2010 IEEE International Conference on Robotics and Automation
Anchorage Convention District
May 3-8, 2010, Anchorage, Alaska, USA

Global Rover Localization by Matching Lidar and Orbital 3D Maps

Patrick J.F. Carle and Timothy D. Barfoot
Institute for Aerospace Studies
University of Toronto
Toronto, Canada M3H 5T6
{pat.carle, tim.barfoot}@utoronto.ca

Abstract— **Current rover localization techniques such as visual odometry have proven to be very effective on short to medium-length traverses (e.g., up to a few kilometres). This paper deals with the problem of long-range rover localization (e.g., 10km and up). An autonomous method to globally localize a rover is proposed by matching features detected from a 3D orbital elevation map and rover-based 3D lidar scans. The accuracy and efficiency of the algorithm is enhanced with visual odometry, and inclinometer/sun-sensor orientation measurements. The methodology was tested with real data, including 37 lidar scans of terrain from a Mars-Moon analogue site on Devon Island, Nunavut. When a scan contained a sufficient number of good topographic features, localization produced position errors of no more than 100m, and as low as a few metres in many cases. On a 10km traverse, the developed algorithm's localization estimates were shown to significantly outperform visual odometry estimates. It is believed that this architecture could be used to accurately and autonomously localize a rover on long-range traverses.**

I. INTRODUCTION

The ongoing Mars Exploration Rover (MER) missions have proven to be historic landmarks in space exploration. However, they are also humbling reminders of the challenges ahead. For example, the MER Opportunity has operated on Mars for over 5 years now, but has only driven a total of about 20km due to mechanical/energy limitations and a lack of autonomy [1]. An important goal for future generations of rovers will be to overcome these deficiencies to allow them to explore sites hundreds of kilometres away from their landers [2]. Rovers will consequently require an autonomous long-range localization system to aid them in their journey.

Currently, a rover employs a variety of techniques to determine its pose at any given time. The MERs were first localized with radio tracking [3], descent trajectory modeling, and by comparing orbital to ground camera imagery [1]. After leaving their landers, localization has been accomplished primarily with dead-reckoning techniques such as wheel odometry, visual odometry (VO) and local bundle adjustment (BA). Wheel odometry is not computationally intensive, but is vulnerable to sensor noise and mechanical disturbances (e.g., wheel slippage) [4]. Computer vision techniques, such as VO and BA, complement wheel odometry when needed.

VO is automated and can work in real-time, but is computationally very demanding. It has yielded impressive results in the past with error as low as 0.1% over a 10km traverse [5]. BA can offer further gains in accuracy [6], but efforts to automate the process are ongoing [7]. Despite significant advances in the technology, such dead-reckoning approaches are not suitable for long-range localization (e.g., more than 10km) since they will always exhibit unbounded error growth with distance traversed [8].

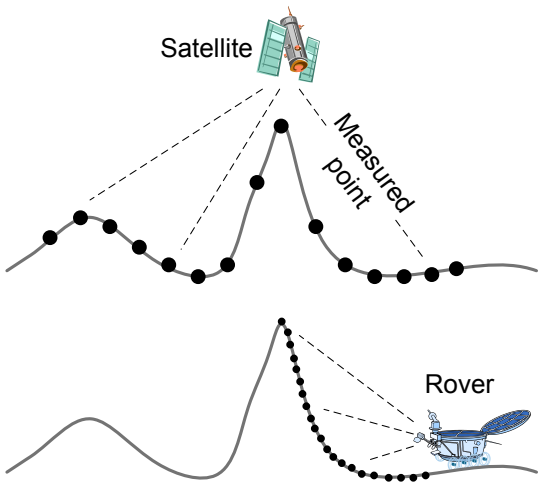

Fig. 1. Matching rover's 3D map (bottom) to a 3D orbital map (top).

Global localization techniques can be used to correct dead-reckoning pose estimates once these become unreliable. On Earth, the Global Positioning System (GPS) is commonly used for this purpose. However, the satellite infrastructure required for such a system is not feasible for non-Earth applications. This paper proposes an alternate solution that aligns a rover-based three-dimensional (3D) *local* map to a satellite-based 3D *global* map, as shown in Figure 1.

In this research, the local map is a point cloud obtained from a time-of-flight lidar (*Li*ght *D*etection *a*nd *R*anging). This instrument can measure distance to far-away objects by rapidly firing a laser and measuring the time for reflected beams to return. In a surveying configuration, a lidar can sample terrain with centimetre-accuracy at a range of up to 1.5km, making it a vital guidance and navigation sensor. Lidars have been tested in numerous applications on Earth [9], [10] and in space [11], [12].

The global map may be acquired from a satellite-based laser altimeter (e.g., LOLA, MOLA2), or by extracting 3D information from a stereo pair of high-resolution satellite images (e.g., LROC, HiRISE) [13]. Current satellites have extensive and accurate coverage of the Moon and Mars (see

Table I). Therefore, relevant map data could be loaded into a rover even before it begins its mission. This would allow the rover to autonomously localize without requiring any additional map data from Earth.

TABLE I
SUMMARY OF SATELLITE MAPS FOR MARS AND THE MOON.

Target	Instrument	Coverage	Horiz. Res.	Vert. Res.
Moon	LOLA [14]	Total	50-100m	10cm
	LROC [14]	>10%	50cm	NA
Mars	MOLA2 [15]	Total	100m	1m
	HiRISE [16], [17]	2%	<1m	<1m

Having obtained local and global maps, a means of matching them is then required. In feature-based matching, interest points are first extracted from the global and local maps, and then matched in search of global-local feature correspondences. If at least three, non-collinear feature correspondences are found, a rigid 3D transformation can be expressed that aligns the two maps [18].

The most prominent features visible to both a low-resolution orbital map and a high-resolution lidar map are topographic peaks. Since one side of a peak will always be hidden from the rover, features from the lidar map will only be partially describable as shown in Figure 1. It will therefore be difficult to compare features between the two maps if a descriptor-based approach is used, such as spin-image matching [19]. A more convenient alternative is to search for similar *feature constellations*, where the spacing between features now effectively acts as the descriptor. This method is known as the Data-Aligned Rigidity-Constrained Exhaustive Search (DARCES) [20].

Previous work has investigated these 3D matching techniques in the context of localization [10], [9]. Localization has been attempted with VO and BA on a network of 2D images and orbital data with some success [7]. However, automation is problematic. The VIPER algorithm [21] matches the skyline in a panoramic image to predicted skylines at various positions on a 3D global map. The final solution is not very accurate (>100m).

This paper investigates a novel global localization algorithm using 3D data. The overall architecture is outlined in Section II, followed by a description of the methodology in Sections III through V. Sections VI to VIII present results from field tests at a Mars/Moon analogue site in the Canadian Arctic.

II. ARCHITECTURE

The architecture is developed for the general case where a rover traverses over some distance and occasionally stops to scan the terrain with the lidar as shown in Figure 2. The general formulation can also be simplified by omitting odometry to examine single-scan localization.

The goal is to determine the rover's pose at each scan site with respect to the global map's reference frame, $\underrightarrow{\mathcal{F}}_o$. The rover pose at scan site ℓ is defined as a transformation, $\mathbf{T}_{\ell o} := \{\mathbf{t}_{\ell o}, \mathbf{S}_{\ell o}\}$, from $\underrightarrow{\mathcal{F}}_o$ to the rover's local frame $\underrightarrow{\mathcal{F}}_\ell$, where $\mathbf{t}_{\ell o}$ and $\mathbf{S}_{\ell o}$ are, respectively, the translation vector and rotation matrix from $\underrightarrow{\mathcal{F}}_o$ to $\underrightarrow{\mathcal{F}}_\ell$. An overview of the procedure is presented in Figure 3 and is summarized as follows:

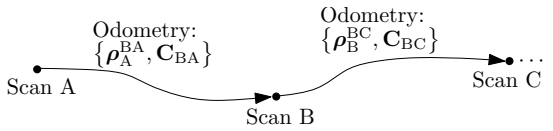

Fig. 2. Rover traverses to lidar scan sites collecting odometry along the way (e.g., the odometry-measured translation and rotation from site A to B are, respectively, $\boldsymbol{\rho}_{\mathrm{A}}^{\mathrm{BA}}$ and $\mathbf{C}_{\mathrm{BA}}$.)

(a) **Feature Detection**: Features are detected from the global and local maps.
(b) **Feature Matching**: DARCES obtains global-local feature correspondences and an initial estimate of the rover's poses.
(c) **Pose Refinement**: Global-local feature correspondences are used in combination with orientation and odometry measurements to refine rover pose estimates.

The sections to follow discuss these in further detail.

III. FEATURE DETECTION

Features must first be detected (box (a) in Figure 3) from the global map (the *global features*) and from the local map (the *local features*). It is assumed a global elevation map is given with resolution L_{xy}, as well as one or more local elevation maps. Generally, the most prominent features common to both maps are topographic peaks. These peaks are detected using a local maxima detector based on morphological dilation[1] [22], [23]. Other feature detectors could also be used within the presented framework.

The scan is leveled using pitch and roll measurements (e.g., inclinometer). This ensures the $+z$-direction of the global and local maps roughly coincide. The local map is gridded to the global map resolution with nearest-neighbour interpolation. This speeds up the detection process and ensures the scale of global and local features is the same.

Morphological dilation replaces lower grid values with neighbouring higher grid values, effectively blurring out low elevations. Once dilation is completed for all points on the grid, the blurred map is compared to the original map. Cells with no change in value are interpreted as local maxima. The dilation window is chosen to be a pixelated circle (Figure 4) to make the window's coverage more uniform in all directions. The radius of this circle limits the size of the detected features, as well as the distance between features. The minimum distance between features, D_{detect}, depends on the global map resolution, L_{xy}, such that $D_{\text{detect}} := n \times L_{xy}$ where n is the circle's cell radius.

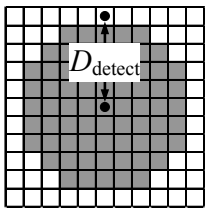

Fig. 4. Pixelated circle with $n = 5$.

[1]This was inspired by code found on the Matlab Central repository as 'localMaximum.m' by Yonathan Nativ.

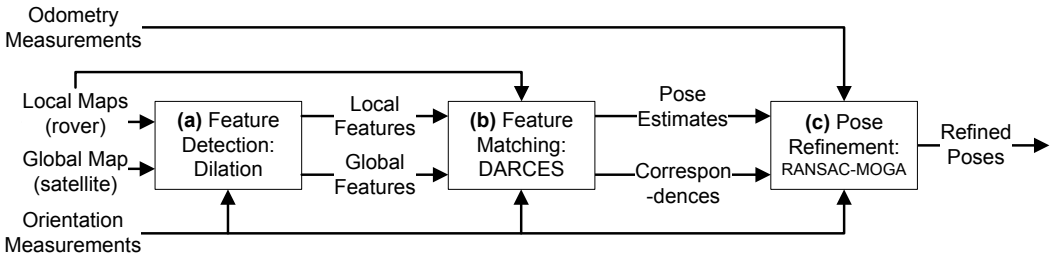

Fig. 3. Overall architecture.

Features must then undergo final processing to eliminate false peaks that might be detected in flat areas. This is done by ensuring the minimum distance between any two features is D_{detect} and by testing areas for flatness. Sample feature detection results are shown in Figure 5.

A good uncertainty model is necessary to correctly assess the quality of individual features and global-local matches. Global feature positional uncertainty is assumed equal to the position uncertainty of a measured 3D point in the global map. The positional uncertainty of a measured 3D point from the lidar is comparatively very small (<1m) even for measurements far from the lidar's origin. However, the effect of occlusions will dominate the uncertainty in local feature measurements. Local feature uncertainty is therefore better estimated by the size of a local feature, D_{detect}.

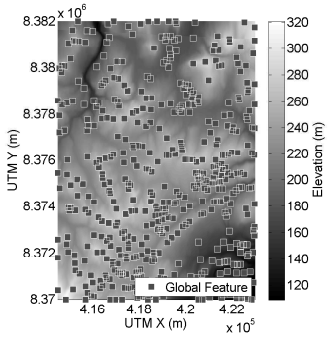

Fig. 5. Detected global features plotted on global map.

IV. FEATURE MATCHING

The feature matching methodology (box (b) in Figure 3) is based on the DARCES algorithm [20]. Lidar scans are processed individually.

A. Hypothesis Search

Hypothesized correspondences must first be generated between global and local features. A *hypothesis* is defined as a group of possible correspondences between three unique local features, called *control points*, and three unique global features. A hypothesis is not guaranteed to be correct due to noise in feature position measurements. Therefore, many control point groups are tested to increase the chances of finding a valid hypothesis.

A hypothesis is generated if the distances between three global features are similar to the distances between three local features. This test can be quantified knowing each feature's position uncertainty.

B. Hypothesis Evaluation

The validity of a hypothesis is evaluated based on the transformation it produces between the global and local frame. For a hypothesis i, the transformation from $\underrightarrow{\mathscr{F}}_o$ to $\underrightarrow{\mathscr{F}}_\ell$, $\mathbf{T}^i_{\ell o} := \{\mathbf{t}^i_{\ell o}, \mathbf{S}^i_{\ell o}\}$, can be obtained using a least-squares point-alignment algorithm to align the three global and three local features that comprise a hypothesis [24].

To improve the efficiency and robustness of DARCES, hypotheses are first screened with a number of simple tests. A map-boundary test rejects hypotheses that estimate the rover's position outside the global map. A z-deviation test ensures that an estimated z-position is within E_{zdev} of the global map's elevation at the estimated xy-position. Finally, orientation tests check that estimates and measurements (e.g., inclinometer, sun sensor) of roll, α, pitch, β, and heading, γ, differ by no more than E_α, E_β and E_γ. These thresholds are set to three standard deviations of the involved measurement's uncertainty to ensure that a high proportion of valid hypotheses fall within the thresholds. For example, $E_{\text{zdev}} := 3\sigma_{r_{G_z}}$, where $\sigma_{r_{G_z}}$ is one standard deviation of the global map z-position uncertainty. These tests are not theoretically necessary, but are practically very helpful in speeding up the search for valid hypotheses.

A measure of fitness is then calculated for each hypothesis by examining how well the hypothesized transformation aligns lidar scan points to the global map. To more efficiently and robustly calculate this metric, the full lidar scan is decimated to half the global map resolution, $L_{xy}/2$. These more evenly-spread points are called the *reference points*.

For a hypothesis i, the fitness metric, f_i, is the average absolute z-error between the reference points, transformed to the global frame with $\mathbf{T}^i_{\ell o}$, and the global map:

$$f_i := -\frac{1}{N_{\text{ref}}} \sum_{j=1}^{N_{\text{ref}}} \left| z^R_{i,j} - z^G_{i,j} \right|, \tag{1}$$

where N_{ref} is the number of reference points, $z^R_{i,j}$ is the z-position of the transformed reference point j in the global frame, and $z^G_{i,j}$ is the interpolated global map elevation at the xy-position of the transformed reference point j. The negative is applied so that a low error corresponds to a high fitness.

C. Hypothesis Selection

Once a fitness is associated with each hypothesis, a search is made for valid hypotheses, which are defined as a group of

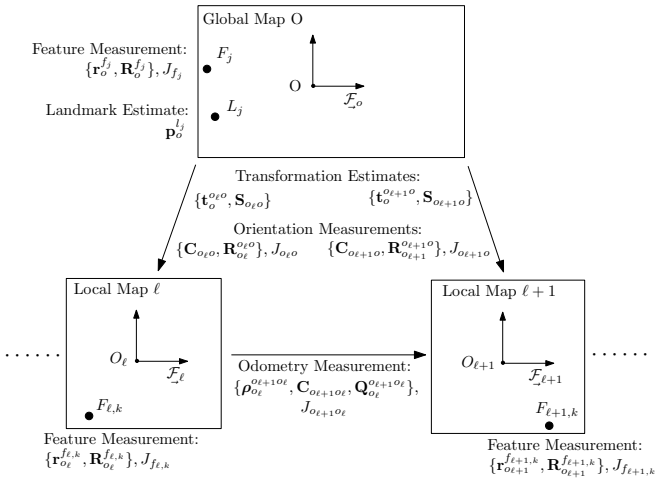

Fig. 6. The M-frame alignment problem.

high-fitness hypotheses that have similar position estimates. The hypothesis with the highest fitness in this valid group is then returned along with its associated transformation estimate and global-local feature correspondences. If no valid group is found after all combinations of control points have been exhausted, then DARCES returns no solution.

V. POSE REFINEMENT

In the pose refinement stage, outlier feature correspondences are first rejected with RANSAC (*Ra*ndom *Sa*mple *C*onsensus) [25]. The remaining, inlier feature correspondences are then combined with odometry and orientation measurements into a Simultaneous Localization And Mapping (SLAM) problem to refine pose estimates (box (c) in Figure 3). The Multi-frame Odometry-compensated Global Alignment (MOGA) algorithm, solves this problem by minimizing the errors between all available measurements and desired estimates. This is essentially a batch SLAM algorithm that fuses relative and absolute pose measurements over an entire rover traverse. MOGA will be summarized here, but full details can be found from [26].

The main goal is to estimate the transformations from the global frame, $\underrightarrow{\mathcal{F}}_o$, to each of the M local frames, $\underrightarrow{\mathcal{F}}_\ell$. This M-frame alignment problem is depicted in Figure 6. Input measurements are assumed to be corrupted with white, zero-mean, Gaussian noise. The four types of measurement are global feature positions $\left(\mathbf{r}_o^{f_j}\right)$, local feature positions $\left(\mathbf{r}_{o_\ell}^{f_{\ell,k}}\right)$, odometry measurements of the rotation $\left(\mathbf{C}_{o_{\ell+1}o_\ell}\right)$ and translation $\left(\boldsymbol{\rho}_{o_\ell}^{o_{\ell+1}o_\ell}\right)$ between adjacent local frames, and orientation measurements $\left(\mathbf{C}_{o_\ell o}\right)$. The covariance matrices associated with these measurements are, respectively, $\mathbf{R}_o^{f_j}$, $\mathbf{R}_{o_\ell}^{f_{\ell,k}}$, $\mathbf{Q}_{o_\ell}^{o_{\ell+1}o_\ell}$, and $\mathbf{R}_{o_\ell}^{o_\ell o}$.

There are $3(2M+N)$ design parameters to be estimated. The three types are estimated rotations $\left(\mathbf{S}_{o_\ell o}\right)$, estimated translations, $\left(\mathbf{t}_o^{o_\ell o}\right)$, and estimated feature positions, $\left(\mathbf{p}_o^{l_j}\right)$.

The optimal design parameters will be obtained by minimizing the sum of squared errors between estimates and measurements. The four types of error are between estimated and globally-measured feature positions $\left(J_{f_j}\right)$, between estimated and locally-measured feature positions $\left(J_{f_{\ell,k}}\right)$, between estimated and odometry-measured frame transformations $\left(J_{o_{\ell+1}o_\ell}\right)$, and between estimated and measured orientation $\left(J_{o_\ell o}\right)$.

To allow all design parameters to be optimized simultaneously, each of these error terms must be expressed as a function of a common design parameter column, $\mathbf{z}$. The overall objective function, $J(\mathbf{z})$, is then

$$J(\mathbf{z}) := \sum_j^N J_{f_j}(\mathbf{z}) + \sum_\ell^M \sum_k^{N_\ell} J_{f_{\ell,k}}(\mathbf{z}) + \sum_\ell^{M-1} J_{o_{\ell+1}o_\ell}(\mathbf{z}) + \sum_\ell^M J_{o_\ell o}(\mathbf{z}), \tag{2}$$

where maximum likelihood is assured by weighting each term with the inverse covariance of the involved measurement. For example, the objective function contribution of global feature F_j is $J_{f_j}(\mathbf{z}) := \frac{1}{2}\mathbf{e}_{f_j}(\mathbf{z})^T \mathbf{R}_o^{f_j{}^{-1}} \mathbf{e}_{f_j}(\mathbf{z})$, where $\mathbf{e}_{f_j}(\mathbf{z})$ is the error between the estimated and globally-measured feature positions.

The optimal design parameter column, $\mathbf{z}^*$, is sought by minimizing the objective function through unconstrained optimization: $\mathbf{z}^* := \text{argmin}_{\mathbf{z}} J(\mathbf{z})$. The Gauss-Newton algorithm [27] is used to solve this nonlinear least-squares problem. When optimizing with Gauss-Newton, the objective function must be quadratic, which is not immediately the case due to the rotations involved. Therefore, it is necessary to linearize the error terms of $J(\mathbf{z})$ at each iteration and solve for $\mathbf{z}^*$, giving the direction to the minimum of the local quadratic approximation of the objective function. This solution is used to update the nominal design parameters. Convergence to a local minimum is achieved when the relative change in the objective function is less than E_{converge}.

Single-frame localization can be obtained by dropping the odometry terms, but otherwise proceeding in the same way.

VI. FIELD DATA

A realistic dataset was collected from a Mars/Moon analogue site on Devon Island, Nunavut [28] at 75°22'N and 89°41'W. In total, 37 lidar scans were used to test the single-scan localization performance of the algorithm. A long-range rover traverse was simulated by connecting 23 lidar scans with real odometry data collected by a pushcart outfitted with a variety of rover engineering sensors (sun sensor, inclinometer, stereo cameras). The resulting 10km path is comparable in length to the distance traversed by the MERs to date. The orientation of the lidar and cart were matched by imaging a reference target in both systems [26].

When localizing, the rover position was assumed to be within the bounds of a 100km^2 global map. The map's x and y resolutions were, respectively, 13m and 24m. Global feature positional uncertainty was estimated to a radius of 23m for one standard deviation. The local maps were collected with an Optech ILRIS3D-ER lidar, which had a range of about 1.5km, and a positional measurement uncertainty of less than a metre. However, occlusions in the scan increased

the uncertainty of a local feature's position to a radius of about 45m for one standard deviation.

The VO algorithm was similar to sparse feature correspondence methods [5]. Accuracy was on par with [29].

Heading measurements for A01-A25 were obtained from a sun sensor [30]. Heading for A26-A37 were computed with similar accuracy knowing the GPS and local positions of a distant target [26]. Roll and pitch were zero since the lidar was leveled before each scan. The uncertainties in heading, roll and pitch were, respectively, $\sigma_\gamma = \sigma_\alpha = \sigma_\beta = 1^\circ$.

Ground-truth *xy*-position measurements were obtained from a consumer-grade GPS. It had positional uncertainty 10m after averaging measurements for about 30 minutes.

VII. SINGLE-FRAME RESULTS

In the single-frame configuration, local frames were localized individually. There was a stochastic aspect to the DARCES algorithm since control points were randomly selected. Therefore, robustness was demonstrated by running 100 randomly-seeded trials on each of the 37 scans. The time to localize a single scan on a Matlab implementation was about 20 min., most of which was spent in DARCES.

Figure 8 shows localization performance for x and y positions. Although the algorithm outputs six-degree-of-freedom localization, only the x and y position performance is discussed since the others are assumed to be well-known from measurements (e.g., heading from the sun sensor). Most trials showed position error well below 50m.

In the 3,700 trials run, the algorithm never produced error greater than 100m. When the algorithm was unconfident about the rover's position, it produced no DARCES solution. This occurred when few features were within the lidar's range (flat areas, A10, A13), or when features were too close to the lidar (canyons, A02, A20-A23). In canyon-like settings, nearby hills occlude much of the view, leaving few features for the rover to detect (see Figure 7). Furthermore, a nearby detected peak is more likely to be a poor representation of the true peak, which would be occluded at close range. Future work could investigate how to automatically detect featureless areas when planning a path.

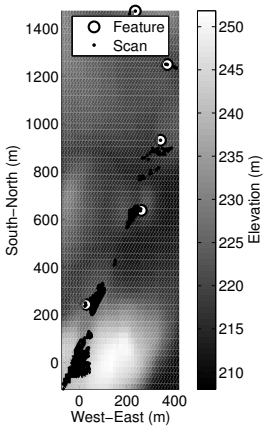

Fig. 7. Canyon scan (A02).

VIII. MULTI-FRAME RESULTS

In this configuration, VO measurements between scans A01 to A23 were available. This allowed for a single, large MOGA optimization to be executed. DARCES transformation estimates served as initial guesses for MOGA. However, if a particular frame did not have a DARCES solution, its transformation was estimated using VO to the next-closest, solved frame. As shown in Figures 9 and 10, MOGA can now produce localization estimates for frames with no DARCES solution using information from all available measurements in the chain of frames. VO estimates are provided in the figures with and without the use of absolute heading measurements (i.e., sun sensor). It should also be made clear that the VO paths made use of ground-truth for frame A01, whereas MOGA used no ground-truth whatsoever.

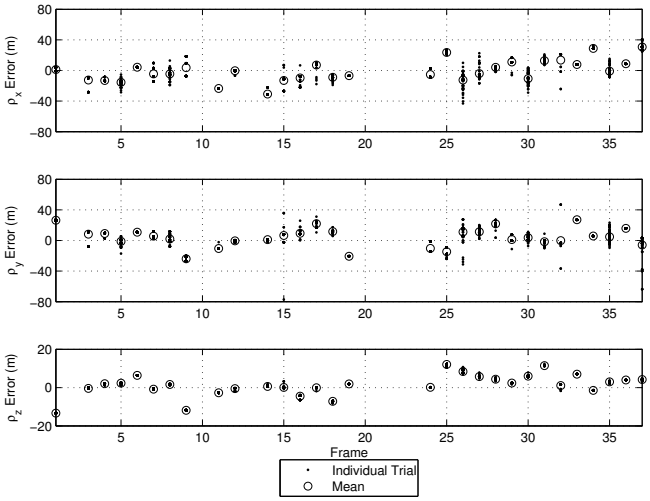

Fig. 8. Position errors for 100 single-frame localization trials on all 37 scans. Mean errors for each frame are also shown.

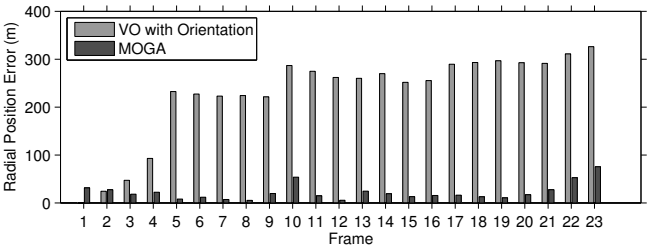

Fig. 9. MOGA, VO performance. VO given true pose for frame A01.

Note the significant improvement in the estimates as more information is given to the estimator. With absolute heading measurements, VO reduces its error from up to several kilometres to a few hundred metres. With absolute position measurements provided by matched features in MOGA, the error can be further reduced. This improvement can be attributed not only to VO, but also to feature-sharing between frames. Sharing features allows a given frame to use local features from other nearby frames with common global correspondences to improve estimates.

Frames with no DARCES solution had higher error since estimates were more dependent on odometry measurements. Despite these cases, MOGA localized a long-range traverse with no dependence on distance traversed.

IX. CONCLUSION

This research has produced a number of novel contributions. A global localization technique was developed that matches rover-based lidar scans to an orbital elevation map

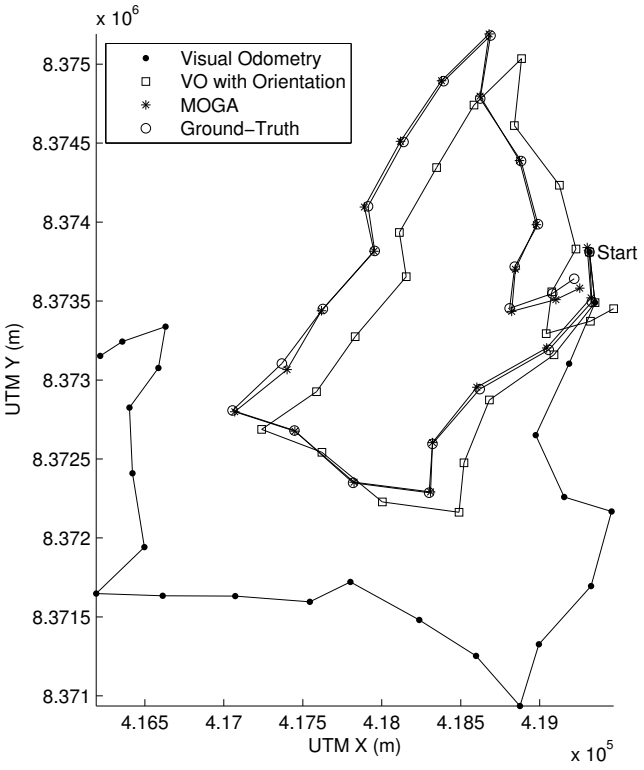

Fig. 10. VO, MOGA and ground-truth positions. Odometry was given ground-truth position and orientation for frame A01, but MOGA was not. Connecting lines do not indicate the intermediate path.

using DARCES feature constellations. MOGA, a multiple-frame, least-squares alignment technique was designed that uses feature position, orientation and odometry measurements to refine pose estimates. The architecture was also validated with a realistic Mars/Moon analogue dataset from Devon Island, Nunavut. With additional future work, it is believed that the architecture presented in this paper could be used to autonomously localize a rover over long-ranges with an accuracy comparable to GPS.

X. ACKNOWLEDGMENTS

Thanks to Paul Furgale for the visual odometry implementation and to Joseph Bakambu for helpful discussions on feature matching. Also thanks to Optech Inc., the Mars Institute, the Canadian Space Agency, the Canada Foundation for Innovation, and the Haughton Mars Project for making our field experiments on Devon Island possible.

REFERENCES

[1] R. Li, R. E. Arvidson, K. Di *et al.*, "Opportunity rover localization and topographic mapping at the landing site of Meridiani Planum, Mars," *Journal of Geophysical Research*, vol. 112, no. E2, February 2007.

[2] A. Behar, J. Matthews, F. Carsey, and J. Jones, "NASA/JPL Tumbleweed polar rover," in *Proc. IEEE Aerospace Conf.*, vol. 1, 2004.

[3] J. Guinn, "Mars surface asset positioning using in-situ radio tracking," in *Proceedings of the 11 th Annual AAS/AIAA Space Flight Mechanics Meeting, Santa Barbara, CA*, 2001, pp. 45–53.

[4] R. Li, S. W. Squyres, and R. E. Arvidson, "Initial results of rover localization and topographic mapping for the 2003 Mars Exploration Rover mission," *Photogrammetric Engineering & Remote Sensing*, vol. 71, no. 10, pp. 1129–1142, October 2005.

[5] K. Konolige, M. Agrawal, and J. Sola, "Large scale visual odometry for rough terrain," in *Proc. International Symposium on Robotics Research*, 2007.

[6] R. Li, B. A. Archinal, R. E. Arvidson *et al.*, "Spirit rover localization and topographic mapping at the landing site of Gusev Crater, Mars," *Journal of Geophysical Research*, vol. 111, no. E2, February 2006.

[7] R. Li, K. Di, and A. B. Howard, "Rock modeling and matching for autonomous long-range Mars rover localization," *Journal of Field Robotics*, vol. 24, no. 3, pp. 187–203, March 2007.

[8] C. F. Olson, L. H. Matthies, M. Schoppers, and M. W. Maimone, "Rover navigation using stereo ego-motion," *Robotics and Autonomous Systems*, vol. 43, no. 4, pp. 215–229, June 2003.

[9] N. Vandapel, R. R. Donamukkala, and M. Hebert, "Unmanned ground vehicle navigation using aerial ladar data," *International Journal of Robotics Research*, vol. 25, no. 1, pp. 31–51, January 2006.

[10] J. N. Bakambu, P. Allard, and E. Dupuis, "3D terrain modeling for rover localization and navigation," in *Proceedings of the The 3rd Canadian Conference on Computer and Robot Vision (CRV'06)*. IEEE Computer Society, 2006, p. 61.

[11] R. Shotwell, "Phoenix - the first Mars Scout mission," *Acta Astronautica*, vol. 57, no. 2-8, pp. 121–134, 2005.

[12] M. Nimelman, J. Tripp, A. Allen, D. Hiemstra, and S. McDonald, "Spaceborne scanning lidar system (SSLS) upgrade path," in *Proceedings of SPIE*, vol. 6201, 2006, p. 62011V.

[13] R. Li, K. Di, J. Hwangbo *et al.*, "Integration of orbital and ground images for enhanced topographic mapping in Mars landed missions," in *Proceedings of the Annual NASA Science Technology Conference (NTSC), College Park, MD*, June 2007.

[14] G. Chin, S. Brylow, M. Foote *et al.*, "Lunar Reconnaissance Orbiter overview: The instrument suite and mission," *Space Science Reviews*, vol. 129, no. 4, pp. 391–419, 2007.

[15] D. Smith, H. Frey, J. Carvin *et al.*, "Mars Orbiter Laser Altimeter - Experiment summary after the first year of global mapping of Mars," *Journal of Geophysical Research*, vol. 106, no. E10, pp. 23–689, 2001.

[16] M. Johnston, J. Graf, R. Zurek, and H. Eisen, "The Mars Reconnaissance Orbiter mission," in *IEEE Conference on Aerospace*, 2005.

[17] A. McEwen, W. Delamere, E. Eliason *et al.*, "HiRISE - The High Resolution Imaging Science Experiment for Mars Reconnaissance Orbiter," in *Lunar and Planetary Science XXXIII*, 2002.

[18] B. Horn, "Closed-form solution of absolute orientation using unit quaternions," *Journal of the Optical Society of America*, vol. 4, no. 4, pp. 629–642, 1987.

[19] A. Johnson, "Spin-images: a representation for 3-D surface matching," *Pittsburgh: Carnegie Mellon University*, 1997.

[20] C.-S. Chen, Y.-P. Hung, and J.-B. Cheng, "RANSAC-based DARCES: A new approach to fast automatic registration of partially overlapping range images," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 21, no. 11, pp. 1229–1234, 1999.

[21] F. Cozman, E. Krotkov, and C. Guestrin, "Outdoor visual position estimation for planetary rovers," *Auton. Robots*, vol. 9, no. 2, pp. 135–150, 2000.

[22] R. van den Boomgaard and R. van Balen, "Methods for fast morphological image transforms using bitmapped binary images," *CVGIP: Graphical Models and Image Processing*, vol. 54, no. 3, 1992.

[23] R. Haralick and L. Shapiro, *Computer and Robot Vision. Vol. 1*. Addison-Wesley, 1992.

[24] K. Arun, T. Huang, and S. Blostein, "Least-squares fitting of two 3-D point sets," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 9, no. 5, pp. 698–700, 1987.

[25] M. Fischler and R. Bolles, "Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography," *Communications of the ACM*, vol. 24, no. 6, 1981.

[26] P. J. Carle, "Long-range rover localization by matching lidar scans to orbital elevation maps," Master's thesis, University of Toronto, 2009.

[27] Å. Björck, *Numerical Methods for Least Squares Problems*. Society for Industrial Mathematics, 1996.

[28] P. Lee, "Mars on Earth: The NASA Haughton-Mars Project," *Ad Astra: The Magazine of the National Space Society*, vol. 14, no. 3, 2002.

[29] A. Howard, "Real-time stereo visual odometry for autonomous ground vehicles," in *Intell. Robot. and Syst., 2008. IROS 2008. IEEE/RSJ Int. Conf. on*, Sept. 2008, pp. 3946–3952.

[30] J. Enright, P. Furgale, and T. Barfoot, "Sun sensing for planetary rover navigation," in *Proceedings of the IEEE Aerospace Conference. Big Sky, MT*, March 2009.