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

On-line Calibration of Multiple LIDARs on a Mobile Vehicle Platform

Chao Gao and John R. Spletzer

Abstract— **In this paper, we examine the problem of extrinsic calibration of multiple LIDARs on a mobile vehicle platform. To achieve fully automated and on-line calibration, the original non-linear calibration model is reformulated as a second-order cone program (SOCP). This provides an advantage over more standard linearized approaches in that *a priori* information such as a default LIDAR calibration, calibration tolerances, *etc.*, can be readily modeled. Furthermore, in contrast to general non-linear methods, the SOCP relaxation is convex, returns a global minimum, and can be solved very quickly using modern interior point methods (IPM). This enables the calibration to be estimated on-line for multiple LIDARs simultaneously. Experimental results are provided where the approach is used to successfully calibrate a pair of Sick LMS291-S14 LIDARs mounted on a mobile vehicle platform. These showed the SOCP formulation yielded a more accurate reconstruction and was 1-2 orders of magnitude faster than the traditional non-linear least-squares approach.**

I. INTRODUCTION & RELATED WORK

Over the past decade, Light Detection and Ranging (LIDAR) sensors – herein referred to as LIDARs – have emerged as the dominant exteroceptive sensor for mobile robotics applications. LIDARs are able to provide both relative angle (elevation and/or azimuth) and range measurements to features in the environment with sub-centimeter range accuracy, and with scan rates reaching 75 Hz or higher [1]. Since LIDARs are active sensors that rely upon capturing backscatter from their own light source, they are far more robust to scene illumination changes than traditional vision sensors.

Some robotics applications employing LIDARs include 3D map generation [2], [3], [4], [5], [6], feature extraction, and obstacle detection for robot localization and navigation [7], [8], [9], [10]. LIDARs were instrumental in solving the 2D SLAM problem [2], [11], and enabling Stanley to complete the DARPA Grand Challenge [8]. In fact, all 11 finalists in the subsequent DARPA Urban Challenge relied upon LIDARs as their primary exteroceptive sensing modality [9]. In each of these works, the acquired LIDAR range scans had to be accurately registered to the body frame of the host vehicle. This required an extrinsic calibration to recover the rigid transformation between the vehicle body frame and the LIDAR frame. While this might be physically measured with a fair degree of accuracy, errors of only 1° in LIDAR orientation translate to position errors of ≈ 0.5 meters at ranges of 30 meters. As a result, a more flexible calibration approach is desirable.

Camera calibration has been well studied [12], [13], and a popular toolbox created by the California Institute of Technology is widely used [14]. With the increasing application of LIDARs in three-dimensional mapping, several recent studies have investigated LIDAR calibration with respect to a camera for data fusion purposes [15], [16], [17]. However, fewer efforts have investigated the specific problem of LIDAR calibration. One specific example is [18], where the authors used retro-reflective tape to cover a thin vertical pole fixed to a flat ground area. Data were then collected by driving the vehicle around the pole with different heading angles, and LIDAR calibration was estimated using a sequential quadratic programming method to solve the resulting non-linear optimization problem. In our previous work [6], we also employed retro-reflective targets to support LIDAR calibration. However, these were modeled as point features so that related camera calibration techniques could be leveraged. With a sufficient number of point correspondences over time, the calibration problem was formulated as a non-linear optimization problem. A shortcoming of both of these approaches was the use of non-linear estimators with running times and computational complexities not conducive to on-line operations.

In this paper, we propose a methodology for automated, on-line calibration of multiple LIDARs on a mobile vehicle platform. The main contribution is the relaxation of the original non-linear calibration problem as a second-order cone program (SOCP) [19]. The SOCP formulation provides advantages over more traditional linearized approaches in that *a priori* information such as a default LIDAR calibration, calibration tolerances, *etc.*, can be readily modeled. Furthermore, in contrast to general non-linear methods, the SOCP formulation is convex and can be solved very quickly using modern interior point methods (IPM). This enables the calibration to be estimated on-line for multiple LIDARs simultaneously. Experimental results are provided where the approach is used to successfully calibrate a pair of Sick LMS291-S14 LIDARs mounted on a mobile vehicle platform.

II. PROBLEM STATEMENT

We consider the general LIDAR calibration problem illustrated in Figure 1, where both vehicle body $\mathcal{V}$ and LIDAR sensor $\mathcal{L}$ frames are shown. To further ground the notation, the $X_\mathcal{V}$ axis points forward (heading direction), the $Y_\mathcal{V}$ axis

C. Gao is with the Computer Science and Engineering, Lehigh University, Bethlehem, PA, USA chg205@lehigh.edu

J. R. Spletzer is an Associate Professor of Computer Science and Engineering, Lehigh University, Bethlehem, PA, USA spletzer@cse.lehigh.edu

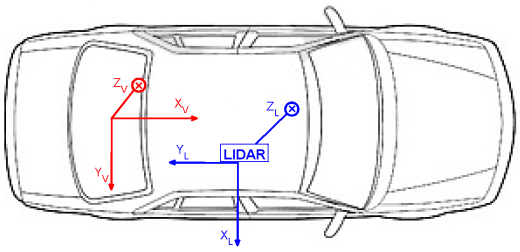

Fig. 1. A general configuration showing vehicle $\mathcal{V}$ and a LIDAR $\mathcal{L}$ coordinate frames. Both $\mathcal{V}$ and $\mathcal{L}$ are right handed.

points right and the $Z_{\mathcal{V}}$ axis points down. The LIDAR sensor frame $\mathcal{L}$ is centered at the origin of the laser beam, the $X_{\mathcal{L}}$ axis points in the right direction (center direction of the laser beam), the $Y_{\mathcal{L}}$ axis points backward and the $Z_{\mathcal{L}}$ axis points down. The extrinsic calibration problem can then be concisely stated as:

Problem 2.1: Develop an automated, on-line approach to recover the translation vector $T_{\mathcal{L}}^{\mathcal{V}} \in R^3$ and the rotation matrix $R_{\mathcal{L}}^{\mathcal{V}} \in SO(3)$ relating the position and orientation of $\mathcal{L}$ in $\mathcal{V}$.

Automation is desired to minimize the need for human involvement, while an on-line approach will support rapid re-calibration at the beginning of vehicle operations.

In solving the calibration problem, we make several assumptions. First, we assume the vehicle platform has the ability to estimate its position $T_{\mathcal{V}}^{\mathcal{W}}$ and orientation $R_{\mathcal{V}}^{\mathcal{W}}$ with respect to a fixed world frame $\mathcal{W}$. Second, we assume that an initial calibration estimate $(\hat{R}_{\mathcal{L}_i}^{\mathcal{V}}, \hat{T}_{\mathcal{L}_i}^{\mathcal{V}})$ for each LIDAR is available. Finally, we assume that the LIDAR systems are capable of simultaneously measuring range r and remission (reflectivity) γ values of features in the environment. While this may appear to be restrictive, there are a wide range of LIDAR systems (*e.g.*, Sick LMS Fast [1] and LD families [20], Hokuyo URG [21] and UTM families [22], Velodyne HDL-64E [23] *etc.*) with such capabilities.

III. TECHNICAL APPROACH

At the highest level, the proposed calibration process can be reduced to four steps:

1) Operate the LIDARs while driving the vehicle through a calibration loop.
2) Automatically segment robust LIDAR features from each calibration loop.
3) Recover inter-loop feature correspondences across successive calibration loops.
4) Estimate the LIDAR calibration using the recovered point correspondences through second-order cone programming.

In this section, we provide details to the process.

A. Recovering Point Correspondences

Given the vehicle pose $(R_{\mathcal{V}}^{\mathcal{W}}, T_{\mathcal{V}}^{\mathcal{W}})$ in $\mathcal{W}$ and the LIDAR extrinsic calibration parameters $(R_{\mathcal{L}}^{\mathcal{V}}, T_{\mathcal{L}}^{\mathcal{V}})$, features segmented from the LIDAR scan as the vehicle traverses the environment can be registered to $\mathcal{W}$ by

$$X_{\mathcal{W}} = R_{\mathcal{V}}^{\mathcal{W}} \cdot (R_{\mathcal{L}}^{\mathcal{V}} \cdot X_{\mathcal{L}} + T_{\mathcal{L}}^{\mathcal{V}}) + T_{\mathcal{V}}^{\mathcal{W}}. \quad (1)$$

From (1) – and assuming that the vehicle pose can be estimated over time – we see that there are 9 unknowns corresponding to the three Euler angles in $R_{\mathcal{L}}^{\mathcal{V}}$, the three translation offsets in $T_{\mathcal{L}}^{\mathcal{V}}$ and the position of the segmented feature $X_{\mathcal{W}} = [x_{\mathcal{W}}, y_{\mathcal{W}}, z_{\mathcal{W}}]^T$ in $\mathcal{W}$. Thus, with a sufficient number of point correspondences, we can solve for the LIDAR extrinsic calibration directly. Such approaches have been used extensively in camera calibration where checkerboard targets allowed point correspondences to be reliably segmented across multiple images using the high contrast of corner features [12], [13]. Inspired by this, we placed several pieces of retro-reflective tape upon vertical poles along a calibration loop to facilitate tracking feature correspondences over time. These targets were reliably segmented from the background using a simple threshold operation $(\gamma > \gamma_{min})$ on the LIDAR remission measurements. This is illustrated at Fig. 2, where a sample landmark pole with two retro-reflective targets is shown on the left. The same pole as viewed by a scanning LIDAR is shown at the right. The locations of two targets as segmented through the threshold operation are clearly visible (highlighted by red crosses).

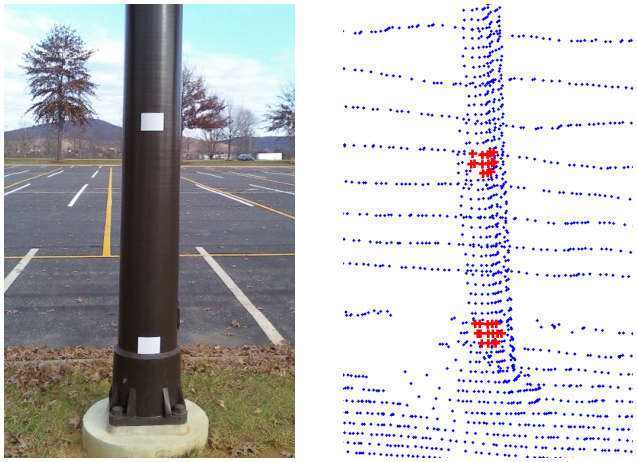

Fig. 2. (Left) A landmark pole used during the LIDAR calibration process. The retro-reflective targets were automatically segmented to track correspondences across multiple calibration loops. These are highlighted by red crosses in the LIDAR scan (right).

Note that the segmented targets were clearly not point features (typical size was 10 cm $\times$ 15 cm). We examined clustering approaches to reduce the target to a point feature, and also treated each valid target point as a single noisy LIDAR measurement so that multiple measurements were associated with each target. Both approaches yielded comparable results. The latter eliminates the clustering step, but dramatically increases the number of correspondences (and thus system equations) artificially. However, since our optimization approach outlined in Section III-B was very efficient computationally, this impact was negligible. The results presented herein employed the latter approach.

We should also note that there is always the potential for noise in the target segmentation process, *i.e.*, that the returned remission value of naturally occurring features exceeds our threshold γ_{min}. The presence of noise can be mitigated to a point by a simple validation gate on target size. We also placed the targets in pairs as shown in Fig. 2 with a known inter-target baseline that could also be used as a validation gate. In practice, these were sufficient to enable reliable target segmentation. An additional benefit of using such pairs is that they provided constraints that also could be used in the calibration process. This is discussed in more details in Section III-B. With the targets reliably segmented for each calibration loop, establishing correspondences across multiple calibration loops was trivial so long as the distance between calibration poles was small compared to error associated with the initial LIDAR calibration.

B. Problem Formulation: An SOCP Approach

Consider a vehicle with N_l LIDARs. Ultimately, we need to register the acquired range scans from each LIDAR frame $\mathcal{L}_i$ to a common world frame $\mathcal{W}$. This registration requires knowledge of extrinsic parameters of both the LIDAR frame relative to the vehicle frame ($R^{\mathcal{V}}_{\mathcal{L}_i}$,$T^{\mathcal{V}}_{\mathcal{L}_i}$) and the vehicle frame versus time with respect to the world frame ($R^{\mathcal{W}}_{\mathcal{V}}(t)$,$T^{\mathcal{W}}_{\mathcal{V}}(t)$). With point correspondences obtained as in Section III-A, the coordinates of these feature points $X_{j,\mathcal{L}_i} = [x_{j,\mathcal{L}_i}, y_{j,\mathcal{L}_i}, z_{j,\mathcal{L}_i}]^T$ in the LIDAR frame can be transformed to $\mathcal{W}$ by

$$X_{j,\mathcal{W}} = R^{\mathcal{W}}_{\mathcal{V}}(t)(R^{\mathcal{V}}_{\mathcal{L}_i} X_{j,\mathcal{L}_i} + T^{\mathcal{V}}_{\mathcal{L}_i}) + T^{\mathcal{W}}_{\mathcal{V}}(t) \tag{2}$$

where $j = 1, 2, \ldots, N_t$ corresponds to the target index. Note from our assumptions in Section II, $R^{\mathcal{W}}_{\mathcal{V}}(t), T^{\mathcal{W}}_{\mathcal{V}}(t)$ are estimated directly by the vehicle pose system, while the LIDAR extrinsic parameters $R^{\mathcal{V}}_{\mathcal{L}_i}, T^{\mathcal{V}}_{\mathcal{L}_i}$ need to be recovered through calibration process.

A straightforward technique to recover the LIDAR calibration would be to employ a non-linear weighted least squares approach of the form

$$\min_{R^{\mathcal{V}}_{\mathcal{L}_i}, T^{\mathcal{V}}_{\mathcal{L}_i}, X^*_{j,\mathcal{W}}} \sum_{i=1}^{N_l} \sum_{j=1}^{N_t} ||X_{j,\mathcal{W}} - X^*_{j,\mathcal{W}}||^2 \tag{3}$$

where, $X^*_{j,\mathcal{W}}$ is the estimated positions of the targets in $\mathcal{W}$. Such an approach was used in our previous paper [6], and solved using a sequential quadratic programming approach [24]. While this provided acceptable calibration results, it could not be implemented on-line as considerable time was required by the non-linear solver. To remedy this, we take a different approach in this work.

As mentioned in Section II, we assume that an initial – albeit potentially poor – calibration estimate ($\hat{R}^{\mathcal{V}}_{\mathcal{L}_\rangle}, \hat{T}^{\mathcal{V}}_{\mathcal{L}_\rangle}$) for each LIDAR is available from physical measurements. This allows us to approximate $R^{\mathcal{V}}_{\mathcal{L}_i}$ as

$$R^{\mathcal{V}}_{\mathcal{L}_i} \approx \hat{R}^{\mathcal{V}}_{\mathcal{L}_i}(I + \Delta R^{\mathcal{V}}_{\mathcal{L}_i}) \tag{4}$$

where $\Delta R^{\mathcal{V}}_{\mathcal{L}_i}$ is given by

$$\Delta R^{\mathcal{V}}_{\mathcal{L}_i} = \begin{bmatrix} 0 & \Delta\Phi_z & -\Delta\Phi_y \\ -\Delta\Phi_z & 0 & \Delta\Phi_x \\ \Delta\Phi_y & -\Delta\Phi_x & 0 \end{bmatrix} \tag{5}$$

and where $\Delta\Phi_x$, $\Delta\Phi_y$ and $\Delta\Phi_z$ are offsets or errors in the physical measurements compared to the actual orientation angles [25]. Applying (4) to (2) we obtain

$$\hat{X}_{j,\mathcal{W}} = R^{\mathcal{W}}_{\mathcal{V}}(t)(\hat{R}^{\mathcal{V}}_{\mathcal{L}_i}(I + \Delta R^{\mathcal{V}}_{\mathcal{L}_i}) X_{j,\mathcal{L}_i} + T^{\mathcal{V}}_{\mathcal{L}_i}) + T^{\mathcal{W}}_{\mathcal{V}}(t)$$

and (3) reduces to a standard linear least squares minimization problem which can be solved directly. In practice, we found shortcomings with such an approach. Because of linearization errors in the rotation matrix and since the problem was unconstrained, we would obtain calibration solutions that while minimizing the reprojection residuals were obviously inconsistent with physical measurements. However, the linear approximation has a significant advantage in that the objective function was now convex. So, to mitigate the impact of linearization errors, we recast the problem as a second-order cone program (SOCP) which can be rewritten equivalently as

$$\begin{aligned} \min_{\Delta R^{\mathcal{V}}_{\mathcal{L}_i}, T^{\mathcal{V}}_{\mathcal{L}_i}, X^*_{j,\mathcal{W}}} \quad & t \\ \text{s.t.} \quad & \textstyle\sum_{i=1}^{N_l} \sum_{j=1}^{N_t} ||\hat{X}_{j,\mathcal{W}} - X^*_{j,\mathcal{W}}||^2 \leq t \end{aligned} \tag{6}$$

Note that as written, (6) and the least-squares formulation are equivalent. However, the strength of the SOCP is that additional constraints can be integrated that exploit *a priori* information regarding system geometry. For example, let us assume that the Euler angles of the initial LIDAR calibration estimate are accurate to within a tolerance of $\pm\delta$. These calibration tolerances can be written as a second-order cone constraint of the form $||\Delta\Phi|| \leq \delta$. Bounds on translation errors could be written similarly as $||T^{\mathcal{V}}_{\mathcal{L}} - \hat{T}^{\mathcal{V}}_{\mathcal{L}}|| \leq \epsilon$. Furthermore, we can also exploit *a priori* knowledge regarding the relative positions of targets. For example, the translation between a pair of targets t_i and t_j on the same calibration pole can be approximated as $T^{t_j}_{t_i} = [0, 0, b]^T$ where b is the measured target baseline. This can also be written as a second-order cone constraint of the form $||X^*_{j,\mathcal{W}} - X^*_{i,\mathcal{W}} - T^{t_j}_{t_i}|| \leq \beta$. Other constraints – such as tolerances relating to the relative pose of LIDAR pairs – can also be imagined. Regardless, the net effect of these constraints is to reduce the feasible set of (6) to a set of physically consistent measurements to limit the effect of linearization on the problem solution. We can now rewrite (6) as

$$\begin{aligned} \min_{\Delta R^{\mathcal{V}}_{\mathcal{L}_i}, T^{\mathcal{V}}_{\mathcal{L}_i}, \hat{X}_{j,\mathcal{W}}} \quad & t \\ \text{s.t.} \quad & \textstyle\sum_{i=1}^{N_l} \sum_{j=1}^{N_t} ||\hat{X}_{j,\mathcal{W}} - X^*_{j,\mathcal{W}}||^2 \leq t \\ & ||\Delta\Phi_{x_i}|| \leq \delta_{x_i},\ i = 1 : N_l \\ & ||\Delta\Phi_{y_i}|| \leq \delta_{y_i},\ i = 1 : N_l \\ & ||\Delta\Phi_{z_i}|| \leq \delta_{z_i},\ i = 1 : N_l \\ & ||T^{\mathcal{V}}_{\mathcal{L}_i} - \hat{T}^{\mathcal{V}}_{\mathcal{L}_i}|| \leq \epsilon_i,\ i = 1 : N_l \\ & ||X^*_{j,\mathcal{W}} - X^*_{i,\mathcal{W}} - T^{t_j}_{t_i}|| \leq \beta_{ij}, \\ & \qquad i = 1 : 2 : N_t,\ j = 2 : 2 : N_t \end{aligned} \tag{7}$$

Note that since each of the constraints we have introduced is a second-order cone constraint, the problem is still a SOCP and can be solved very quickly. However, we must recognize that although we obtain a globally optimal solution to (7), this is only a linearized relaxation of the original problem. As a result, two shortcomings must be addressed. First, because of the approximation in (4), the recovered rotation matrices $R^{\mathcal{V}}_{\mathcal{L}_i} \notin SO(3)$. In other words, they will not be valid rotation matrices. To remedy this, we use Singular Value Decomposition (SVD) to decompose $R^{\mathcal{V}}_{\mathcal{L}_i} = U_i D_i V_i^T$, and then find the closest valid rotation matrix $R^{*\mathcal{V}}_{\mathcal{L}_i} = U_i I V_i^T$ where I is the identity matrix. A second side effect of our convex relaxation is that we cannot guarantee a globally optimal solution. Furthermore, because the recovered rotation matrices $R^{\mathcal{V}}_{\mathcal{L}_i}$ were modified, the solution was no longer *locally* optimal. To address this, we solved (7) a second time using the recovered values $(R^{*\mathcal{V}}_{\mathcal{L}_i}, T^{\mathcal{V}}_{\mathcal{L}_i})$ as our new "initial" calibration estimate. While in theory additional iterations have the potential to refine the solution further, in practice we found that only 2 iterations were typically required to achieve a stable solution.

C. Implementation Specific Model Refinements

For our specific implementation, two additional refinements to the calibration model were necessary. The first relates to our means for estimating the vehicle pose, which relied heavily upon GPS. GPS measurements are subject to both random error as well as deterministic shifts associated with changes in the number and geometry of visible satellites. Since these shifts are strongly correlated in time, we augmented (7) by associating a GPS bias $S_k = [s_{x_k}, s_{y_k}, s_{z_k}]^T,\ k = 1 \dots N_c$ with each calibration loop, where N_c is the number of loops. Note that since this is attempting to compensate for inter-loop shifts, by convention $S_1 \equiv [0,0,0]^T$. Also note that localization solutions not relying upon GPS would not require this refinement.

The second modification relates to our approach for segmenting features. As discussed previously in Section III-A, we treat multiple LIDAR hits on a single target as independent, noisy measurements. One downside with such an approach is that targets that receive a larger number of LIDAR hits (*e.g.*, as a result of being closer to the vehicle) have the potential to bias results. To mitigate this effect, the reprojection residuals are weighted where the weight $W_{i,j,k}$ is inversely proportional to the number of hits by the i^{th} LIDAR on the j^{th} target during the k^{th} calibration loop. With these two refinements, the first constraint in (7) becomes

$$\sum_{i=1}^{N_l}\sum_{j=1}^{N_t}\sum_{k=1}^{N_c} W_{i,j,k}||\hat{X}_{i,j,k,\mathcal{W}} - X^*_{j,\mathcal{W}} - S_k||^2 \leq t \qquad (8)$$

Note that should the targets be reduced to a single point in a pre-processing stage, such a weighting would also be unnecessary.

Fig. 3. Experimental test bed for LIDAR calibration experiments. Ben integrates an OXTS RT-3050 for 6DoF vehicle pose estimation, and a pair of roof mounted vertically scanning Sick LMS291-S14 LIDARs (circled red) to capture relative range measurements.

IV. EXPERIMENTAL RESULTS

A. Experimental Setup

The mobile-platform used in our experiments was "Little Ben," which previously had served as the Ben Franklin Racing Team's entry in the DARPA Urban Challenge [26]. Vehicle pose was provided by an Oxford Technical Solutions (OxTS) RT-3050, which uses a Kalman filter based algorithm to fuse inertial measurements, GPS updates with differential corrections, and odometry information from the host vehicle. It provides 6-DoF pose updates at 100 Hz with a stated accuracy of 0.5 meters circular error probable (CEP). A pair of vertically scanning Sick LMS291-S14 were mounted on the roof of "Little Ben". The two LIDARs are highlighted (circled red) in Fig. 3. The LMS291s offer centimeter-level range accuracy with a 90° field of view, and a sufficiently fast scan rate (75 Hz) to facilitate data capture at reasonable driving speeds.

All experiments were conducted in the parking lots adjacent to Lehigh's Stabler Arena. During data acquisition, Ben was driven multiple times around an L-shaped calibration loop shown at Fig. 4 at speeds of $\approx$ 10-15 km/hr. This loop contained 4 calibration poles, for a total of 8 targets per loop. The targets were 15 cm $\times$ 10 cm in size, and with an intra-pole vertical spacing of 1.0-1.5 meters as shown previously at Figure 2. CVX [27] was used to solve the SOCP, which integrated constraints on translation and orientation tolerances for the initial LIDAR calibration estimates, as well as the relative translation between the two LIDARs, and target geometry constraints as discussed in Section III-B. In all of our experiments, 2 SOCP iterations were run to mitigate the linearizaton error.

B. Calibration Results

Summary quantitative calibration results are shown at Table I. In this, we contrast the performance of our SOCP approach with a non-linear least squares (NLLS) implementation [24] using mean absolute error (MAE) of the target reprojection residuals and run time as performance metrics. In each case, the optimization processes were run on an Intel Core2 Quad CPU Q9300 @ 2.5 GHz. Note that the time

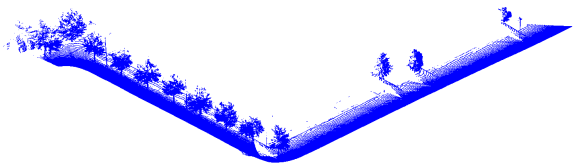

Fig. 4. Reconstruction of the calibration loop used during testing. A total of 8 targets were observed during each loop traversal.

reported for the SOCP is that needed for both iterations. From these results, we make several observations. First,

TABLE I
SUMMARY CALIBRATION RESULTS FOR SOCP AND NLLS ESTIMATORS

Loops	Points	SOCP MAE (cm)	NLLS MAE (cm)	SOCP time (s)	NLLS time (s)
1	131	13.48	15.25	0.69	14.92
2	257	20.05	21.34	0.76	32.11
3	392	21.35	23.17	0.82	40.28
4	515	20.22	21.83	0.94	64.43
5	636	21.14	22.52	1.07	66.83
6	762	20.79	22.36	1.23	80.97

SOCP outperforms the NLLS estimator in terms of the MAE metric for all 6 of the data sets examined. Moreover, the SOCP if also far more efficient computationally. While the difference in MAE performance was $\approx 8\%$, run times for the SOCP formulation were nearly 2 orders of magnitude faster than the NLLS implementation. We should also emphasize that these results are presented for only 2 LIDARs, and we would expect these times to scale similarly to the trend shown in Table I as the number of LIDARs increased. We also note that the MAE in both cases is roughly the same for 2-6 calibration loops. As such, we would expect only incremental benefits from increasing the loop number.

We were admittedly surprised initially that the magnitude of the reprojection errors was not smaller. However, upon further review we note that a portion of this error can be attributed to using each LIDAR target hit as an independent "noisy" measurement for the position of the target center. Even with perfect estimates of the vehicle pose and the LIDARs' extrinsic calibration, the MAE would still be nearly 5 cm due solely to target geometry. We expect the majority of the remaining error was related to uncertainty in vehicle pose. The OxTS RT-3050 has a stated accuracy of 50 cm CEP, meaning that 50% of the x-y position errors would fall within a 50 cm radius. Since the remaining residual error was small in comparison to the vehicle pose error, we were satisfied with the results.

In an attempt to characterize the impact of linearization on solution stability, we repeated the above calibration procedures while varying the initial LIDAR calibration estimate. Specifically, errors of up to $\pm$ 5° were simultaneously added to each of the Euler angles of $\hat{R}^{\mathcal{V}}_{\mathcal{L}_i}$ to gain insights into how sensitive our recovered calibration solution was to the initial calibration estimate. Results from these experiments are summarized at Figure 5. These show that despite the introduction of errors up to 5° , the recovered Euler angle estimates deviated by less than 0.028° from the original calibration estimate.

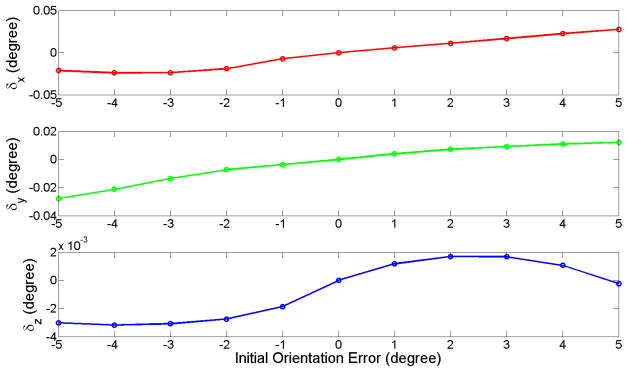

Fig. 5. Difference in Euler angles vs. initial calibration accuracy. Despite the introduction of errors of up to $\pm$ 5° , the recovered Euler angle estimates deviated by less than 0.028° .

We also evaluated the calibration performance qualitatively. Fig. 6 shows a representative scene reconstruction obtained by reprojecting front and back LIDAR range measurements using the initial calibration estimate (top), and the extrinsic parameters recovered using the proposed calibration approach (bottom). In this example, correspondences were collected from 6 calibration loops, and the reprojection was generated from another independent loop. Improvements in data fusion between the two LIDARs from the calibration process are clearly visible.

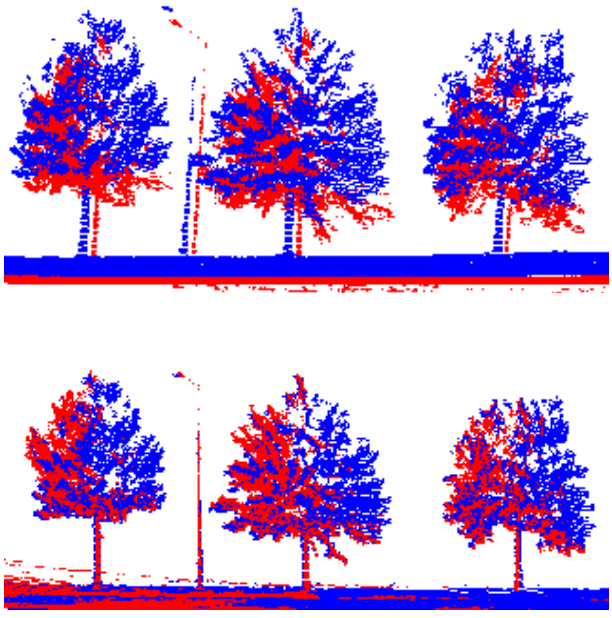

Fig. 6. Reconstruction results before (top) and after LIDAR calibration. Improvements in reprojection and data fusion between the front (red) and back (blue) LIDARS are clearly visible.

V. DISCUSSION

In this paper, we proposed an automated, on-line approach to calibrating multiple LIDARs mounted on a mobile vehicle platform. By using retro-reflective targets to facilitate feature segmentation, the entire calibration process can be fully automated. By reformulating the original non-linear least squares problem as a SOCP, calibration results for multiple LIDARs with hundreds of point correspondences can be obtained in approximately 1 second using a contemporary desktop PC. This convex relaxation is significantly more computationally efficient than more standard NLLS approaches, and supports "real-time" calibration operations. The flexibility of an automated, on-line approach allows LIDARs to be removed and replaced as necessary without fear of an arduous recalibration process. We successfully validated the approach by calibrating a pair of Sick LMS291-S14s LIDARs mounted on a mobile vehicle platform.

Despite these successes, there is still room improvement. To this point, we only evaluated the consistency of calibration results when the initial LIDAR orientation could be estimated to a tolerance of $\pm 5^\circ$. A more significant sensitivity analysis is required to better characterize error limits. We also note that in our problem formulation, we were able to model inter-loop GPS shifts as a result of changes in satellite coverage/geometry. However, intra-loop changes could not be captured. This limitation could be mitigated by employing a RANSAC based approach [28] to filter "outlier loops" which contain significant intra-loop GPS shifts. Note that a RANSAC approach would only be feasible since hypotheses could be evaluated very quickly as a result of the SOCP formulation. This is an additional benefit of the proposed approach.

VI. ACKNOWLEDGMENTS

This work was supported by the National Science Foundation (NSF) Robust Intelligence Program under Grant No. 0844585. Any opinions, findings, conclusions, or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of NSF.

REFERENCES

[1] Sick AG, *LMS200/LMS211/LMS220/LMS221/LMS291 Laser Measurement Systems Technical Description*, Jun 2003.

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

[3] C.-C. Wang, C. Thorpe, and S. Thrun, "Online simultaneous localization and mapping with detection and tracking of moving objects: Theory and results from a ground vehicle in crowded urban areas," in *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 2003.

[4] P. Newman, D. Cole, and K. Ho, "Outdoor slam using visual appearance and laser ranging," in *IEEE International Conference Robotics Automation*, Orlando, Florida, USA, May 2006.

[5] D. Borrmann, J. Elseberg, K. Lingemann, A. Nüchter, and J. Hertzberg, "Globally consistent 3d mapping with scan matching," *Journal of Robotics and Autonomous Systems (JRAS)*, vol. 56, no. 2, pp. 130–142, February 2008.

[6] C. Gao, M. Sands, and J. Spletzer, "Towards autonomous wheelchair systems in urban environments," in *Proceedings of the 2009 International Conference on Field and Service Robotics (FSR '09)*, Cambridge, Massachusetts, July 2009.

[7] S. Thrun, D. Fox, W. Burgard, and F. Dellaert, "Robust monte carlo localization for mobile robots," *Artificial Intelligence*, vol. 128, no. 1-2, pp. 99–141, 2000.

[8] S. Thrun, M. Montemerlo, H. Dahlkamp, D. Stavens, A. Aron, J. Diebel, P. Fong, J. Gale, M. Halpenny, G. Hoffmann, K. Lau, C. Oakley, M. Palatucci, V. Pratt, P. Stang, S. Strohband, C. Dupont, L.-E. Jendrossek, C. Koelen, C. Markey, C. Rummel, J. van Niekerk, E. Jensen, P. Alessandrini, G. Bradski, B. Davies, S. Ettinger, A. Kaehler, A. Nefian, and P. Mahoney, "Winning the darpa grand challenge," *Journal of Field Robotics*, 2006.

[9] M. Buehler, K. Iagnemma, and S. Singh, "Special issues on the 2007 darpa urban challenge," *The Journal of Field Robotics*, vol. 25, no. 8-10, Aug-Oct 2008.

[10] W. Burgard, M. Moors, D. Fox, R. Simmons, and S. Thrun, "Collaborative multi-robot exploration," in *Proceeding of the IEEE Int. Conference on Robotics and Automation*, San Francisco, CA, April 2000, pp. 476–481.

[11] M. Dissanayake, P. Newman, S. Clark, H. Durrant-Whyte, and M. Csorba, "A solution to the simultaneous localization and map building (slam) problem," *IEEE Trans. on Robotics and Automation*, vol. 17, no. 3, pp. 229–241, June 2001.

[12] J. Heikkilä and O. Silvén, "A four-step camera calibration procedure with implicit image correction," in *Proceedings of the 1997 Conference on Computer Vision and Pattern Recognition (CVPR '97)*, San Juan, Puerto Rico, 1997, pp. 1106–1112.

[13] Z. Zhang, "Flexible camera calibration by viewing a plane from unknown orientations," in *International Conference on Computer Vision (ICCV99),*, 1999, p. 666673.

[14] J. Bouguet, "Camera calibration toolbox for matlab," http://www.vision.caltech.edu/bouguetj/calib_doc.

[15] Q. Zhang and R. Pless, "Extrinsic calibration of a camera and laser range finder (improves camera calibration)," in *IEEE International Conference on Intelligent Robots and Systems*, Orlando, Florida, USA, May 2004, pp. 532–537.

[16] C. Mei and P. Rives, "Calibration between a central catadioptric camera and a laser range finder for robotic applications," in *Proceedings of the 2006 IEEE International Conference on Robotics and Automation, ICRA 2006*, Orlando, Florida, USA, May 2006, pp. 532–537.

[17] D. Scaramuzza, A. Harati, and R. Siegwart, "Extrinsic self calibration of a camera and a 3d laser range finder from natural scenes," in *Proceedings of the IEEE International Conference on Intelligent Robots and Systems*, San Diego, CA, USA, October 2007.

[18] J. Underwood, A. Hill, and S. Scheding, "Calibration of range sensor pose on mobile platforms," in *Proceedings of the 2007 IEEE/RSJ International Conference on Intelligent Robots and Systems*, San Diego, CA, USA, October 2007.

[19] M. Lobo, L. Vandenberghe, S. Boyd, and H. Lebret, "Applications of second-order cone programming," *Linear Algebra and Applications, Special Issue on Linear Algebra in Control, Signals and Image Processing*, 1998.

[20] Sick AG, *LD-OEM Laser Measurement System*, March 2006.

[21] Hokuyo Automatic Co. Ltd, *Range-finder Type Laser Scanner URG-04LX Specifications*, July 2005.

[22] Hokuyo Automatic Co. Ltd, *Scanning Laser Range Finder UTM-30LX/LN Specification*, April 2008.

[23] Velodyne, *HDL-64E User's Manual*.

[24] The MathWorks, "Matlab optimization function: fmincon," http://www.mathworks.com.

[25] Oleg S. Salychev, *Applied Inertial Navigation: Problems and Solutions*, BMSTU Press, 2004.

[26] J. Bohren, J. Derenick, T. Foote, J. Keller, A. Kushleyev, D. Lee, B. Satterfield, J. Spletzer, A. Stewart, and P. Vernaza, "Little Ben: The Ben Franklin Racing Team's Entry in the 2007 DARPA Urban Challenge," *Journal of Field Robotics*, 2008, accepted for publication.

[27] M. Lobo, L. Vandenberghe, and S. Boyd, "Socp: Software for second-order cone programming," http://www.stanford.edu/∼boyd/old_software/SOCP.html.

[28] M. Fischler and R. Bolles, "Random sample consensus: A paradigm for model fitting with applications to image analysis and automated cartography," in *Communications of the ACM*, 1981.