Real-Time Map Building with Uncertainty using Colour Camera and Scanning Laser

James Underwood, Steve Scheding and Fabio Ramos
ARC Centre of Excellence for Autonomous Systems, Australian Centre for Field Robotics
The Rose St. Building, J04, University of Sydney, NSW, Australia 2006
Email: j.underwood/scheding/f.ramos @acfr.usyd.edu.au

Abstract

This paper presents a real-time data fusion technique to produce three dimensional colour point clouds of the environment from a mobile platform. Range from a scanning laser and colour from a video camera are locally fused using deterministic transforms. Platform pose information is used to provide global coordinates. Measurement uncertainty is maintained throughout the transformations to enable the fusion of multiple data sets obtained at different times or from different locations, and to enable applications that use the data to assess its quality. The complete system is implemented on an autonomous outdoor ground vehicle and vast colour models of outdoor terrain are built in real time. Although similar techniques have been described in the literature, they have not been applied to a large scale outdoor environment from a mobile platform. The contribution of this paper is to show that very high quality three dimensional colour point clouds can be built in real-time, with a relatively simple approach.

1 Introduction

Three dimensional map building is important in many robotic applications including terrain mapping for path planning and feature detection for pose estimation. In some cases enough data is available directly from local sensors such as [Simmons *et al.*, 1996] which uses local stereo point clouds to determine traversability in the field of view of the sensor. However, when only limited local sensory data is available, building consistent three dimensional maps becomes important as the local data must be compiled for larger coverage or accuracy. Examples of this include [Miller and Campbell, 2006] and [Pfaff and Burgard, 2005] in which elevation maps are built from two dimensional scanning lasers for path planning and obstacle avoidance. Techniques have been developed which augment these types of sensor with colour to build three dimensional colour maps. Several applications benefit greatly from this combination. Applications where human interpretation is important, including archeological, architectural or urban planning where mapping is the specific goal are presented in [Frueh *et al.*, 2004] and [Abmayr *et al.*, 2004]. The technique presented in [Ohno and Tadokoro, 2005] uses three dimensional maps with SLAM to improve pose estimation. At the same time it is combined with colour camera data to make it more appropriate for human interpretation of accident scenes. All of these techniques use a stationary platform and achieve dense three dimensional maps by panning the sensor platform. Although these techniques are highly successful for certain applications, they are not particularly suited for real-time use, particularly from a moving platform.

In outdoor environments it is not possible to use geometry alone to accurately determine regions of traversability. As discussed in [Durrant-Whyte and Scheding, 2003] it is necessary that the sensory information 'spans the decision space' and colour and visual texture information dramatically increase the available information for a perceptual system. Techniques using vision and laser have been applied successfully to tasks such as road following with [Thrun *et al.*, 2006] and [Rasmussen, 2002] where in both cases the classifications use raw sensory information as input. If classification models could instead be learnt from platform independent models of the environment (such as globally located colour/spatial information), then the systems would be more robust, more modular, and even sharable across different robotic platforms.

If the uncertainty of the data is well modeled, then separate sources of information can be fused, regardless of differences in viewpoint, quality of sensor or time when data was collected. In this paper, colour point clouds are built using a colour camera and a scanning laser, but the technique to augment the spatial information

with additional sensory properties is extensible to other sensor modalities, such as radar and infrared.

This paper is organised as follows: Section 2 describes the choice of navigation system to best support outdoor mapping, Section 3 describes the coordinate frames used and the transformations between them. Section 4 describes the preservation of uncertainty through the transformations. Section 5 presents the results and Section 6 concludes.

2 Navigation for Outdoor Mapping

Platform pose information is required to keep the map consistent over time. This could consist of dead reckoning, optimised pose estimation from the map data using a 'sharpness measure' such as [Chandran and Newman, 2006], Global Positioning System (GPS) based techniques that provide a global coordinate frame, or a combination of these techniques. Although GPS based systems are not as locally consistent (discontinuities occur when fusing the global information) as either dead reckoning or data centric optimisations, they do provide several benefits. Data centric techniques require multiple views of the same environment to provide good results and in an outdoor application, often only a single traversal of the environment is made. Pure dead reckoning techniques provide the smoothest local environment maps, but in situations where multiple passes of the same area are made, they are inconsistent due to drift. Therefore in this work a GPS/IMU (Inertial Measurement Unit) navigation system is used to obtain platform location. To prevent inconsistencies in the final map, the GPS/IMU uncertainty is transformed with the data so that the uncertainty of every data point in the map is known, as discussed in detail in Section 4.

3 Coordinate Frames and Transformations

The physical layout of the sensors is shown in Figure 1. The flow of information from collecting raw laser and video data to the final three dimensional colour map is illustrated in Figure 2. In this figure, the pictures represent the coordinate frames and the arrows are the transforms between them. This process consists of two pathways. The first transforms the laser data into the camera image frame to obtain the colour information for each scan point. The second transforms the laser data to the navigation frame and from there to the global frame using the navigation information. The second pathway is similar to the process described in [Miller and Campbell, 2006]. This section describes the transformations in detail.

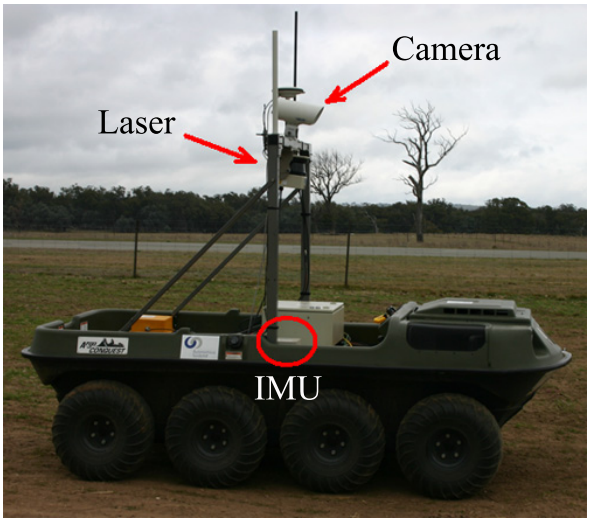

Figure 1: Photo of autonomous ground vehicle showing locations of laser scanner, colour camera and Inertial Measurement Unit (IMU)

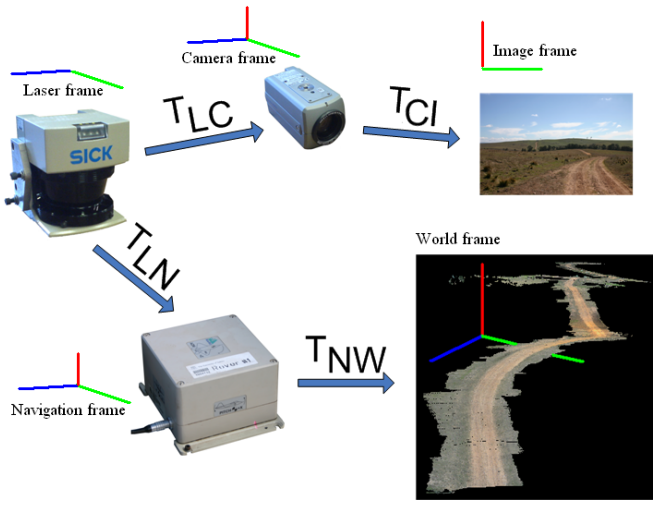

Figure 2: Data Transformations: Pictures represent different coordinate frames, arrows represent transformations between them

3.1 General Coordinate Transforms

The coordinate transform notation and rotation order (yaw, pitch, roll) from [Nettleton, 2003] is used. With sin() and cos() replaced with s and c respectively, the transformation of a point in frame a to coordinate frame b is given by:

$$\begin{aligned} C_a^b &= \begin{bmatrix} c\psi c\theta & c\psi s\theta s\phi - s\psi c\phi & c\psi s\theta c\phi + s\psi s\phi \\ s\psi c\theta & s\psi s\theta s\phi + c\psi c\phi & s\psi s\theta c\phi - c\psi s\phi \\ -s\theta & c\theta s\phi & c\theta s\phi \end{bmatrix} \\ P^b &= P_a^b + C_a^b P^a \end{aligned} \quad (1)$$

P^a is the original point in frame a, P^b is the transformed point in frame b, P_a^b is the offset between the two frames (the location of frame a in frame b) and C_a^b is the rotation matrix that encodes the yaw, pitch and roll needed to rotate frame b to align with frame a.

3.2 Laser to Camera

The data from the laser is collected in polar coordinates and transformed to cartesian coordinates aligned and centred at the laser device. It is transformed to the camera centered cartesian frame using:

$$P^C = P_L^C + C_L^C P^L \quad (2)$$

represented by T_{LC} in Figure 2. This requires that the offset P_L^C and rotation C_L^C from the camera to the laser are known and although it is possible to accurately measure the offset, it is often practically impossible to measure the rotation. For this reason, a data centric calibration technique such as [Pless and Zhang, 2004] is necessary. This technique uses observations of a planar chess board in both the laser and camera data to perform an optimisation to obtain the rotation and offset values.

With the data aligned to the cartesian camera frame, it must be transformed into the camera image frame. The standard pinhole model as presented in [Pless and Zhang, 2004] and [Stoyanov, 2006] is used. This transformation requires that the intrinsic properties of the camera (focal length $f = [f_1, f_2]$, principal point $p = [p_1, p_2]$ and skew coefficient α) be known and these should be determined by calibrating the camera with a technique such as [Stoyanov, 2006]. The transform is a two stage process. First the point P^C is normalised by the depth axis (the axis that represents the concept of further into the image) to produce $P^{C'}$:

$$\begin{aligned} P^C &= \begin{bmatrix} X & Y & Z \end{bmatrix}^T \\ P^{C'} &= \begin{bmatrix} \frac{X}{Z} & \frac{Y}{Z} & 1 \end{bmatrix}^T \end{aligned} \quad (3)$$

Figure 3: Laser scan drawn in image frame: Illustrates calibration quality of offset and rotation between camera and laser frames

The normalised $P^{C'}$ is then transformed to an image pixel location $P^I = [u, v, 1]^T$ using:

$$\begin{aligned} K &= \begin{bmatrix} f_1 & \alpha f_2 & p_1 \\ 0 & f_2 & p_2 \\ 0 & 0 & 1 \end{bmatrix} \\ P^I &= K P^{C'} \end{aligned} \quad (4)$$

where the matrix K is formed using the intrinsic camera parameters mentioned above.

Figure 3 shows points from a single laser scan transformed with T_{LC} and T_{CI} into an image. This can be used to inspect the quality of the calibration results as there should be no visibly incorrect scan points at the sharp edges in the image.

3.3 Laser to Navigation

The navigation solution is provided for a particular point located somewhere on the platform which defines the navigation frame. If the laser is not physically located and aligned at the centre of this frame then the data must be transformed to this location. This is a standard transformation:

$$P^N = P_L^N + C_L^N P^L \quad (5)$$

and is represented by T_{LN} in Figure 2.

Similar to the transformation between the cartesian laser and camera frames, this transformation requires the offset and rotation between the laser and navigation frame on the platform. Although the offset can easily be measured, measuring the rotation accurately is physically much more difficult. The complete transforma-

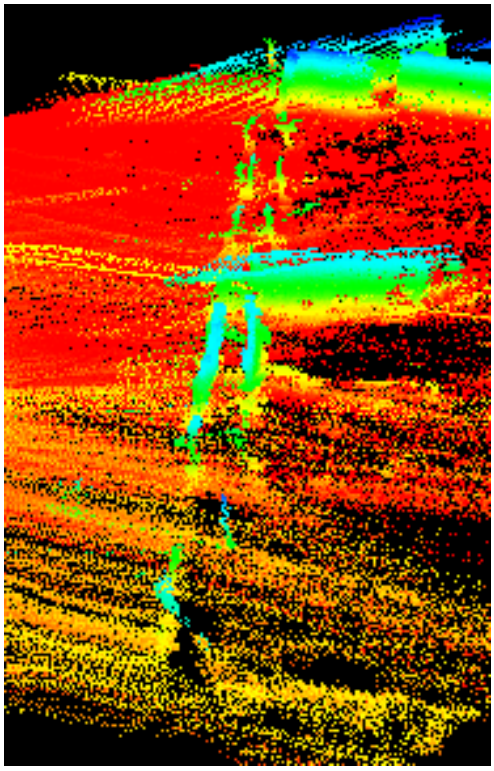

Figure 4: Yaw between laser and navigation system uncalibrated. Scan of fence single fence line appears duplicated in map

tion should be determined from a data centric calibration technique.

Figures 4 and 5 show data (coloured by height for contrast) obtained by driving the vehicle adjacent to a fence line with multiple passes at different headings. Figure 4 shows the result of incorrect measurement of the yaw between the laser and the navigation frame and Figure 5 shows the result after correction. Before correction it can be seen that the individual posts appear to be leaning, and the whole fence line is duplicated. The correction was done by applying the data centric calibration technique of [Underwood *et al.*, 2007].

3.4 Navigation To World

The transformation from the navigation frame to the global coordinate frame (T_{NW} in Figure 2) is given by:

$$P^W = P_N^W + C_N^W P^N \tag{6}$$

The attitude of the vehicle provides the yaw, pitch and roll for C_N^W and the position (given by the right handed coordinate frame of northings, eastings, and down) provides the offset P_N^W. The covariance matrix representing the uncertainty of all of these measurements is also obtained and used in Section 4 to maintain the uncertainty of the map.

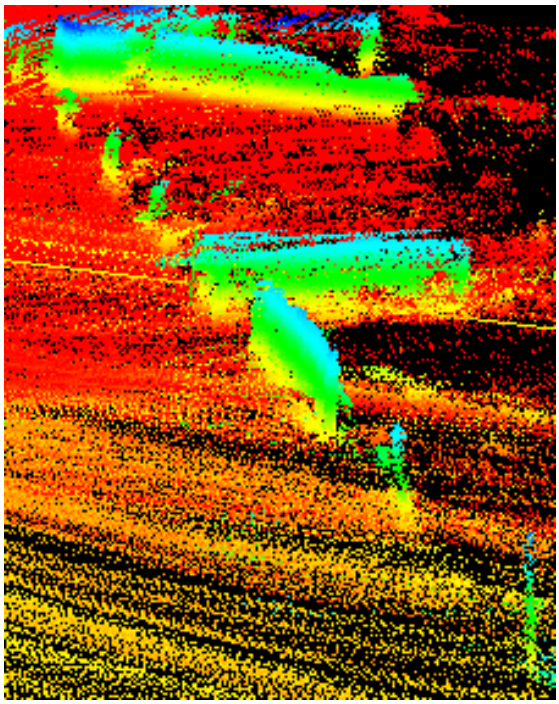

Figure 5: Yaw between laser and navigation system calibrated correctly. Scan of single fence line appears as a single focussed fence line in map

4 Tracking Uncertainty through Transformations

Uncertainty is present at several stages in the map building process. There is uncertainty in the pose measurements from the navigation system and in the measurements given by the scanning laser and colour camera. There is also uncertainty in the rigid translation and rotation offsets due to calibration error.

Uncertainty in the dynamic transform T_{NW} from navigation to world can be maintained, provided that the navigation system tracks pose uncertainty. It must provide a covariance matrix describing the uncertainty of position and attitude. The uncertainty from the laser can also be incorporated by using a sensor model that accounts for the measurement error in both range and bearing (often modeled as two independent gaussian distributions centred at the measured values).

For the rigid transformations T_{LC}, T_{CI}, T_{LN}, it is possible to provide bounds on the calibrated offsets and rotation angles but the error is constant. Unlike the uncertainty in the dynamic transformation, it is not valid to use these bounds as they form a constant bias.

The camera uncertainty for this system is highly non-linear due to the non-zero offset between the bearing-only camera and the laser. Even if perfect calibration were achieved there is the potential for misalignment of

range and colour information due to occlusion. In cases where the laser beam passes near an object which is an obstruction from the camera perspective, the colour from the object in the foreground can be attributed to the location of the background. This problem is reduced by physically minimising the offset between the camera and laser. Colour errors due to this phenomenon have been observed empirically but are rare.

The following describes the process of linearising the transformations from the range scanner through the navigation system to the world coordinate frame. In this way the measurable uncertainty in the navigation pose and the range sensor can be convolved to produce a three dimensional gaussian representing the uncertainty of each point in the global colour map. The process is described in detail for the two dimensional case in [Leal, 2003] and is extended here for three dimensions. The following equation combines T_{LN} and T_{NW}, together with the initial range scanner polar to cartesian conversion into one equation representing the complete transform from laser to world coordinates:

$$\begin{aligned} P^L &= \begin{bmatrix} R_l \times cos(\theta_l) \\ R_l \times sin(\theta_l) \\ 0 \end{bmatrix} \\ P^W &= P_N^W + C_N^W \times (P_L^N + C_L^N P^L) \end{aligned} \tag{7}$$

The Jacobian matrix J is calculated for the three components of P^W (x,y,z) with respect to the eight terms for which uncertainty is known:

$$J = \begin{bmatrix} \frac{\partial x}{\partial n} & \frac{\partial x}{\partial e} & \frac{\partial x}{\partial d} & \frac{\partial x}{\partial y} & \frac{\partial x}{\partial p} & \frac{\partial x}{\partial r} & \frac{\partial x}{\partial R_l} & \frac{\partial x}{\partial \theta_l} \\ \frac{\partial y}{\partial n} & \frac{\partial y}{\partial e} & \frac{\partial y}{\partial d} & \frac{\partial y}{\partial y} & \frac{\partial y}{\partial p} & \frac{\partial y}{\partial r} & \frac{\partial y}{\partial R_l} & \frac{\partial y}{\partial \theta_l} \\ \frac{\partial z}{\partial n} & \frac{\partial z}{\partial e} & \frac{\partial z}{\partial d} & \frac{\partial z}{\partial y} & \frac{\partial z}{\partial p} & \frac{\partial z}{\partial r} & \frac{\partial z}{\partial R_l} & \frac{\partial z}{\partial \theta_l} \end{bmatrix} \tag{8}$$

In the order used above, the variables are the northings, eastings, down, yaw, pitch and roll of the vehicle (n,e,d,y,p,r) and the range and bearing of the laser return (R_l, θ).

The Jacobian can be formed by calculating the partial derivatives numerically for every new point, or for simplicity and speed of the final implementation, equation 7 can be expanded and analytically differentiated. The complete expanded Jacobian is too large to print, but the calculation is trivial. The covariance matrix Q representing the system uncertainty is formed:

$$Q = \begin{bmatrix} \sigma_{nn}^2 & \sigma_{ne}^2 & \sigma_{nd}^2 & \sigma_{ny}^2 & \sigma_{np}^2 & \sigma_{nr}^2 & 0 & 0 \\ \sigma_{en}^2 & \sigma_{ee}^2 & \sigma_{ed}^2 & \sigma_{ey}^2 & \sigma_{ep}^2 & \sigma_{er}^2 & 0 & 0 \\ \sigma_{dn}^2 & \sigma_{de}^2 & \sigma_{dd}^2 & \sigma_{dy}^2 & \sigma_{dp}^2 & \sigma_{dr}^2 & 0 & 0 \\ \sigma_{yn}^2 & \sigma_{ye}^2 & \sigma_{yd}^2 & \sigma_{yy}^2 & \sigma_{yp}^2 & \sigma_{yr}^2 & 0 & 0 \\ \sigma_{pn}^2 & \sigma_{pe}^2 & \sigma_{pd}^2 & \sigma_{py}^2 & \sigma_{pp}^2 & \sigma_{pr}^2 & 0 & 0 \\ \sigma_{rn}^2 & \sigma_{re}^2 & \sigma_{rd}^2 & \sigma_{ry}^2 & \sigma_{rp}^2 & \sigma_{rr}^2 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & \sigma_{R_l R_l}^2 & \sigma_{R_l \theta_l}^2 \\ 0 & 0 & 0 & 0 & 0 & 0 & \sigma_{\theta_l R_l}^2 & \sigma_{\theta_l \theta_l}^2 \end{bmatrix} \tag{9}$$

The values for the upper left block are obtained directly from the navigation system. The values for the lower right block are from the laser sensor model. The final three dimensional gaussian point uncertainty Σ is calculated using:

$$\Sigma = JQJ^T \tag{10}$$

The three dimensional linearised uncertainty for the world location P^W of each point is assumed to have:

$$P^W \sim N(\mu, \mathbf{\Sigma}) \tag{11}$$

5 Results

The sensors were mounted on an autonomous ground vehicle and driven along a dirt road. All of the data produced is stored as a list of points with mean location P^W, uncertainty covariance matrices Σ, and colour. For visualisation, the mean location of each point is displayed and the point is either coloured by the stored colour or the largest orthogonal standard deviation of the point uncertainty in metres:

$$\sigma = \sqrt{max(eigenvalue(\Sigma))} \tag{12}$$

Figure 6 shows a map of a road coloured by uncertainty. In the lower left corner of the image where the vehicle started, the largest standard deviation is only 1.7cm. As the vehicle descended down the road, radio contact was lost with the Differential GPS (DGPS) station and the navigation uncertainty increased. This is reflected in the uncertainty of the mapped points and the maximum standard deviation is nearly 85cm. Figure 7 shows a closer view of a section where the navigation uncertainty increased. In this figure, a fence can clearly be seen, but as the uncertainty becomes larger, there is a reduction in quality of the reconstruction of the fence as displayed by the mean point locations. Although in some regions the quality of the data is inevitably low, the accuracy is measurable for every individual point.

Situations like this are very common because the quality of most navigation solutions is highly variable. Maintaining uncertainty in the map is therefore critical for

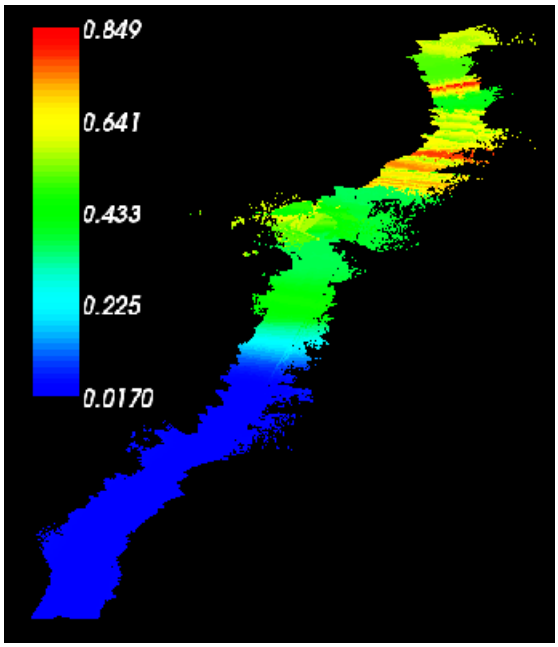

Figure 6: Map of dirt road coloured by uncertainty as largest orthogonal standard deviation in metres, calculated by equation 12.

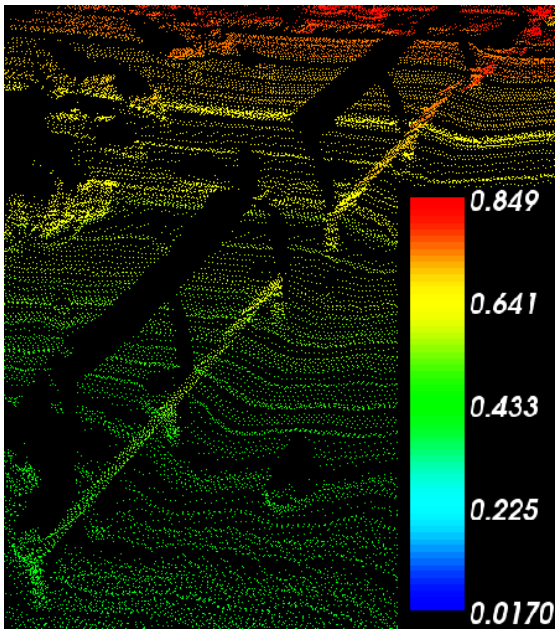

Figure 7: View of road-side fence coloured by uncertainty as largest orthogonal standard deviation in metres, calculated by equation 12.

any downstream process that uses the information including obstacle avoidance, region classification or map fusion, where techniques use uncertainty to determine the confidence that can be placed in the data.

To quantitatively assess the point uncertainty, a reference data set is compared to a set of overlapping sample data using the iterative closest point (ICP) algorithm [Besl and McKay, 1992]. The reference sets were obtained by traversing a route in one direction (shown in Figure 6) and splitting the data into seven consecutive subsets. Overlapping sample sets were obtained by driving back along a similar route. For this data, the transformation produced by ICP is most accurate in the vertical (Z) axis. The horizontal convergence is weaker due to the lack of salient vertical features. To compare the ICP results with the point uncertainty, Equation 12 is calculated for each point in the matching sample and reference data sets and averaged to produce σ_{sample} and σ_{ref}. These standard deviations are convolved using Equation 13 to produce σ_{total}, which reflects the combined uncertainty of the reference and sample data. The offsets produced by the ICP algorithm ($\triangle Z_{icp}$) correspond to measurements of a random variable with a gaussian distribution with a standard deviation of σ_{total}. The results are shown in Table 1. The ICP offsets are normalised with respect to σ_{total} to produce $\triangle Z_{norm}$. This value is the number of standard deviations of error of the ICP offset given the reference and sample point cloud uncertainty. Although the number of samples is small, the offsets do agree with the transformed uncertainty.

$$\sigma_{total} = \sqrt{\sigma_{ref}^2 + \sigma_{sample}^2} \tag{13}$$

Figures 8 and 9 show a three dimensional colour map taken from a single traversal of a road. Figure 9 shows a detailed view of a particular section of the map, and a camera image of the same region is shown in Figure 10. The quality of the colour mapping is evident from visual inspection. The accuracy of the three dimensional structure is sufficient to show the individual bars of the gate and surrounding posts, and the colouring accuracy is high enough to colour them appropriately. The colour mapping also preserves the visual texture of the ground as seen in the camera. (Note the tyre patterns on the road in Figure 8 and the gravel in Figure 9, and the colour and texture detail preserved in the vehicle point clouds in Figures 11 and 12).

Table 1: Comparison of ICP and Transformed Point Uncertainty

	*Sample*1	*Sample*2	*Sample*3	*Sample*4	*Sample*5	*Sample*6	*Sample*7
σ_{sample}	0.3385	0.4295	0.5292	0.6491	0.4730	0.6424	0.7139
σ_{ref}	0.5428	0.5927	0.5797	0.4148	0.3125	0.0607	0.0574
σ_{total}	0.6397	0.7320	0.7849	0.7703	0.5669	0.6453	0.7162
$\triangle Z_{icp}$	1.1889	1.3020	0.4922	0.2095	0.0985	0.7455	0.6106
$\triangle Z_{norm}$	1.8585	1.7788	0.6271	0.2720	0.1739	1.1553	0.8526

Figure 8: Colour map of dirt road. Mean locations plotted as a point cloud with colour. Colour mapping preserves visual texture of the ground. Only points in the field of view of both the laser and camera are shown.

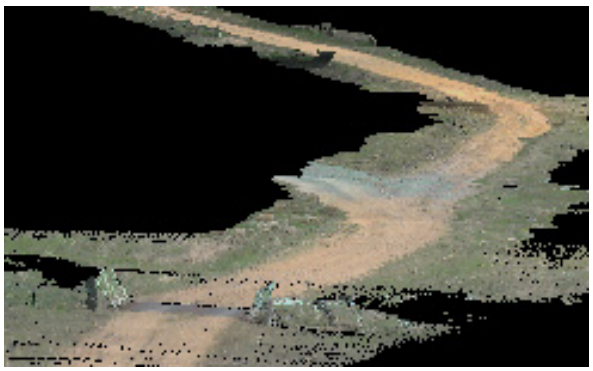

Figure 9: Colour map of dirt road and gate. Mean locations plotted as point cloud with colour. View also visible in camera image in Figure 10

Figure 10: View of roadside gate from colour camera, also visible in three dimensional map in Figure 9.

Figure 11: Three dimensional colour point cloud of one side of a vehicle and surrounding tufts of grass.

Figure 12: Three dimensional colour point cloud of one side of a vehicle.

6 Conclusion

This paper has combined a set of transformations to fuse colour camera, range scanning and navigation data to produce three dimensional colour point clouds in real-time. The results show that the point clouds are accurate when the navigation system is working optimally but that the quality is dependent on having good navigation information. Maintaining the measurement uncertainty in the whole system is critical if any reliable processing, fusion or visualisation of the data is required, and this technique provides a mechanism for calculating the uncertainty due to the data gathering process. This paper has shown that very high quality three dimensional colour maps can be built in real-time, with a relatively simple approach.

7 Future Work

Point cloud matching algorithms such as ICP typically optimise the transform between two clouds to minimise the Euclidean distance between nearby points. Future work will determine whether the uncertainty of the individual points can be explicitly incorporated into the point cloud matching optimisation, and whether this can reduce the processing time, or increase the accuracy of the result. In addition, if colour information is available, it could be used to improve the association of points between different data sets, rather than relying on distance only.

8 Acknowledgements

This work is supported by the ARC Centre of Excellence programme, funded by the Australian Research Council (ARC) and the New South Wales State Government.

References

[Abmayr *et al.*, 2004] T. Abmayr, F. Härtl, M. Mettenleiter, I. Heinz, A. Hildebrand, B. Neumann, and C. Fröhlich. Realistic 3D Reconstruction Combining Laserscan Data with RGB Color Information. *Information Sciences*, 35(Part B):198–203, 2004.

[Besl and McKay, 1992] P. J. Besl and N. D. McKay. A method for registration of 3-D shapes. *IEEE Trans. Pattern Analasys and Machine Intelligence*, 14:239–256, 1992.

[Chandran and Newman, 2006] M. Chandran and P. Newman. Motion Estimation from Map Quality with Millimeter Wave Radar. *Intelligent Robots and Systems, 2006 IEEE/RSJ International Conference on*, pages 808–813, 2006.

[Durrant-Whyte and Scheding, 2003] H. Durrant-Whyte and S. Scheding. A model for machine perception in natural environments. *International Symposium of Robotics Research (ISSR)*, 2003.

[Frueh *et al.*, 2004] C. Frueh, R. Sammon, and A. Zakhor. Automated texture mapping of 3D city models with oblique aerial imagery. *3D Data Processing, Visualization and Transmission, 2004. 3DPVT 2004. Proceedings. 2nd International Symposium on*, pages 396–403, 2004.

[Leal, 2003] J. Leal. *Stochastic Environment Representation.* PhD thesis, The University of Sydney, January 2003.

[Miller and Campbell, 2006] Isaac Miller and Mark Campbell. A Mixture-Model Based Algorithm for Real-Time Terrain Estimation. *Journal of Field Robotics*, pages 755–775, September 2006.

[Nettleton, 2003] E. Nettleton. *Decentralised Architecture for Tracking and Navigation with Multiple Flight Vehicles.* PhD thesis, The University of Sydney, 2003.

[Ohno and Tadokoro, 2005] K. Ohno and S. Tadokoro. Dense 3D map building based on LRF data and color image fusion. *Proceedings of IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 2792–2797, 2005.

[Pfaff and Burgard, 2005] P. Pfaff and W. Burgard. An Efficient Extension of Elevation Maps for Outdoor Terrain Mapping. *Proceedings of the International Conference on Field and Service Robotics (FSR)*, pages 165–176, 2005.

[Pless and Zhang, 2004] R. Pless and Q. Zhang. Extrinsic calibration of a camera and laser range finder. *IEEE International Conference on Intelligent Robots and Systems (IROS)*, 2004.

[Rasmussen, 2002] C. Rasmussen. Combining laser range, color, and texture cues for autonomous road following. *Proceedings of IEEE/RSJ International Conference on Robotics and Automation (IROS)*, 2002.

[Simmons *et al.*, 1996] R. Simmons, L. Henriksen, L. Chrisman, and G. Whelan. Obstacle Avoidance and Safeguarding for a Lunar Rover. *AIAA Forum on Advanced Developments in Space robotics*, august 1996.

[Stoyanov, 2006] D. Stoyanov. Camera Calibration Toolbox for Matlab. http://www.vision.caltech.edu /bouguetj/calib_doc/index.html#ref, 2006. Website as at 15/08/2006.

[Thrun *et al.*, 2006] S. Thrun, M. Montemerlo, H. Dahlkamp, D. Stavens, A. Aron, J. Diebel, P. Fong, J. Gale, M. Halpenny, G. Hoffmann, et al. Winning the darpa grand challenge. *Journal of Field Robotics*, 2006.

[Underwood *et al.*, 2007] J. Underwood, A. Hill, and S. Scheding. Calibration of Range Sensor Pose on Mobile Platforms. *IEEE International Conference on Intelligent Robots and Systems (IROS)*, 2007.