Simultaneous Local and Global State Estimation for Robotic Navigation

David C. Moore, Albert S. Huang, Matthew Walter, Edwin Olson*,
Luke Fletcher, John Leonard and Seth Teller
MIT Computer Science and Artificial Intelligence Laboratory
Cambridge, MA
Email: {dcm, ashuang, mwalter, lukesf, jleonard, teller}@mit.edu
*University of Michigan
Ann Arbor, MI
Email: ebolson@umich.edu

Abstract— **Recent applications of robotics often demand two types of spatial awareness: 1) A fine-grained description of the robot's immediate surroundings for obstacle avoidance and planning, and 2) Knowledge of the robot's position in a large-scale global coordinate frame such as that provided by GPS. Although managing information at both of these scales is often essential to the robot's purpose, each scale has different requirements in terms of state representation and handling of uncertainty. In such a scenario, it can be tempting to pick either a body-centric coordinate frame or a globally fixed coordinate frame for all state representation. Although both choices have advantages, we show that neither is ideal for a system that must handle both global and local data.**

This paper describes an alternative design: a third coordinate frame that stays fixed to the local environment over short time-scales, but can vary with respect to the global frame. Careful management of uncertainty in this local coordinate frame makes it well-suited for simultaneously representing both locally and globally derived data, greatly simplifying system design and improving robustness. We describe the implementation of this coordinate frame and its properties when measuring uncertainty, and show the results of applying this approach to our 2007 DARPA Urban Challenge vehicle.

I. INTRODUCTION

Mobile robots generally must solve two problems:

1) What is my location?
2) What do I do now, given my location?

Although these tasks are often addressed separately in the literature, real systems must employ mutually compatible solutions for both tasks. This constraint often forces the designer to make fundamental compromises, such as in the choice of coordinate frame used by the localization algorithm (Problem 1) and the planning algorithm (Problem 2).

A. Localization and the Global Frame

A localization algorithm, for the purposes of this paper, provides a coordinate representation of the robot's position in some globally consistent coordinate frame. This position usually consists of a point (x, y, z) and may include additional information such as velocity, orientation, or angular velocity. The localization algorithm may be a commodity solution such as the Global Positioning System (GPS) which provides coordinates directly in an Earth-centered frame, or an approach such as Simultaneous Localization and Mapping (SLAM) that provide coordinates within a dynamically-constructed map of the environment. In both cases, we consider the position estimates to be from a *global coordinate frame* because each attempts to measure the robot's position relative to some truth which is known to be unique, even though the position estimate itself is noisy.

B. Planning and the Body Frame

When a mobile robot plans its next action, one of its most important considerations is its immediate surroundings, so that it can ensure that the action is carried out safely and accurately. For those robots that have sensing mounted on-board, the most straightforward way to represent this sensor data is to use a *body coordinate frame* where the current robot position is defined exactly as the origin of the frame and sensor data is represented relative to the robot position, possibly with a noise model. A powerful advantage of the body frame is that by definition, *uncertainty about the sensor data is conditional with respect to the robot's position*. In other words, relative motions of the robot from its current position can be predicted with minimal uncertainty against the sensor data.

C. Combining Global and Local Information

This paper considers systems that must combine both global and local information simultaneously. For example, a ground robot may have GPS coordinates of known reference points (global state) and sensor data that tells it where obstacles are located (local state). For efficient operation of the planning algorithms, both types of information need to be combined into a single coordinate frame.

The first possibility is to use the *global frame* for all data. Since the GPS reference points are already represented in the global frame, their coordinates are unmodified and their uncertainty, if any, is not further transformed. The robot's position and its uncertainty are estimated directly by the localization algorithm. The sensor data, which started with some uncertainty, is transformed into the global frame where
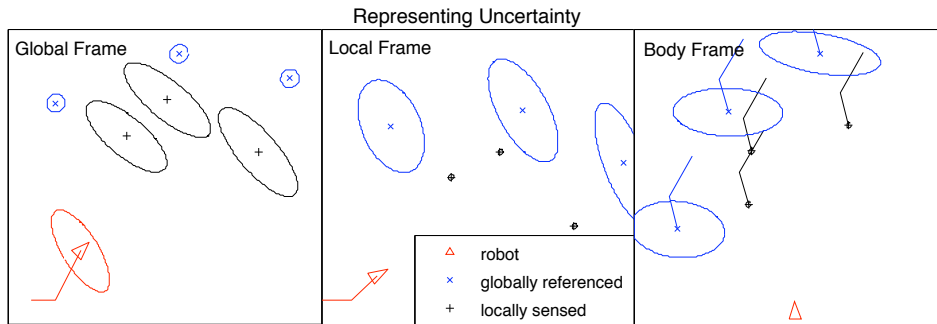
Fig. 1. Graphical representation of the uncertainties (using an ellipsoidal confidence interval) in a typical robot scenario that involves representing global and local state in a single coordinate system. On the left, the global frame is used such that globally referenced points are in their native coordinate system and other quantities, the robot position and its sensor data, are transformed into the global frame. On the right, globally referenced points and local features are transformed into the body frame. In the center, our "local frame" representation is used. Straight lines represent the motion of the features over time in their respective coordinate systems.

it becomes *even more uncertain* since it is conditional with respect to the noisy robot position. These properties of the uncertainty are depicted in the left panel of Figure 1, using ellipses to approximate the 95% confidence interval of the uncertainty.

Besides this additional uncertainty, another reason a global frame is poorly suited for representation of all state is that global localization is often subject to *multi-modal* uncertainty distributions. For example, with GPS, the maximum likelihood estimate of position may change as much as several meters as different satellites go in and out of view or obstructions in the environment create multipath interference. A real example of such a discontinuity is shown in Figure 2. Localization methods such as SLAM also exhibit similar behavior over time as the maximum likelihood position estimates may temporarily lock on to incorrect local features.

There is a real-life example of state estimation in the global frame going awry. In the 2005 DARPA Grand Challenge, Alice, the vehicle from Team Caltech, stored all obstacle information in the global frame, registered against GPS. During the final event Alice experienced GPS interference from overhead power lines. As the GPS interference corrected itself, a large position discontinuity took place in the global frame that caused previously measured sensor data to become invalid with respect to the new position estimate. Alice continued to respect this previous, incorrect data, and eventually drove into a concrete barrier [1]. For the 2007 race, Caltech instead adopted an approach much like the one described herein to avoid similar problems.

For these reasons, we believe the global frame is unsuitable for simultaneous state estimation of local and global data.

A second possible approach for state estimation is to use the *body frame* exclusively. The robot position is represented exactly in this frame since it defines the origin. The sensor data is represented relative to this robot position with some noise, only from the sensor itself or from its time evolution. Globally represented data, such as GPS reference points, are transformed into the body frame using the localization algorithm, and take on the uncertainty of the robot position
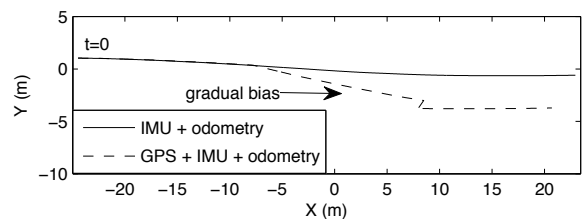
(a) Gradual GPS bias
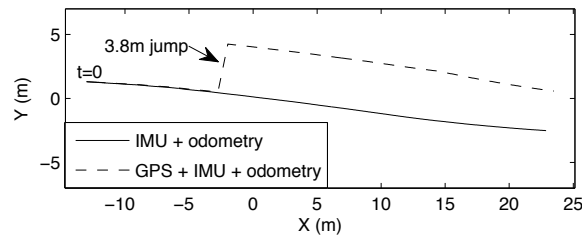
(b) Sudden, persistent GPS bias

Fig. 2. (a) A position estimate affected by a GPS unit experiencing a gradually increasing bias error followed by a small jump. (b) A position estimate affected by a GPS unit experiencing a strong and persistent bias error. Data is taken with an Applanix POS-LV 220 during a drive through Cambridge, MA. Despite a very high-end navigation system, the robot's relative motion estimates are corrupted by GPS error, which would result in significant errors for maps constructed in a global reference frame.

as measured in the global frame. Compared to the global frame, the body frame has the advantage that both the robot position and sensor data are represented with minimal noise, making it ideal for accurate motion planning.

The primary disadvantage of the body frame representation is that all data must be recomputed whenever the robot moves, since sensor data and global data will change position relative to the robot as seen in the right panel of Figure 1. This update can pose a challenge in a working system in terms of computational efficiency, code complexity, and synchronization between different logical modules.

D. The Local Frame

To solve these problems, we propose an alternative coordinate frame for global and local state representation—the "local frame". The local frame is a smoothly varying Euclidean

coordinate system, with arbitrary origin, into which sensor information is projected. The local frame has the desirable property that the vehicle always moves smoothly through this coordinate system; in other words, spatial relationships in the local frame are very accurate over short time scales but may drift over longer time scales.

This smoothness property makes the local frame ideal for registering the sensor data for the vehicle's immediate environment. An estimate of the coordinate transformation between the local frame and the GPS reference frame is updated continuously. This transformation is needed only when projecting a GPS feature, such as a waypoint, into the local frame. All other navigation and perceptual reasoning is performed directly in the local frame.

In the local frame, the robot position is represented without uncertainty, but unlike the body frame, is not necessarily fixed at the origin. Instead, the robot moves in the local frame according the best available relative motion estimate. This estimate could come from an Inertial Measurement Unit (IMU) or Visual Odometry [2], both of which provide smooth and accurate relative motion estimates over short time scales. Although sensor data should still be updated with each movement of the robot, the maximum likelihood position of an older measurement will not change from its previous position in the local frame. Thus, in many cases, no update is necessary. A graphical representation of these uncertainties is shown in the center panel of Figure 1.

II. RELATED WORK

Other groups have independently arrived at the same concept of a smoothly varying local frame, and briefly mention its robustness against GPS discontinuities [3], [4], [5]. At present, we are not aware of a detailed analysis of the local frame in the published literature, such as a probabilistic analysis or its amenability to system simplification.

Coordinate frames are a central concept in sensor fusion algorithms such as the widely used Extended Kalman Filter (EKF) [6] or more advanced approaches such as FastSLAM [7] and pose graph optimization [8]. However, these algorithms solve the more complicated problem of representing the robot trajectory and sensor observations for all time in a single, fully-consistent coordinate system, which we consider a *global frame*. In our concept for the local frame, we seek to represent only enough state to address the robot's immediate path planning and obstacle avoidance needs.

We also distinguish our local frame from some approaches to SLAM that make use of multiple "local maps". A popular technique for efficient large-scale mapping is to incrementally construct multiple sub-maps of a region while estimating geometric relationships across sub-maps [9], [10]. These sub-maps are often called local reference frames, but are based on a significantly different concept from what we propose in this paper.

Castellanos et al. proposed a map-building technique in which feature poses are always represented in the robot's body frame, and showed reduced linearization errors when compared to map-building in a global frame [11]. This idea shares some of our motivations, such as eliminating uncertainty about the robot's pose, but operates exclusively in the body frame. The primary benefit of the local frame proposed here is to provide improved robustness in performing *short-term* processing of sensor data, to support operations such as real-time path planning in the presence of static and moving obstacles.

III. APPROACH

Using the local frame in a robotic vehicle requires two data sources:

- A localization algorithm that computes vehicle position in the global frame, preferably including a measurement of uncertainty.
- A relative motion estimator such as one obtained from an Inertial Measurement Unit (IMU), wheel odometry, or visual odometry.

The localization algorithm can be used to constantly update the transformation between the local frame and the global frame. This transformation is needed whenever globally registered data needs to be placed in the local frame for use by planning algorithms. Over short timescales, this transformation will stay relatively static because the vehicle's position moves through the local frame much the same way as it moves through the global frame. Because the transformation changes slowly, the localization algorithm can be run relatively infrequently, say at 1Hz, unlike traditional inertial navigation systems which run at 100Hz or more.

The relative motion estimator is used to update the vehicle's position in the local frame. There are several ways to get this relative motion estimate. An IMU could be used, integrating the velocities and accelerations it measures. Clever filtering of this data can produce a very smooth estimate of relative motion. Alternatively, if one already has a fully integrated inertial navigation and GPS system, the output velocities from such a system can be integrated to produce the relative motion estimate, provided they are not overly influenced from GPS data. This approach was used by the authors for the MIT Urban Challenge vehicle. Lastly, pure wheel odometry or visual odometry could be used for a relative motion estimate.

A. Tracked State

Formally, we denote the local frame at time t as L_t. The vehicle's pose with respect to its local frame is denoted as $\mathbf{x}_t^{L_t}$. The essential principle of the local frame is that $\mathbf{x}_t^{L_t}$ is a *constant* with no noise. The vehicle pose with respect to the global frame is $\hat{\mathbf{x}}_t^G$, as computed by the localization algorithm. This quantity *is noisy* and should be tracked with some model of its uncertainty. Then $\mathbf{x}_t^{L_t}$ can be used together with $\hat{\mathbf{x}}_t^G$ to define a transformation $T_{L_t}^G$ that transforms a point $\mathbf{p}^G$, some coordinate in the global frame, into a coordinate in the local frame.

We also have to track local sensor measurements over time in the local frame. This data, which we denote as $\mathbf{M}^{L_t}$, may constitute an obstacle map for motion planning, task-specific

goals, or any other quantity that is measured relative to the robot's position. This data is noisy, and should be tracked with some model of uncertainty.

In contrast to traditional sensor fusion techniques, we do not use current sensor observations to refine the position estimates of previously sensed features. For the DARPA Urban Challenge we did not find this extra refinement necessary. However, it would be possible to devise a more advanced approach that used an EKF or other Bayesian model to refine the position estimates for local features over short timescales, while still using the concept of a local frame.

B. Time Evolution

The relative motion estimator, described earlier, measures the change in robot pose over time. For some time t we will denote this change as $\mathbf{u}_t$. We also define a function F that, given the robot's current pose and $\mathbf{u}_t$, will compute the pose at time $t+1$:

$$\mathbf{x}_{t+1}^{L_t} = F\big(\mathbf{x}_t^{L_t}, \mathbf{u}_t\big) + \mathbf{w}_t^{L_t} \tag{1}$$

where $\mathbf{w}_t^{L_t}$ represents a noise term with covariance Q_t. Since the current vehicle pose must always be a noiseless constant in the local frame, we define a new local frame L_{t+1} (instead of using L_t) that does not include the noise term from the time evolution:

$$\mathbf{x}_{t+1}^{L_{t+1}} = F\big(\mathbf{x}_t^{L_t}, \mathbf{u}_t\big). \tag{2}$$

Note that this equation is not an approximation of Equation 1. Rather, it is the *definition* of L_{t+1}.

Since we have migrated the robot pose from L_t to L_{t+1}, we must also do the same for the map $\mathbf{M}^{L_t}$. To do so, we construct a rigid-body transformation $T_{L_{t+1}}^{L_t}$ from the difference in pose between $\mathbf{x}_{t+1}^{L_t}$ and $\mathbf{x}_{t+1}^{L_{t+1}}$ (which is noisy). This transformation updates a measurement in L_t to be in L_{t+1}. Conveniently, the maximum likelihood values of $\mathbf{x}_{t+1}^{L_t}$ and $\mathbf{x}_{t+1}^{L_{t+1}}$ are identical, since the only difference between them is the noise term. Thus, $T_{L_{t+1}}^{L_t}$ has a maximum likelihood value which is the identity function and the maximum likelihood values of $\mathbf{M}^{L_t}$ are updated to be $\mathbf{M}^{L_{t+1}}$ unchanged. However, this transformation still introduces noise. If covariances are being used to model the noise of $\mathbf{M}^{L_t}$, these covariances must be updated.

C. Tracking Uncertainty

One way to update covariances is by linearizing $T_{L_{t+1}}^{L_t}$ about its mean and the map mean, similar to the method described by Smith *et al.* [6]. We decompose the transformation $T_{L_{t+1}}^{L_t}$ into a rotation about the vehicle followed by a translation, and then describe the updated position of a single map feature $\mathbf{m}_i$ with covariance $\Sigma_{\mathbf{m}_i}^{L_t}$ as:

$$\begin{aligned}\mathbf{m}_i^{L_{t+1}} &= T_{L_{t+1}}^{L_t}(\mathbf{m}_i^{L_t}) \\ &= \mathrm{R}_{L_{t+1}}^{L_t}(\mathbf{m}_i^{L_t} - \mathbf{x}^{L_t}) + \mathbf{t}_{L_{t+1}}^{L_t}\end{aligned} \tag{3}$$

We update the covariance of the map feature according to:

$$\Sigma_{\mathbf{m}_i}^{L_{t+1}} = \mathrm{J}_1 \Sigma_{\mathbf{m}_i}^{L_t} \mathrm{J}_1^\top + \mathrm{J}_2 Q_t \mathrm{J}_2^\top \tag{4}$$

where J_1 and J_2 denote the Jacobians corresponding to the transformation operation. They are simplified by the assumption that the mean of $T_{L_{t+1}}^{L_t}$ is the identity transformation. In the case where map features are 2D points $(\hat{x}_{m_i}, \hat{y}_{m_i})$ and the robot pose is described as (x_t, y_t, θ_t), the individual Jacobians evaluate to $\mathrm{J}_1 = \mathrm{I}_{2\times2}$ and

$$\mathrm{J}_2 = \begin{bmatrix} 1 & 0 & -\big(\hat{y}_{m_i} - y_t\big) \\ 0 & 1 & \big(\hat{x}_{m_i} - x_t\big) \end{bmatrix} \tag{5}$$

The cross covariances of map features can be updated in a similar fashion, although in many cases we require only an understanding of the individual covariances.

Approximating the covariance growth in this fashion has the disadvantage of increasing linearization error for each time step. Another approach is to build a single non-linear transformation that updates each sensor measurement from its captured form directly into the latest local frame, reducing linearization error. These covariance estimates are often not needed at every time step, and can be computed only when a particular quantity is needed.

In summary, this time evolution exhibits a valuable property of the local frame: after each prediction step, *the robot's own pose has no uncertainty and the maximum likelihood poses of map features do not change.* This distinguishes it from a pure body frame approach where the poses of map features change after each time evolution step.

D. Globally Referenced Features

The mapping of a globally referenced feature to the local frame is straightforward if we assume Gaussian probabilities. Consider an uncertain 2D global waypoint $\hat{\mathbf{p}}_{\mathrm{wp}}^G$ with covariance Σ_{wp}^G, and represent the uncertainty of the robot's global pose $\hat{\mathbf{x}}^G$ as the covariance matrix $\Sigma_{\mathbf{x}_t}^G$. The mapping is typically a rigid body transformation of the form,

$$\hat{\mathbf{p}}_{\mathrm{wp}}^{L_t} = \mathrm{R}_{L_t}^{G_t}\big(\hat{\mathbf{p}}_{\mathrm{wp}}^G - \hat{\mathbf{x}}^G\big) + \mathbf{t}_{L_t}^{G_t} \tag{6}$$

where $\mathrm{R}_{L_t}^{G_t}$ is a rotation matrix of θ_L^G radians. The covariance matrix is transformed by linearizing about the robot's mean estimates:

$$\Sigma_{\mathrm{wp}}^{L_t} = \mathrm{J}_x \Sigma_{\mathbf{x}_t}^G \mathrm{J}_x^\top + \mathrm{J}_{\mathrm{wp}} \Sigma_{\mathrm{wp}}^G \mathrm{J}_{\mathrm{wp}}^\top \tag{7}$$

where J_x and J_{wp} are the Jacobians corresponding to the transformation.

$$\begin{aligned}\mathrm{J}_x &= \begin{bmatrix} -\cos(\theta_L^G) & \sin(\theta_L^G) & -\hat{y}_{wp}^{L_t} - t_{y_{L_t}}^{G_t} \\ -\sin(\theta_L^G) & -\cos(\theta_L^G) & \hat{x}_{wp}^{L_t} - t_{x_{L_t}}^{G_t} \end{bmatrix} \\ \mathrm{J}_{\mathrm{wp}} &= \begin{bmatrix} \cos(\theta_L^G) & -\sin(\theta_L^G) \\ \sin(\theta_L^G) & \cos(\theta_L^G) \end{bmatrix}\end{aligned} \tag{8}$$

E. New Map Elements

New elements can be added to the local frame map $\mathbf{M}^{L_t}$. In most cases, one already has a function G^{B_t} that, given a raw sensor measurement $\mathbf{z}$, projects this measurement into the body frame B_t:

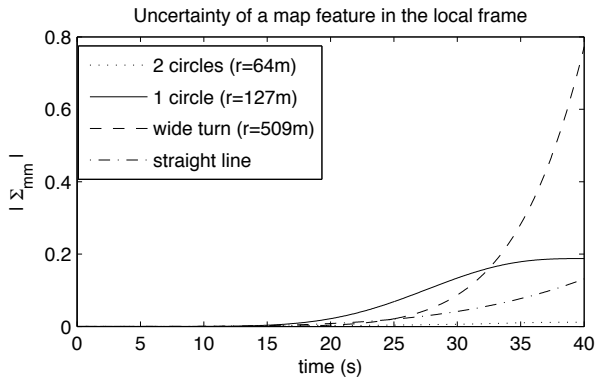$$\mathbf{m}^{B_t} = G^{B_t}(\mathbf{z}) + \mathbf{v} \tag{9}$$

Fig. 3. The change in uncertainty of a map feature is dependent on the robot's path and the uncertainty of its relative motion estimates. Here we show simulated uncertainty estimates for a single map feature observed only once and coincident with the robot's initial pose. Parameters for relative motion uncertainty were chosen to mimic a low-cost MEMS IMU and odometer, with the vehicle traveling at a constant speed of 10m/s. The determinant of the feature's covariance is used as a general measure of its uncertainty.

where $\mathbf{v}$ is a noise term. Since the transformation between body frame and local frame is noiseless, it is easy to construct a new function G^{L_t} directly from G^{B_t} and $\mathbf{x}_t^{L_t}$ that transforms the raw measurement directly into the local frame:

$$\mathbf{m}^{L_t} = G^{L_t}(\mathbf{z}) + \mathbf{v} \qquad (10)$$

As in the time evolution step, covariance can be estimated using a compounding operation that makes use of the Jacobian of G^{L_t}.

IV. DISCUSSION

Figure 3 shows how the uncertainty of a prototypical map feature evolves over time as the vehicle moves. In this case, a single map feature was measured once at $t = 0$ in simulation and then its covariance was tracked over time as the vehicle moves in various simulated paths. As described previously, the maximum likelihood estimate of the feature's position remains constant throughout this simulation—only the covariance changes. This simulation shows that the local frame is able to capture some of the nuance of feature uncertainty even though the feature's position is unchanging.

Of special note is the fact that over short time periods, the accumulated map uncertainty is insignificant relative to the scale at which the vehicle is operating. With sufficiently accurate motion and sensor models, this allows the vehicle to construct essentially noise-free maps suitable for local motion planning over short time horizons.

The local frame was put to the test in the 2007 DARPA Urban Challenge when it was used for all state estimation purposes on Team MIT's vehicle [12]. One critical module of our team's architecture was the drivability map, which served as the primary interface between perception and path planning. By performing all processing for creation of the drivability map in the local frame, the robustness of the overall system was significantly improved.

In our system, an Applanix POS-LV 220 INS system was used for GPS and inertial measurement. Translational and rotational velocities output by the Applanix were integrated to produce the relative position estimates needed for local frame maintenance. A single process was responsible for maintaining and broadcasting the vehicle's pose in the local frame as well as the most recent local-to-GPS transformation. These messages were transmitted at 100Hz.

Figure 2 shows some sample data from the MIT vehicle during a drive in Cambridge in which GPS interference caused errors in the localization algorithm. Nevertheless, the relative position estimates were accurate over short timescales. Figures 4 and 5 show the result of these trajectories when used to build obstacle maps in the local and global frames. Using the global frame causes previously measured obstacles to move with GPS discontinuities. Then, when the obstacles are detected again they appear at a different place in the map, causing smearing. In contrast, the local frame does not exhibit these problems because it is based only on relative position estimates, which remain smooth.

V. CONCLUSION

With the local frame, we have shown that local and global data can be fused seamlessly in a single map with minimal noise. The local frame does not suffer from the increase in uncertainty that results when the global frame is used for state estimation. Furthermore, multi-modal noise characteristics common to global state estimators such as GPS do not cause problems for maps generated in the local frame.

We have demonstrated the local frame on a real system during the 2007 DARPA Urban Challenge without suffering any navigation problems due to GPS discontinuities. We have also applied the local frame to smaller indoor robots as well as an autonomous forklift.

We have also discussed how the local frame can be cheaper to maintain than an equivalent map in the body frame because maximum likelihood estimates for the location of sensed features do not change as the vehicle moves, and probabilistically justified assumptions about map uncertainty allow for simple map updates.

Discussions with other teams that participated in the Urban Challenge revealed that many have come to the same conclusions we have about the disadvantages of the global and body frames. Most of the top-performing teams in the Urban Challenge used a state representation very similar to the local frame, although we believe this paper is the first to describe such an approach in any detail.

By documenting the local frame approach in this paper, we hope that others will consider using it when developing ground robots that require fusion of global and local state. We also hope that manufacturers of Inertial Navigation Systems will consider providing data from their systems, in particular, the relative position estimate, that enable easy implementation of the local frame. It would be especially helpful if these systems could provide a relative position estimate even in the complete absence of a GPS signal.

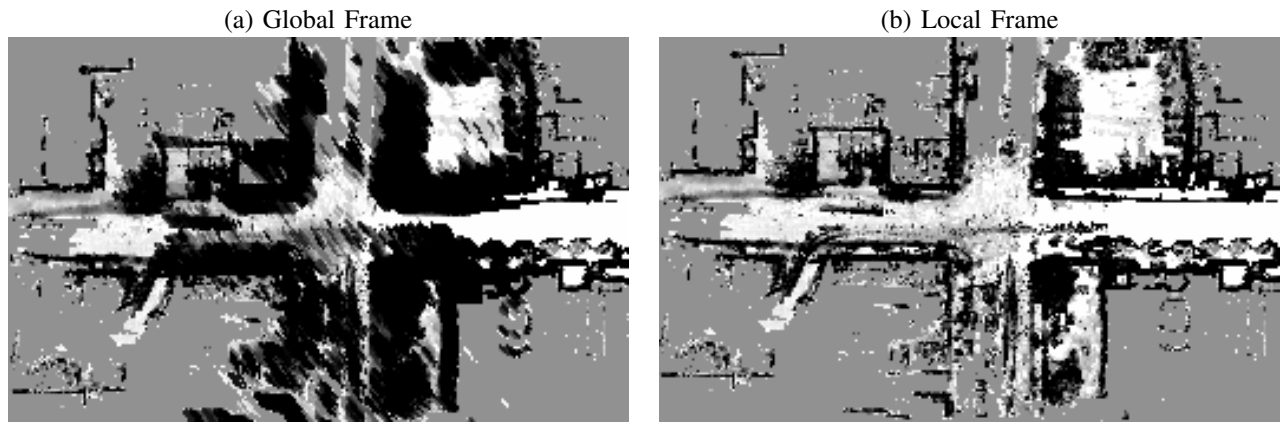(a) Global Frame   (b) Local Frame

Fig. 4. An obstacle map built from vehicle sensor measurements (laser scans) during the vehicle trajectory of Figure 2a. Darker areas represent high cost. (a) The map built directly in the global frame. Obstacles appear "smeared", highlighting the disadvantage of using the global frame due to GPS drift. (b) The map built using the local frame. Despite the GPS drift, the obstacle map is uncorrupted.
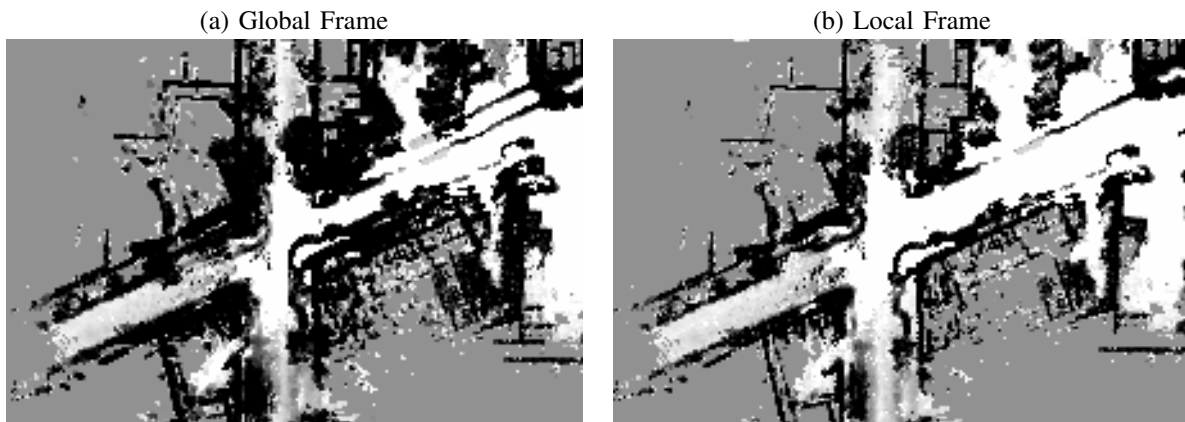
(a) Global Frame   (b) Local Frame

Fig. 5. An obstacle map built from vehicle sensor measurements (laser scans) during the vehicle trajectory of Figure 2b. Darker areas represent high cost. (a) The map built directly in the global frame. Obstacles appear doubly, highlighting the disadvantage of using the global frame due to GPS discontinuities. (b) The map built using the local frame. Despite the GPS discontinuity, the obstacle map is uncorrupted.

REFERENCES

[1] L. Cremean *et al.*, "Alice: An information-rich autonomous vehicle for high-speed desert navigation," *Journal of Field Robotics*, vol. 23, no. 9, pp. 777–810, September 2006.

[2] D. Nister, O. Naroditsky, and J. Bergen, "Visual odometry," in *Int. Conf. Computer Vision and Pattern Recognition*, June 2004.

[3] M. Montemerlo *et al.*, "Junior: The Stanford entry in the Urban Challenge," *Journal of Field Robotics*, vol. 25, no. 9, Aug 2008.

[4] A. Bacha *et al.*, "Odin: Team VictorTango's entry in the DARPA Urban Challenge," *Journal of Field Robotics*, vol. 25, no. 8, Aug 2008.

[5] I. Miller *et al.*, "Team Cornell's Skynet: Robust perception and planning in an urban environment," *Journal of Field Robotics*, vol. 25, no. 8, Aug 2008.

[6] R. Smith, M. Self, and P. Cheeseman, "Estimating uncertain spatial relationships in robotics," in *Autonomous Robot Vehicles*, I. Cox and G. Wilfong, Eds. Springer-Verlag, 1990, pp. 167–193.

[7] M. Montemerlo, S. Thrun, D. Koller, and B. Wegbreit, "FastSLAM: A factored solution to the simultaneous localization and mapping problem," in *Proceedings of the AAAI National Conference on Artificial Intelligence*. Edmonton, Canada: AAAI, 2002.

[8] E. Olson, J. Leonard, and S. Teller, "Fast iterative optimization of pose graphs with poor initial estimates," in *Proc. IEEE Int. Conf. Robotics and Automation*, 2006, pp. 2262–2269.

[9] S. Simhon and G. Dudek, "A global topological map formed by local metric maps," in *Proceedings of the IEEE/RSJ Conference on Intelligent Robots and Systems (IROS)*, vol. 3, Victoria, Canada, October 1998, pp. 1708–1714.

[10] M. Bosse, P. Newman, J. Leonard, and S. Teller, "Simultaneous localization and map building in large-scale cyclic environments using the Atlas framework," *Int. J. Robotics Research*, vol. 23, no. 12, pp. 1113–1139, December 2004.

[11] J. A. Castellanos, J. Neira, and J. Tardós, "Limits to the consistency of EKF-based SLAM," in *Proceedings of the Fifth IFAC Symposium on Intelligent Autonomous Vehicles*, Lisbon, Portugal, 2004.

[12] J. Leonard *et al.*, "A perception-driven autonomous vehicle," *Journal of Field Robotics*, vol. 25, no. 10, Oct 2008.