Autonomous Off-Road Driving in the DARPA Grand Challenge

Eagle Jones, Brian Fulkerson, and Emilio Frazzoli
Henry Samueli School of Engineering and Applied Science
University of California, Los Angeles
Los Angeles, California, USA
{eagle@cs.,bfulkers@cs.,frazzoli@}ucla.edu

Deepak Kumar, Robert Walters,
Jim Radford, and Richard Mason
The Golem Group
Santa Monica, California, USA
{deepak,rwalters,jim,mason}@golemgroup.com

***Abstract*— The Golem Group/UCLA team was a finalist in the 2005 DARPA Grand Challenge, and traveled 22 miles on race day. The Golem Group was also one of the most successful teams in the 2004 Grand Challenge, traveling 5.2 miles on a shoestring budget. We present the strategies, challenges, outcomes, and lessons learned from two years of autonomous vehicle development. Autonomous navigation in the off road environment is a challenging problem, which requires the successful integration of many different sensors and systems. We discuss the integration of GPS, IMU, and odometry together with various models for vehicle dynamics. In particular, we find it useful to incorporate the non-holonomic property of the ground vehicle, but also point out situations in which this constraint is violated and how to compensate. We also explore our approach to selection of quantities to use for control, taking into account which pieces of the system are most sensitive to errors in position and pose estimates. Finally, we address the performance of our system, and future possibilities for improvement.**

I. Introduction

In 2004, the U.S. Defense Advanced Research Projects Agency (DARPA) held its first Grand Challenge. A million-dollar prize was offered for the individual or team that could build an autonomous ground vehicle capable of traversing a 150 mile off-road course in the desert in under 10 hours. The event was described as "a field test to accelerate research and development in autonomous ground vehicles." No vehicle was able to travel more than seven miles during the first race, and a second event was held (for twice the prize money) in 2005. We qualified for and competed in both races, completing five miles of the 2004 event, and 22 miles of the 2005 race course.

A. Prior Work

From a scientific standpoint, the general localization problem in the DARPA Grand Challenge (DGC) could arguably be considered "solved". GPS navigation is common, and filtering frameworks which combine multiple sensors for state estimation are well-understood[1], [2], [3], [4]. Unmanned ground vehicles such as those built by Dickmanns[5] have been capable of high-speed autonomous driving in the structured highway environment for years. However, it is clear that there is an unusually large divide between theory and practice in off-road autonomous vehicle navigation. Prior to the DGC, the U.S. Army's eXperimental Unmanned Vehicle[6] represented the state of the art for autonomous operation in unstructured off-road environments. It could navigate an obstacle-filled environment at just over 6 kph.

B. Vehicle

Our vehicle was a standard heavy duty pickup truck. We installed actuators to control the accelerator, brakes, and steering. Sensors used to determine the vehicle state included: a high-quality NovAtel GPS receiver with decimeter-level DGPS corrections from the Omnistar satellite, a BEI C-MIGITS III integrated INS/GPS system, a quadrature encoder which utilizes Hall sensors on the rear differential to measure velocity, a high-precision optical encoder to measure steering angle, and feedback (including velocity) from the vehicle's On-Board Diagnostics system. We forego a complete description of the system architecture here and emphasize state estimation; for a more general discussion, see [7].

II. Sensor Fusion

In theory, it is relatively simple to estimate a vehicle's state (position, orientation, and velocity) given measurements from the variety of sensors available to us. The standard approach, and one which we adopted, is to use a Kalman filter. It would make sense to combine all sensor data in a single filter, and one would expect excellent performance from such a filter (better than that of any individual sensor). However, the assumptions of such an approach are that the system follows a particular model, and that the noise in the system is reasonably well-behaved. Unfortunately, this is far from true in the real world. Our best models fail to adequately describe the complexities of the vehicle's interactions with its environment. Sensors exhibit reasonable noise characteristics most of the time, but occasionally depart substantially from reason. Multipath error in GPS is the worst offender, sometimes creating consistent, non-zero mean biases which are difficult to differentiate from valid data. Additionally, the system must operate flawlessly for many hours in a harsh environment where sensors may fail completely.

Given these realities, a simple monolithic filter actually increases the system's sensitivity to noise and component failure. While it may show optimal performance most of the time, a single departure from normal operations can lead to complete failure. We address this problem by simultaneously

employing a multitude of state estimation techniques. We developed two different Kalman filters which are independently capable of estimating the vehicle's entire state. Each filter relies on a different model and different sensors. Also, individual sensors and subgroups of sensors are incorporated into simple estimators for individual components of the state. For example, velocity can be obtained from the IMU, from the GPS, from the quadrature encoder, or from the vehicle's network. All estimators run in parallel, and each estimate is associated with a measure of health and a priority. Depending on the health estimates and priorities, the final state reported to the rest of the system can be provided entirely by a single filter, or each component (position, velocity, etc.) can come from a different source. In fact, in typical operations, we achieved the best results by combining the output of three different state estimation techniques in the final state. This system is robust to failures and strong outliers from individual sensors as long as at least one of the estimators for each critical quantity remains viable.

In the following sections, we describe our two main Kalman filters. They represent substantially different approaches. Both filters rely on GPS in the long term, but attempt to maintain good estimates through periods of long GPS outage, poor signal, and multipath (conditions which were all present the the DGC). The first state estimator uses a model analogous to a bicycle for the vehicle and relies heavily on the state history. The second estimator is based on a six-degree-of-freedom rigid-body model, with the addition of a "soft" non-holonomic constraint on the vehicle's velocity enforced as a virtual heading measurement.

III. The Bicycle Estimator

The bicycle estimator is a discrete time extended Kalman filter, with the following inputs:

1) GPS-provided latitude and longitude at 20 Hz.
2) Rear axle velocity at 30 Hz.
3) Steering wheel angle at 20 Hz.

The state propagation equations for the bicycle estimator are:

$$\begin{aligned} \tilde{x}_{k+1} &= \hat{x}_k + \hat{v}_k \Delta t \frac{\cos(\hat{\phi}_k + \hat{\sigma}_k)}{\cos \hat{\sigma}_k} \\ \tilde{y}_{k+1} &= \hat{y}_k + \hat{v}_k \Delta t \frac{\sin(\hat{\phi}_k + \hat{\sigma}_k)}{\cos \hat{\sigma}_k} \\ \tilde{\phi}_{k+1} &= \hat{\phi}_k + \frac{\hat{v}_k \Delta t}{d} \tan \hat{\sigma}_k \end{aligned} \quad (1)$$

where, as illustrated in Fig. 1, (x, y) are the local cartesian coordinates of the center of the front axle of the car. The GPS sensor is located at this point. The angle ϕ is the heading of the car with respect to the local x axis. The angle σ is the steering angle of the car as shown in the figure. v is the rear axle velocity. The quantity d is the distance between the front and rear axles of the vehicle. A caret over a variable implies that the variable is an estimate, and a tilde over a variable implies that the variable is predicted or propagated forward. The time elapsed since the last state update is Δt.

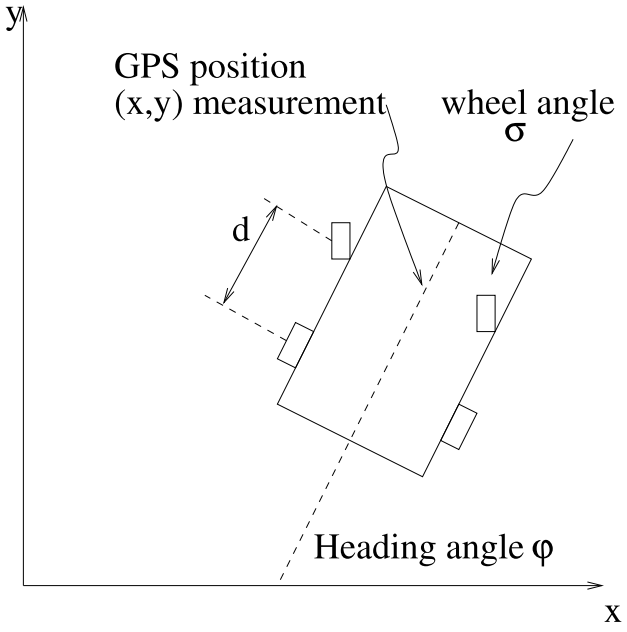

Fig. 1. The "bicycle" model.

The subscripts in the above equations refer to successive time indices in state propagation.

The bicycle model works well when the steering angle is small. For large steering angles however, the model is inaccurate and exhibits considerable lag between the state estimate and the GPS measurement. Also, the model is clearly violated at high velocities, when the vehicle slips as it turns. Thus, the steering angle, σ, is calculated as:

$$\hat{\sigma}_k = \hat{\gamma}_k(\sigma_{\text{measured}} - \hat{\sigma}_{\text{bias } k}) \quad (2)$$

where γ is a steering calibration factor which compensates for model inaccuracies. The steering bias estimate σ_{bias} is used to compensate for bias in the measured steering angle. The rear axle velocity, v, is scaled:

$$\hat{v}_k = \hat{S}_k(v_{\text{measured}}) \quad (3)$$

where v_{measured} is the velocity input to the bicycle estimator. The slip factor S compensates for slip (while accelerating or driving on a slope) as well as variations in the vehicle's apparent wheel radius. This approach generally was effective for tracking slip. However, in one instance when the vehicle was stuck with the wheels spinning, we found that the slip factor was not able to track long periods of constant slipping. We were not greatly concerned by this edge case.

The state variables are latitude and longitude (translated to local Cartesian coordinates x,y), heading ϕ, slip factor S, and either steering bias σ_{bias} or steering calibration factor γ. The steering calibration factor and the steering bias cannot be estimated simultaneously, as the state variables become unobservable. Therefore, the bicycle estimator operates in the following two modes:

1) When the vehicle is going straight ($|\sigma_{\text{measured}}| < 3^\circ$), σ_{bias} is estimated.

2) When the vehicle turns ($|\sigma_{measured}| > 3^\circ$), γ is estimated. This allows us to compensate for model inaccuracies in hard turns.

A. Modeling System Noise

The state is propagated and updated with every measurement. We associate no noise with the state propagation equations (1), and additive white Gaussian noise with the inputs (v_{measured} and σ_{measured}). To efficiently track the constants in the state variables (S, σ_{bias}, γ), we assume an additive white Gaussian noise term in their propagation. We also associate noise with the GPS measurements (latitude and longitude). The variances assigned to the noise processes described above were tuned to ensure uncorrelated innovations, stability, and reasonably quick convergence. Adjusting the variances alters the "trust" associated with the history of the vehicle state in relation to the new measurements. This is a delicate balance in particular for GPS, as different noise modes must be handled explicitly.

B. Determining the Appropriate GPS Measurement Noise Variance

The data received from the GPS consists of latitude, longitude and horizontal dilution of precision (HDOP). HDOP is a figure of merit which is directly related to the number of satellites visible to the GPS antenna, and indirectly to the noise in each measurement. However, our attempts at mapping the HDOP values to actual variances were futile, as we did not observe a monotonic relationship. We noticed that an HDOP of more than 5 usually corresponded to multipath reception, which is associated with large, non-zero mean noise. In the absence of any ad-hoc relationship between HDOP and noise variance, we opted to use this HDOP threshold to switch between a small and large variance for the GPS noise. By introducing this nonlinearity into the GPS noise model, we achieved very satisfactory results.

C. Projection of the Innovations

The bicycle estimator is based on a very intuitive model of the vehicle, which motivates us to consider the GPS innovations in a physically relevant reference frame rather than any arbitrary reference frame. It is beneficial to project the innovations into the local frame of the vehicle: parallel to the direction of motion and perpendicular to it.

For an ideal estimator, the innovations will be uncorrelated with each other. However, we tuned the estimator to just achieve innovations with zero mean. While tuning the estimator it was very useful to consider this physical interpretation of the innovations. For example, a DC bias in the parallel innovation implied that we were not tracking the slip factor (S) adequately. Thus, to ensure a zero mean parallel innovation, the variance associated with the propagation of slip factor was increased.

D. Adaptive Shaping of the Innovations

The noise in the GPS data is highly correlated and we have very little a priori knowledge of the variance. Very often, especially when the vehicle drives near a wall, or approaches a tunnel, we observe highly erratic jumps in the GPS measurements due to multipath reflections. Without any a priori knowledge of the variance in such cases, the state estimate jumps around, with corresponding erratic behavior in control, obstacle detection, and path planning.

To counter these "unphysical" jumps, once the estimator is converged, we clip the innovations to a certain maximum absolute value. For example, a GPS measurement corresponding to a perpendicular innovation of 2 meters in 0.05 sec while driving straight is unphysical. In this case, perpendicular innovations should be clipped to a nominal value (in our case, six inches). When we added this behavior, we found that it did prevent large jumps in the state estimate, but at a great cost. We observed that if the innovations were clipped to a fixed range, then in certain situations the estimate lagged far behind a "good" set of GPS measurements and took a long time to reconverge. To prevent this, we clip the innovations adaptively. The limit is determined as the minimum of either a fixed limit or the mean of the innovations in the last two seconds, scaled by a factor slightly greater than unity. The parallel and perpendicular components of the innovation are clipped separately with different numerical constants.

E. Countering Time Delays

All the software for our vehicle runs on a single laptop computer. While this has some great benefits in terms of ease of development and simplicity, it also leads to some complications. In particular, we experience non-constant delays between the time sensor data appears on the bus and the time it is processed by the control program, due to other tasks like obstacle detection or path planning running. Usually this delay is nominal ($\sim$50-200 μs), but it is sporadically very large. A large delay in GPS data processing manifests itself as a large negative parallel innovation. However, these innovations are clipped effectively by the method just described, and do not significantly affect the state estimate. In the future, we plan to explore various options to enforce harder real-time performance.

F. Special Modes of Operation

Two other model violations require special handling:

- The GPS signal drifts considerably when the vehicle is at rest. This is unphysical, and was particularly bad for our path planning algorithm. Therefore, when the vehicle is going extremely slowly or is at rest, the variance assigned to the GPS measurements is increased significantly. In such a case, the bicycle estimator essentially works like an integrator rather than a filter.
- We observed that the GPS signal occasionally jumps discretely. These jumps usually correspond to the presence of a power transmission line nearby. This was troublesome because the GPS takes a while ($>$2 s) to

reconverge after the jumps. These unphysical jumps are easily detected from corresponding jumps in the parallel and perpendicular innovations. After the detection of such jumps, the GPS variance is increased until it can be considered trustworthy again, i.e., until the innovations are within a certain limit.

G. Relationship with the IMU

One advantage of this model is that it converges very quickly while driving straight, usually in approximately $5\,\mathrm{s}$. Once the bicycle model converges, the heading estimate can be used to initialize the Inertial Measurement Unit (IMU). While driving straight, the heading estimate from the bicycle estimator is extremely good, typically within 0.5° of the short-term IMU heading. However, on sharp turns, the heading estimate is up to $\sim 3^\circ$ from the IMU computed heading. Thus, IMU heading is given a higher priority in the state arbitration system, especially when GPS signals are bad. In the future, we may extend the bicycle model to include the angular rotation and linear displacement data from the IMU.

IV. The Six-Degree-of-Freedom Estimator

We now describe the of the six-degree-of-freedom estimator (6DOF). Like the bicycle estimator, 6DOF is implemented as a discrete-time Extended Kalman Filter. The estimator is designed using fairly standard techniques for strap-down inertial navigation systems. Rather than a detailed model of the vehicle's dynamics, the filter relies mainly on the rigid-body kinematic equations. However, due to the absence of a magnetometer or other means to measure the vehicle's orientation, and the need to be able to ensure convergence of the non-linear filter without requiring an initial calibration procedure, we exploit knowledge of the vehicle's dynamics in the form of virtual heading measurements from inertial velocity data.

The vehicle is modeled as a rigid body moving in three-dimensional space; the state of the vehicle can hence be described by a position vector $p \in \mathbb{R}^3$, representing the location with respect to an Earth-fixed reference frame of the on-board IMU, a velocity vector $v = dp/dt$, and a rotation matrix $R \in SO(3)$, where $SO(3)$ is known as the Special Orthogonal group in the three-dimensional space, and includes all orthogonal 3 by 3 matrices with determinant equal to $+1$. The columns of R can be thought of as expressing the coordinates of an orthogonal triad rigidly attached to the vehicle's body (body axes).

The inputs to the estimator are the acceleration and angular rate measurements from the IMU, provided at 100 Hz, and the GPS data, provided at 20 Hz. In addition, the estimator has access to the velocity measured at the rear differential, and to the steering angle measurement.

In the following, we will indicate the acceleration measurements with $z_\mathrm{a} \in \mathbb{R}^3$, and the angular rate measurements with $z_\mathrm{g} \in \mathbb{R}^3$. Moreover, we will indicate with Z_a the unique skew-symmetric matrix such that $Z_\mathrm{a} v = z_\mathrm{a} \times v$, for all $v \in \mathbb{R}^3$. A similar convention will be used for z_g and other three-dimensional vectors throughout this section.

The IMU accelerometers measure the vehicle's inertial acceleration, in body-fixed coordinates, minus the gravity acceleration; in other words,

$$z_\mathrm{a} = R^T(a - g) + n_\mathrm{a},$$

where a and g are, respectively, the vehicle's acceleration and gravity acceleration in the inertial frame, and n_a is an additive, white Gaussian measurement noise. Since the C-MIGITS IMU estimates accelerometer biases, and outputs corrected measurements, we consider n_a as a zero-mean noise.

The IMU's solid-state rate sensors measure the vehicle's angular velocity, in body axes. In other words,

$$z_\mathrm{g} = \omega + n_\mathrm{g},$$

where ω is the vehicle's angular velocity (in body axes), and n_g is an additive, white Gaussian measurement noise. As in the case of acceleration measurements, n_g is assumed to be unbiased.

The kinematics of the vehicle are described by the equations

$$\begin{aligned} \dot{p} &= v, \\ \dot{v} &= a, \\ \dot{R} &= R\Omega, \end{aligned} \tag{4}$$

in which Ω is the skew-symmetric matrix corresponding to the angular velocity ω, and we ignore the Coriolis terms for simplicity. This is justified in our application due to the relatively low speed and short range, and the fact that errors induced by vibrations and irregularities in the terrain are dominant with respect to the errors induced by ignoring the Coriolis acceleration terms.

We propagate the estimate of the state of the vehicle using the following continuous-time model, in which the hat indicates estimates:

$$\begin{aligned} \dot{\hat{p}} &= \hat{v}, \\ \dot{\hat{v}} &= \hat{R} z_\mathrm{a} + g, \\ \dot{\hat{R}} &= \hat{R} Z_g. \end{aligned} \tag{5}$$

An exact time discretization of the above, under the assumption that the (inertial) acceleration and angular velocity are constant during the sampling time is:

$$\begin{aligned} p^+ &= p + v\Delta t + \tfrac{1}{2}\left(\hat{R} z_\mathrm{a} + g\right)\Delta t^2, \\ v^+ &= v + \left(\hat{R} z_\mathrm{a} + g\right)\Delta t, \\ R^+ &= R \exp(Z_g \Delta t). \end{aligned} \tag{6}$$

The matrix exponential appearing in the attitude propagation equation can be computed using Rodrigues' formula. Given a skew-symmetric 3 by 3 matrix M, write it as the product $M = \Omega\theta$, such that Ω is the skew-symmetric matrix corresponding to a unit vector ω; then

$$\exp(M) = \exp(\Omega\theta) = I + \Omega \sin\theta + \Omega^2(1 - \cos\theta). \tag{7}$$

The error in the state estimate is modeled as a 9-dimensional vector $\delta x = (\delta p, \delta v, \delta \phi)$, where

$$\begin{aligned} p &= \hat{p} + \delta p, \\ v &= \hat{v} + \delta v, \\ R &= \hat{R} \exp(\delta \Phi). \end{aligned} \tag{8}$$

Note that the components of the vector $\delta\phi$ can be understood as the elementary rotation angles about the body-fixed axes that make $\hat{R}$ coincide with R; such a rotation, representing the attitude error, can also be written as $\delta R = \exp(\delta\Phi) = \hat{R}^T R$.

The linearized error dynamics are written as follows:

$$\frac{d}{dt}\delta x = A\delta x + Fn, \tag{9}$$

where

$$A := \begin{bmatrix} 0 & I & 0 \\ 0 & 0 & -RZ_{\mathrm{a}} \\ 0 & 0 & -Z_{\mathrm{g}} \end{bmatrix}, \qquad F := \begin{bmatrix} 0 & 0 \\ R & 0 \\ 0 & I \end{bmatrix}. \tag{10}$$

When no GPS information is available, the estimation error covariance matrix $P := \mathrm{E}[\delta x(\delta x)^T]$ is propagated through numerical integration of the ODE

$$\frac{d}{dt}P = AP + PA^T + FQF^T.$$

Position data from the GPS is used to update the error covariance matrix and the state estimate. The measurement equation is simply $z_{\mathrm{gps}} = p + n_{\mathrm{gps}}$. In order to avoid numerical instability, we use the UD-factorization method described in [8] to update the error covariance matrix and to compute the filter gain K.

Since the vehicle's heading is not observable solely from GPS data, and we would like to keep calibration and initialization procedures to a minimum (e.g., to allow for seamless resets of the filter during the race), we impose a soft constraint on the heading through a virtual measurement of the inertial velocity, with the form:

$$z_{\mathrm{nhc}} = \arctan\left(\frac{v_{\mathrm{East}}}{v_{\mathrm{North}}}\right) - \lambda\sigma,$$

where σ is the measured steering angle and λ is a factor accounting for the fact that the origin of the body reference frame is not on the steering axle.

In other words, we effectively impose a non-holonomic constraint on the motion of the vehicle through a limited-sideslip assumption. This assumption is usually satisfied when the vehicle is traveling straight, but may not be satisfied during turns; moreover, when the vehicle is moving very slowly, the direction of the velocity is difficult to estimate, as the magnitude of the velocity vector is dominated by the estimation error. Hence, the virtual measurement is applied only when the steering angle is less than 10 degrees, and the vehicle's speed is at least 2.5 mph (both these values were determined empirically).

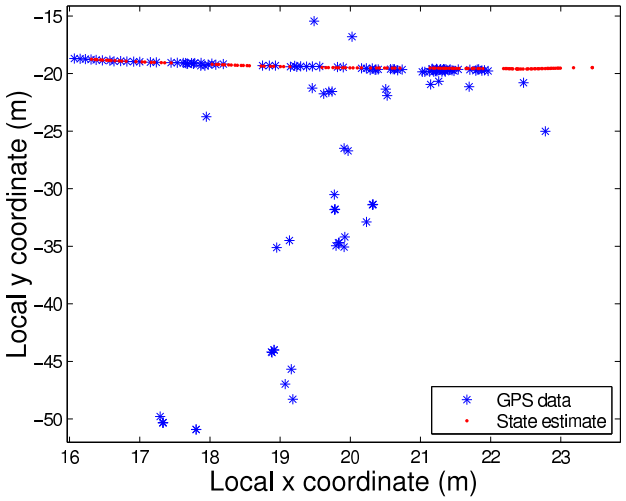

Fig. 2. Performance of state estimator while going under a bridge.

V. Results

Performance of the bicyle estimator while driving under a bridge is shown in Fig. 2. The track of the vehicle is from left to right in this figure. Note that the GPS data has a huge variance due to weak signal and multipath. Also note that the GPS noise is highly correlated; in this case the variance is much larger in the direction perpendicular to the motion as compared to the parallel direction. As seen from the figure, the state estimator is immune to the large "unphysical" jumps in GPS data. Ocassional jumps of length $\sim 0.3m$ are observed in the successive updates of estimated state. These correspond to a delay in received GPS data and large negative parallel innovations as discussed in Section III-E.

During the 2005 Grand Challenge, we successfully traversed 22 miles of challenging terrain in just under one hour. Our race ended due to a memory allocation failure unrelated to the state estimation system, and in general we were very pleased with our localization performance during the race. In analyzing the data collected during the race, we confirmed that the state estimation system was immune to multipath and interference. While our vehicle had "hiccups" due to failures in obstacle detection and path planning, our estimated state was reliable and consistent throughout the race. Typical performance for both estimators was position accuracy within 10cm and heading accuracy within $0.1°$.

VI. Conclusions

It is interesting to compare the two state estimators. During practice, qualification, and the race, each estimator was observed to outperform the other in different situations. With good GPS reception, we generally found the performance of 6DOF to be superior due to the high precision of the integrated IMU information. However, in real-world conditions, the bicycle model was far more robust due to its stronger reliance on a physical model of the vehicle. Thus, the bicyle model was given a higher priority in the state arbitration system. Both estimators did exhibit catastrophic failure modes

where estimates were unreliable for extended periods of time, sometimes requiring a full reset of the filter. We conclude that employing both and implementing our dynamic arbitration system were critical to our success.

ACKNOWLEDGMENTS

In addition to the authors of this paper, the Golem Group included Josh Arensberg, Bill Caldwell, David Caldwell, Kerry Connor, Jeff Elings, Jerry K. Fuller, Izaak Giberson, Maribeth Mason, Jason Meltzer, Brent Morgan, Stefano Soatto and Jim Swenson. Our industry partners included Mobileye (Yaniv Alon and Amnon Shashua) and Toshiba (Hiroshi Hattori and Nobuyuki Takeda). Without the invaluable technical assistance of these individuals, constructing the Golem vehicles would not have been possible.

The financial support of the Henry Samueli School of Engineering and Applied Sciences at the University of California, Los Angeles, is gratefully acknowledged. In addition, we received financial support and in-kind donations from a number of other organizations, including BEI Technologies and NovAtel.

REFERENCES

[1] R. E. Kalman, "A new approach to linear filtering and prediction problems," *Transactions of the ASME*, vol. 82, pp. 35–45, 1960.

[2] R. E. Kalman and R. S. Bucy, "New results in linear filtering and prediction theory," *Transactions of the ASME*, vol. 83, pp. 95–107, 1961.

[3] A. Gelb, *Applied Optimal Estimation*. Cambridge, MA: MIT Press, 1974.

[4] G. Welch and G. Bishop, "An introduction to the kalman filter," University of North Carolina at Chapel Hill, Chapel Hill, NC, Tech. Rep. TR 95-041, 1995.

[5] E. D. Dickmanns, "Dynamic vision-based intelligence," *AI Magazine*, vol. 25, no. 2, pp. 10–30, 2004.

[6] J. A. Bornstein and C. M. Shoemaker, "Army ground robotics research program," in *Unmanned Ground Vehicle Technology V*, 2003, pp. 303–310.

[7] R. Mason, J. Radford, D. Kumar, R. Walters, B. Fulkerson, E. Jones, D. Caldwell, J. Meltzer, Y. Alon, A. Shashua, H. Hattori, N. Takeda, E. Frazzoli, and S. Soatto, "The Golem Group / UCLA autonomous ground vehicle in the DARPA Grand Challenge," *Journal of Field Robotics, Special Issue on the DARPA Grand Challenge 2005*, to appear.

[8] R. M. Rogers, *Applied Mathematics in Integrated Navigation Systems*, 2nd ed., ser. AIAA Education Series. Reston, VA: AIAA, 2003.