Utilization of Position and Orientation Data for Preplanning and Real Time Autonomous Vehicle Navigation

Dr. William "Red" Whittaker, *Carnegie Mellon University, Pittsburgh, Pennsylvania*
Louis Nastro, *Applanix Corporation, Richmond Hill, Ontario, Canada*

BIOGRAPHIES

Dr. William "Red" Whittaker leads the Robot City development effort. He is the Fredkin Research Professor of Robotics at Carnegie Mellon University 's Robotics Institute and the founder of the Field Robotics Center and the Robotics Engineering Consortium. Dr. Whittaker received his BS from Princeton. MS and PhD are from Carnegie Mellon University.

Louis Nastro received his BA (Hons) degree from York University in 1994. He also received MA in Strategic Studies and MS. from the Odegard School of Aerospace Sciences, University of North Dakota. He is Director of Land Products at Applanix Corporation.

ABSTRACT

Pre-planning information about terrain is as important as real time navigation for achieving peak performance in autonomous driving. Both of these rely on accurate, trustworthy estimates of position and orientation. Pre-planning uses position-registered terrain data to detail an intended route and set desired speeds to achieve elapsed times that are otherwise impossible. Navigation and path tracking also rely on position estimation for effective driving. This paper profiles technologies that successfully guided Sandstorm and H1ghlander, two robots from the Carnegie Mellon Red Team, through a 132 mile course. Fusion of data from various onboard sensors provides accurate real time perception. To control driving, position estimation must persist even during prolonged periods of GPS outages. The application of the Applanix POS LV for both preplanning and real time operation of the vehicle is outlined to illustrate how the system provides highly accurate position and orientation which is crucial in maximizing the performance of autonomous vehicles.

This paper first describes the hardware and architecture which comprises the POS LV system and how the data from the POS LV was utilized. This is followed by analysis of test results which highlight the system's robust positioning capabilities in GPS adverse environments and demonstrates how position and orientation information is used for pre-planning.

INTRODUCTION

This paper addresses the problem of how to achieve reliable and repeatable positioning data and how to use that for maximizing the performance of autonomous vehicles. Robust positioning (which is the ability of a positioning system to maintain accurate position information even during GPS outages), is a necessary component of successfully navigating the vehicle. However, accurate orientation of the vehicle to derive very precise measures of vehicle dynamics for both pre-planning functions and real time navigation are absolutely essential to provide onboard sensors with relevant data to steer autonomous vehicles on their intended track, and deal with unanticipated conditions upon routes.

POS LV DESCRIPTION

The POS LV system is a tightly coupled inertial/GPS system which is shown in Figure 1. T*ightly-coupled* implementation, optimally blends the inertial data with raw GPS observables from individual satellites (ranges and range rates). In this case if the number of visible satellites drops below four, the inertial navigator is still aided by the GPS. The result is improved navigational accuracy when compared to the *free-inertial* operation.

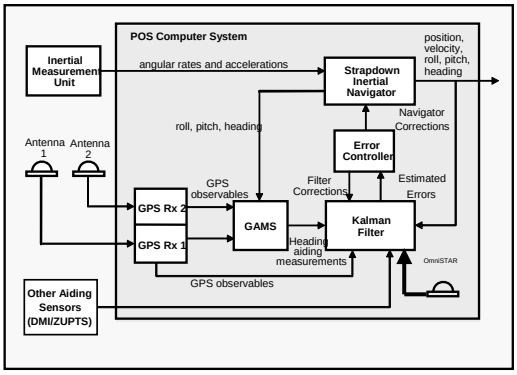

Figure 1: POS LV Tightly Coupled System Architecture

An additional advantage of *tightly-coupled* integration is the improved re-acquisition time to recover full RTK position accuracy after satellite signal loss (see [1]). The inherent benefits of *tightly-coupled* data blending become readily apparent in the accuracy and integrity of the resulting navigation solution. By contrast, *loosely-coupled* implementation blends the inertial navigation data with the position and velocity output from the GPS. If the number of visible satellites is sufficient for the GPS to compute its position and velocity, i.e. four or more satellites, then GPS position and velocity are blended with the inertial data. Otherwise, if the GPS data is not available, the system will operate without any GPS aiding. The inertial navigator computes position, velocity and orientation of the IMU. The Kalman filter estimates the errors in the inertial navigator along with IMU, distance measurements instrument (DMI) and GPS receivers. System components are shown in Figure 2.[2] The only addition to this system setup for the Red Team was a Trimble Ag 252 receiver which provided RTCM corrections for position information. Typical position accuracies for open sky conditions are in the order of 0.5m RMS.

Figure 2: POS LV System Components

The GPS Azimuth Measurement Subsystem (GAMS) integrates the IMU with a 2-antenna heading measurement system. As long as there is GPS coverage GAMS continuously calibrates the IMU and azimuth does not drift. A single-antenna configuration, in comparison, requires dynamic heading alignment and delivers heading measurements that suffer from drift. GAMS uses a carrier phase differential GPS algorithm to measure the relative position vector between the two antennas. GAMS uses carrier phase measurements from five or more satellites to estimate and, eventually, to identify a set of integer phase ambiguities for each satellite being tracked by both receivers. For the ambiguity resolution algorithm to work, both receivers must track at least five common satellites. Once tracking has been obtained, GAMS will continue to operate with as few as four satellites. The GAMS heading system will not provide measurements when fewer than 4 GPS satellites are available. During GPS outages, POS LV will continue to provide accurate heading measurements drifting at the rate of about 1 arc min/min. Accurate heading is critical for robotic vehicle navigation especially when intermittent or non existent GPS conditions occur over extended periods of time.

The distance measurement instrument (DMI) is another essential piece of the POS LV hardware which outputs pulses representing fractional revolutions of the instrumented wheel. These pulses are converted by the POS LV into measurements of incremental distance travelled by the vehicle when no GPS is available.

For both Red Team vehicles, the DMI serves not only to bridge GPS outages and provide POS LV with incremental distance estimation, but as an input into the velocity controller for detection of when the vehicle may be stuck. Wheel slippage is monitored by comparing the DMI output to the velocity reported by the POS LV system. When the system reports speeds over 5m/sec and a velocity of 0 m/sec., the vehicles execute a set of protocols utilizing the perception system and POS LV data to find an alternate path to the next pre-programmed point.

AUTONOMOUS VEHICLES

Autonomously navigating at high speeds for long distances across the desert, as was required the DARPA Grand Challenge, necessitates robust and capable robots. The Red Team developed a pair of racing robots (H1ghlander and Sandstorm) which are more fully described elsewhere [6].

Figure 3: Sandstorm (left) and H1ghlander (right) were developed to navigate at high-speed in desert terrain.

USE OF POSE ESTIMATION FOR DRIVING

Autonomous vehicles sense, plan, and drive without the benefit of onboard human skill. Position estimation is essential for each of these. In a path-centric architecture, the fundamental action is to follow a path. A path data structure is pervasive through our approach. Pre-planned routes are provided to the navigation system and planning operations act as filters on the path. Positioning is also used to steer sensor focus and allow the perception system to handle incompletely sensed terrain.

The Red Team's path-centric architecture provides a simple method for incorporating a pre-planned route. It

reduces the search space for a planning algorithm from the square of the path length to linear in the path length, since planning is performed in a corridor around the pre-planned route. The path-centric approach avoids problems with arc-based arbitration such as discontinuities in steering commands (due to contradictory information) and jerky control (due to discrete arc-sets).

To use terrain evaluation data from multiple sources, the architecture uses a map based data fusion approach. To provide this functionality the architecture defines a second fundamental data type; the map. In this system, a map is a rectilinear grid aligned with the world coordinate system and centered on the robot. Each of the sensor processing algorithms produces its output in the form of a cost map. Cost maps are a specific map type that represent the traversability of a cell with a numeric value.

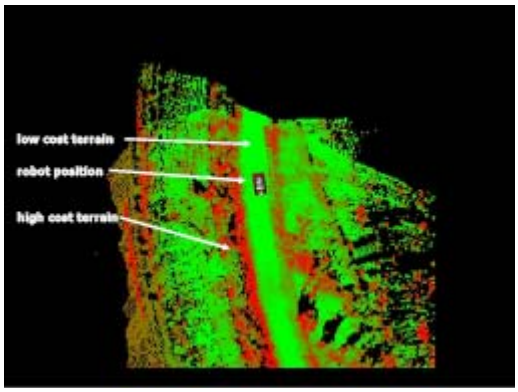

Figure 4: An example cost map showing low and high cost terrain.

Position estimation is critical to map fusion. Map fusion is critical to the robustness of the navigation system, as it enables the system to cope with sensor failures and missing data. To use the data from the various sensor processing algorithms it is necessary to combine it into a composite world model (either implicitly or explicitly). This relies heavily on position estimation. Sensor fusion generates a composite map using a weighted average of each of the input maps, each registered using position estimation. Each of the processing algorithms specifies a confidence for the output map it generates. The fusion algorithm then combines the maps with these weightings to generate the composite expected cost map. This design allows the sensor processing algorithms to adjust their contribution to the composite map if they recognize that they are performing poorly. In practice a set of static weights, based on a heuristic sense of confidence in the algorithm's ability to accurately assess the safety of terrain, worked well. With calibrated sensors, and with good position data, this approach produces usable composite terrain models. Figure 5 shows various input maps and the resulting, fused composite map.

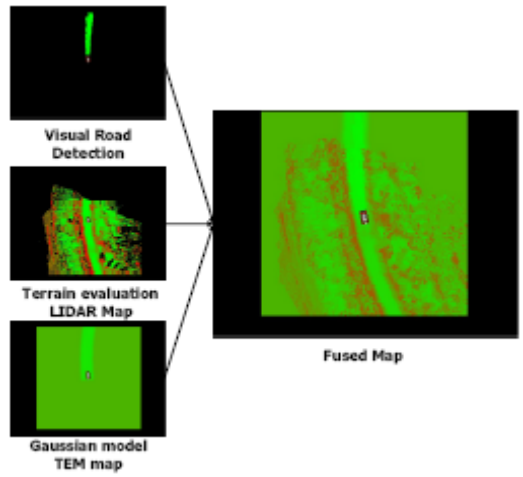

Figure 5: An illustration of fused sensor maps.

PLANNING AND DRIVING

Pure pursuit tracking utilizes POS LV data in conjunction with the vehicle's drive-by-wire system to ensure waypoints are reached within the programmed speed / time parameters. Accurate path tracking is dependent upon vehicle position and dynamics. The pure pursuit path tracking algorithm utilizes the information provided by the POS which ensures that the maximum speed and curve apex trajectory is within the constraints of the vehicle's performance so as to not allow it to skid or veer off course.

Trajectory planning algorithms attempt to find an optimal path from a starting point to a goal point. Many approaches have been developed to solve various planning problems. The most popular are deterministic, heuristic-based algorithms and randomized algorithms. Some work has been done to set speeds and curvatures reactively. In general, the search space for a mobile robot is large, so search is computationally expensive. Deterministic searches typically sample the search space at a resolution that allows fast search, but decreases efficiency of solutions. Randomized algorithms do not sample the search space, but tend to generate somewhat random trajectories.

In the Grand Challenge, a prescribed route consisting of a centerline with a set of bounds was provided. This information can be exploited to significantly improve planning speeds. The bounds and centerline did not exactly define a road, but instead kept vehicles near terrain that DARPA desired the vehicles to traverse.

Prior to the race it was not known how accurately the centerline would track roads and trails.

MAPPING AND PREPLANNING

Good human drivers adjust radii, favor lanes and inherently set speeds while racing. They gracefully enter and exit turns, and "read the terrain" or use foreknowledge of the course to slow down for harsh terrain features. Robots can utilize terrain data for preplanning. Terrain data can be from maps or aerial imagery, but highest fidelity and accuracy are achieved by using POS data to register laser range scans into models that we call “drive-by topography.” These models are obtained by driving a vehicle equipped with laser scanner and POS system over terrain and recording topographic imagery. The method is broadly applicable for detailed surveys that are unachievable from satellite or aerial flyover.

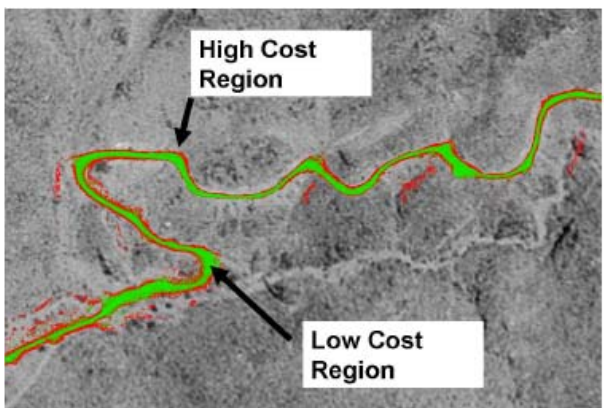

Figure 6: Topography data overlaid on imagery.

Detailed terrain topography can be acquired by collecting range scanner and vehicle position measurements while driving. This data is combined to generate a height map reconstructed by solving for the position of each range measurement in 3-D space. The resulting surface models provide resolution and accuracy that are unobtainable from satellites or from traditional maps.

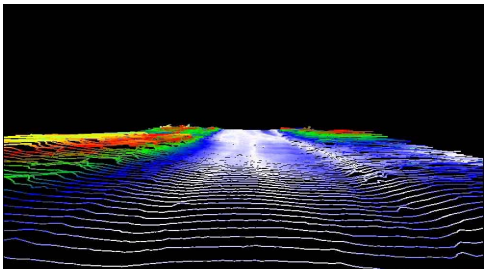

Figure 7: Cost regions with road

An example of the detail of topography is shown in Figure 7. In this figure, the high and low cost regions are identified, where the low cost region exists on the actual location of the road.

Figure 8: The Topographer with sensors and POS LV.

As shown in Figure 8, an H1 Hummer is outfitted with a laser scanner and POS LV system to derive laser scans of the terrain typically with .25m resolution and 1.5m accuracy. The POS LV records position and orientation of the scanner relative to the vehicle frame of reference enabling seamless drive by topography data.

PREPLANNING

Preplanning uses terrain data, including topography data, to detail paths and set speeds. The pre-planning phase requires highly accurate data of vehicle dynamics and terrain conditions. Path detailing is a process that transforms a set of coarse waypoints and speed limits to a preplanned path with one meter spaced waypoints. The resulting path defines location and speed, and is a smooth trajectory for robots to follow with the goal of increasing the probability of successfully navigating a route. As part of this process a smooth path is generated by interpolating between an initial set of coarse waypoints using splines. The splines can then be adjusted by human editors to smooth tight radius curves and to bias the path away from areas of high risk. The splines are then converted to one meter spaced waypoints that are followed by the vehicles. The interpolation process produces a route of curved splines from waypoint to waypoint defined by a series of control points and spline angle vectors. Human editors can alter splines by shifting a series of spline control points, and spline angle vectors that adjust to specify the location and orientation of a path. The generated splines are constrained to ensure continuity to the second order which prevents discontinuities in both the position and heading of a path. The preplanning process is illustrated in Figure 9.

During pre-planning, a speed setting process specifies the target speeds for an autonomous vehicle given a target elapsed time to complete a preplanned path. Speed setting is performed by assessing the risk for a

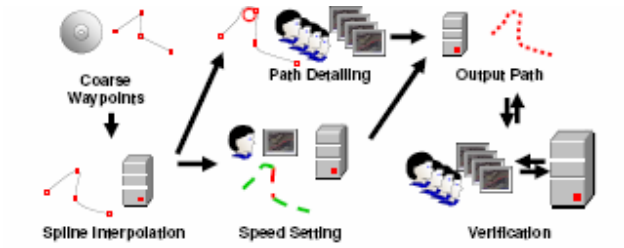

Figure 9: Pre Planning Process

given robot to traverse a section of terrain based on available information. An automated process then uses a speed policy generated by combining the risk assessment with any speed limits imposed on the course to assign speeds to each waypoint in the path.

The result of preplanning is the generation of two high performance, successful routes for two autonomous robots in the 2005 Grand Challenge traverse of 132 miles in about 7 hours. Aerial imagery was used for most of the foreknowledge. Drive-by topography covered only 3% of the actual race route, since the team performed little reconnaissance in the race region. POS data was utilized for both pre-planning and real time vehicle navigation.

POS LV SYSTEM PERFORMANCE

While the DARPA Grand Challenge was run in mostly clear sky conditions, the abilities of the POS LV to provide robust position and orientation data can be best illustrated when used in unfavorable GPS conditions. The following is a brief summary of POS LV system performance utilizing the same OmniSTAR VBS


Figure 10: POS LV Test Route Downtown Toronto

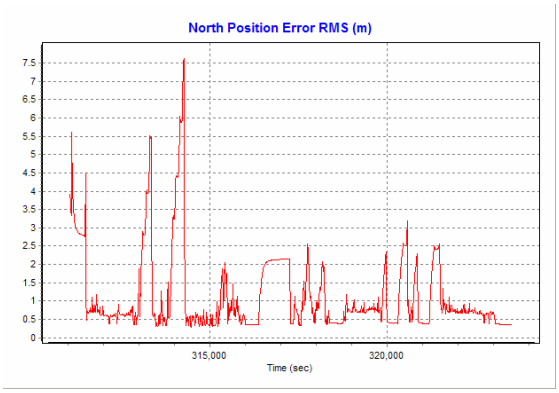

Figure 11: North Position Error RMS

Position data utilized during the race. Figure 10 details the test route which took place in downtown Toronto, and Figure 11 summarizes the real time performance of the system. Accuracies varied from 0.5 to 7.5m RMS real time, with more than 75% of the mission recording accuracies of between 0.5 to 3.2m RMS. GPS was unavailable or severely degraded throughout the entire test as illustrated in Figure 12.

What is interesting to note is the real-time performance of GPS alone as compared to the POS LV. The blue dots represent real-time GPS while the yellow lines show the real-time trajectory as calculated by the POS LV. For the majority of the course, the GPS data

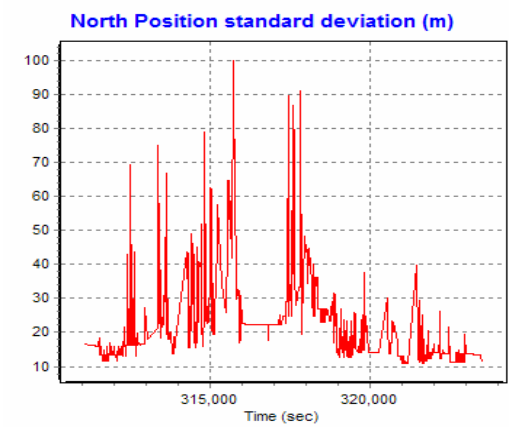

Figure 12: GPS Availability – North Position Standard Deviation

is spurious at best and even when available, position errors of over 50m were recorded.

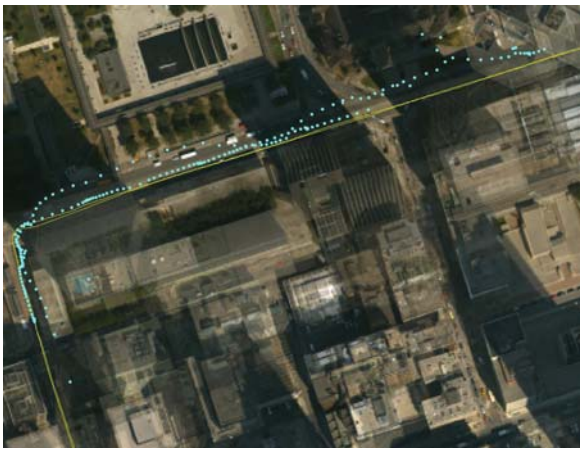

Figure 13: Repeat Trajectories

Figure 13 is magnified to illustrate how the POS unit tracks precisely, enabling repeatable results along a one way street, with differing levels of GPS signal reception. Of note here is how one set GPS data is quite close to the actual. However, on the second run, the GPS is off the actual track by over 15 meters, yet the POS LV trajectory tracked right over 95% of its previous course.

SUMMARY

Robust, accurate, timely position data provided by the POS LV system was utilized in a path-centric architecture to autonomously sense, plan and drive with unprecedented performance. Preplanning was predicated on topography models that were obtained by driving a vehicle equipped with laser scanner and POS system over terrain. During tests Sandstorm and H1ghlander each drove over 1000 autonomous miles. Both robots completed race length runs at race pace on routes more difficult than the 2005 DARPA Grand Challenge race route.

The robustness of this approach was clearly demonstrated by Sandstorm and H1ghlander's performances at the 2005 Grand Challenge. Despite a failing engine and the loss of a perception sensor, both robots completed the course within 20 minutes of the winning time. Without the careful design and implementation of a robust positioning system, high performance autonomous driving would not have been possible.

REFERENCES

[1] Scherzinger, B. *Precise Robust Positioning with Inertial/GPS RTK* Proceedings of ION-GPS-2000, Salt Lake City UH, September 20-23, 2000

[2] Peterson, K, *DARPA Grand Challenge 2005 Technical Paper*, Pittsburgh, PA 24 August 2005.

[3] J. Lalonde, N. Vandapel, and M. Hebert, *Data Structure for Efficient Processing in 3-D*, Robotics: Science and Systems 1, MIT Press, June, 2005.

[4] Kelly, A., et al., *Toward Reliable Off-Road Autonomous Vehicle Operating in Challenging Environments*, International Symposium on Experimental Robotics, June, 2004, Singapore.

[5] A. Talukder, R. Manduchi, R. Castano, K. Owens, L. Matthies, A. Castano, R. Hogg, *Autonomous Terrain Characterisation and Modelling for Dynamic Control of Unmanned Vehicles"* Proceedings of the 2002 IEEE/RSJ Intl. Conference on Intelligent Robots and Systems, EPFL, Lausanne, Switzerland, October 2002.

[6] Urmson, C., et al, *A Robust Approach to High Speed Navigation for Unrehearsed Desert Terrain*, Pittsburgh, PA 2006.

[7] Gutierrez, A., Galatali, T., Gonzalez, J.P., Urmson, C., Whittaker, W., *Preplanning for High Performance Autonomous Traverse of Desert Terrain Exploiting a Priori Knowledge to Optimize Speeds and to Detail Paths*, Pittsburgh, PA December 2005.

[8] Bartz, D., Whittaker, W., *Dynamic Considerations for Achieving High Performance Unmanned Driving by Integrating Terrain Models, Vehicle Physics, and Planning* Pittsburgh, PA March 2006.

[9] Koon, P., Whittaker, W., *Evaluation of Autonomous Ground Vehicle Skills*, Pittsburgh, PA March 2006