] P. Besl and N. McKay, "A method for registration of 3-D shapes," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 14, no. 2, pp. 239–256, 1992.

[3] Y. Chen and G. Medioni, "Object modelling by registration of multiple range images," *Image and Vision Computing*, vol. 10, no. 3, pp. 145–155, 1992.

[4] M. Magnusson, "3D scan matching for mobile robots with application to mine mapping," Ph.D. dissertation, 2006.

[5] M. Magnusson, A. Lilienthal, and T. Duckett, "Scan registration for autonomous mining vehicles using 3D-NDT," *Journal of Field Robotics*, vol. 24, no. 10, pp. 803–827, 2007.

[6] A. Vettore, A. Guarnieri, M. Pontin, and J. Beraldin, "Digital 3D reconstruction of scrovegni chapel with multiple techniques," in *XXth ISPRS Congress*, vol. 35, 2004, pp. 159–164.

[7] N. Zimmermann and G. Eßer, "Showing the invisible documentation and research on the roman domitilla catacomb, based on image laser scanning and 3D modelling," in *35th International Conference on Computer Applications and Quantitative Methods in Archaeology (CAA)*, A. Posluschny, K. Lambers, and I. Herzog, Eds., 2008.

[8] S. Thrun, W. Burgard, and D. Fox, *Probabilistic Robotics*. Cambridge, MA: MIT Press, 2005.

[9] F. Pomerleau, B. Lescot, F. Colas, M. Liu, and R. Siegwart, "Dataset acquisitions for USAR environments," 2011.

[10] A. Nüchter, K. Lingemann, J. Hertzberg, and H. Surmann, "6D SLAM3D mapping outdoor environments," *Journal of Field Robotics*, vol. 24, no. 8-9, pp. 699–722, 2007.

[11] M. Magnusson, "The three-dimensional normal-distributions transform: an efficient representation for registration, surface analysis, and loop detection," Ph.D. dissertation, 2009.

[12] M. Magnusson, A. Nuchter, C. Lorken, A. Lilienthal, and J. Hertzberg, "Evaluation of 3D registration reliability and speed - a comparison of ICP and NDT," in *IEEE International Conference on Robotics and Automation (ICRA), 2009*. IEEE, pp. 3907–3912.

[13] J. Minguez, F. Lamiraux, and L. Montesano, "Metric-based scan matching algorithms for mobile robot displacement estimation," in *IEEE International Conference on Robotics and Automation*, vol. 4. Citeseer, 2005, pp. 35–57.

[14] A. Milstein, M. McGill, T. Wiley, R. Salleh, and C. Sammut, "A method for fast encoder-free mapping in unstructured environments," *Journal of Field Robotics*, 2011.

[15] J. Talaya, R. Alamus, E. Bosch, A. Serra, W. Kornus, and A. Baron, "Integration of a terrestrial laser scanner with GPS/IMU orientation sensors," in *XXth ISPRS Congress, Commision 5*, pp. 990–995.

[16] D. Cole and P. Newman, "Using laser range data for 3D slam in outdoor environments," in *IEEE International Conference on Robotics and Automation (ICRA), 2006*. IEEE, pp. 1556–1563.

[17] A. Harrison and P. Newman, "High quality 3D laser ranging under general vehicle motion," in *IEEE International Conference on Robotics and Automation*. IEEE, 2008, pp. 7–12.

[18] D. Holz, C. Lorken, and H. Surmann, "Continuous 3d sensing for navigation and slam in cluttered and dynamic environments," in *XIth International Conference on Information Fusion*. IEEE, pp. 1–7.

[19] P. Henry, M. Krainin, E. Herbst, X. Ren, and D. Fox, "Rgb-d mapping: Using kinect-style depth cameras for dense 3D modeling of indoor environments," *The International Journal of Robotics Research*, vol. 31, no. 5, pp. 647–663, 2012.

[20] O. Wulf, A. Nüchter, J. Hertzberg, and B. Wagner, "Benchmarking urban six-degree-of-freedom simultaneous localization and mapping," *Journal of Field Robotics*, vol. 25, no. 3, pp. 148–163, 2008.

[21] P. Biber, S. Fleck, and W. Strasser, "A probabilistic framework for robust and accurate matching of point clouds," *Pattern Recognition*, pp. 480–487, 2004.

[22] M. Whitty, S. Cossell, K. Dang, J. Guivant, and J. Katupitiya, "Autonomous navigation using a real-time 3D point cloud," in *2010 Australasian Conference on Robotics and Automation*.

[23] J. Guivant and E. Nebot, "Optimization of the simultaneous localization and map-building algorithm for real-time implementation," *Robotics and Automation, IEEE Transactions on*, vol. 17, no. 3, pp. 242–257, 2001.

[24] G. Dissanayake, S. Sukkarieh, E. Nebot, and H. Durrant-Whyte, "The aiding of a low-cost strapdown inertial measurement unit using vehicle model constraints for land vehicle applications," *IEEE Transactions on Robotics and Automation*, vol. 17, no. 5, pp. 731–747, 2001.

[25] J. Guivant, S. Cossell, M. Whitty, and J. Katupitiya, "Internet-based operation of autonomous robots: The role of data replication, compression, bandwidth allocation and visualization," *Journal of Field Robotics (to appear)*, 2012.

[26] G. Sibley, C. Mei, I. Reid, and P. Newman, "Vast-scale outdoor navigation using adaptive relative bundle adjustment," *The International Journal of Robotics Research*, vol. 29, no. 8, pp. 958–980, 2010.

[27] J. Dongzhen, T. Y. Khoon, Z. Zheng, and Z. Qi, *Indoor 3D Modeling and Visualization with a 3D Terrestrial Laser Scanner*, ser. Lecture Notes in Geoinformation and Cartography. Springer Berlin Heidelberg, 2009, ch. 15, pp. 247–255.