

POSITION AND ORIENTATION DATA REQUIREMENTS FOR PRECISE AUTONOMOUS VEHICLE NAVIGATION

Louis Nastro

Director, Land Products Applanix Corporation 85 Leek Crescent Richmond Hill, Ontario CANADA L4B3B3
lnastro@applanix.com

KEYWORDS: Position and Orientation, DARPA Urban Grand Challenge, Autonomous Vehicles, Sensor Fusion, POS LV

ABSTRACT

The challenge of navigating an autonomous vehicle over large distances was illustrated in 2005 at the DARPA Grand Challenge when 4 out of 23 Teams successfully completed a 132 mile course within a 10 hour time limit. What the Grand Challenge revealed is that one of the most critical components of a successful autonomous vehicle was the reliability of accurate pose (positioning and orientation estimation). Data from the Applanix POS LV provided critical vehicle dynamics, navigation and planning data. Pre-planning information is as important as real time navigation for achieving peak performance in autonomous driving as demonstrated by the Carnegie Mellon Red Team and their approach. With the third iteration of the DARPA Grand Challenge, autonomous vehicles were required to navigate an urban course which contained dynamic obstacles a host of other impediments, providing the most realistic operational environment to date for autonomous vehicles. This paper will outline the uses of positioning and orientation data for autonomous vehicle operations at the 2005 and 2007 events and how the Applanix POS LV system was an integral part of the Tartan Racing and Stanford University top finishing results at the DARPA Urban Grand Challenge.

1. INTRODUCTION

This paper addresses the problem of how to achieve reliable and repeatable positioning data and maximizing the performance of autonomous vehicles. Robust positioning (which is the ability of a positioning system to maintain accurate position and orientation 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.

2. POS LV DESCRIPTION

The POS LV system is a tightly coupled inertial/GPS system which is shown in Figure 1. *Tightly-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 *free-inertial* operation. 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

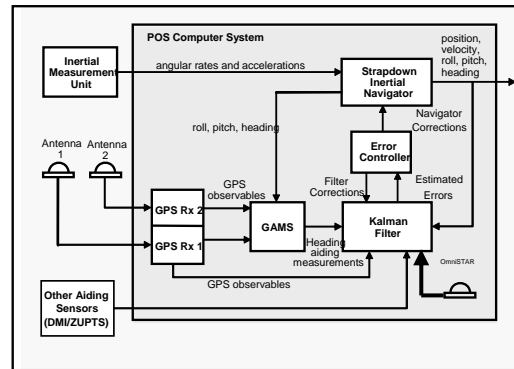


Figure 1: POS LV Tightly Coupled System Architecture

measurements instrument (DMI) and GPS receivers. System components are shown in Figure 2. The only addition to this system setup for the Carnegie Mellon Red Team at the 2005 DARPA Grand Challenge was a Trimble Ag 252 receiver which provided OmniSTAR VBS corrections for position information. Typical position accuracies for open sky conditions are in the order of 0.5m RMS. For the DARPA Urban Grand Challenge Ag 332 units were utilized and Teams had a choice to complete the course with OmniSTAR XP or HP corrections in order to achieve, in open sky conditions, 10 to 20 centimeter accuracy.



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 and the rate of drift is heavily dependent on vehicle dynamics. GAMS uses a carrier phase differential GPS algorithm to measure the relative position vector between the two antennas. These carrier phase measurements from five or more satellites are used 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. In the 2005 DARPA Grand Challenge both Red Team vehicles Highlander and Sandstorm utilized DMI data 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.

3. THE 2005 DARPA GRAND CHALLENGE – A CHALLENGE IN SENSOR FUSION

For the 2005 DARPA Grand Challenge Applanix collaborated with Carnegie Mellon's Red Team for the position and orientation component of their two entries into the race, Highlander and Sandstorm [2]. Both vehicles utilized a series of LIDAR and radar systems to sense terrain and feed that information into onboard computers which would modify pre-planned route information to avoid obstacles and deal with changes in terrain. The data provided by the POS LV is



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

essential in governing vehicle dynamics to safely navigate the course for real time operation. The Red Team's approach involved a methodical analysis of the course terrain and modification of the RDDF (the DARPA defined route definition file) in order to provide both vehicles the optimum path. LIDAR data (provided through a gimbal located on the roof of the vehicle which provided medium and long range terrain data) and supplemental lasers (scanning the boundaries of the

navigable track) in addition to the short range radar (vital for detecting targets in the immediate vicinity of the vehicle), were incorporated to form a view of the world within which the robots would sense and evaluate terrain. Position information from the POS LV is critical in determining the direction of rotation of the gimbal in order to sense the oncoming terrain and georeferencing point cloud data [3].

The Red Team utilized a path-centric architecture which provided a simple method for incorporating a pre-planned route. The primary reason in utilizing this approach was to reduce 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 avoided problems with arc-based arbitration such as discontinuities in steering commands (due to contradictory information) and jerky control (due to discrete arc-sets) [4].

With data derived from the LIDAR and radar systems, it is fused into a composite model of the terrain as illustrated in figure 4. The data is processed and is assigned a value dictating its 'cost'. Lower elevations (shown in green) are assigned a lower cost whereas higher elevation (shown in red) is assigned a higher cost. The autonomous vehicle is 'trained' to navigate on the low cost sections of terrain and make modifications to its pre-programmed route in the event obstacles or terrain anomalies block its intended path. Accurate position and orientation estimation is essential to this process. Map fusion is critical to the robustness of the navigation process, as it enables the system to cope with sensor failures and missing data. In addition to this, deriving data from multiple sources compares sensor input to account for anomalies. If a sensor is damaged and not providing accurate data, the processing algorithms will accord that sensor input a lower degree of confidence and adjust its contribution to the overall weighting of the data characterization and mapping process.




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

Errors in terrain characterization can, in most cases, be attributed to errors not in the data acquired by the sensor, but by errors in position and orientation estimation. As demonstrated in Figure 6a which shows test data from Stanford University's 2005 entry 'Stanley', inaccurate pose can cause the vehicle to stop (Figure 6b) as all oncoming terrain will be perceived as not being traversable. The illustrations mark red terrain as not traversable, white is low cost terrain and grey areas are not known. The blue corridor is the DARPA assigned route. This pose error of less than 0.5 degrees in roll/pitch is enough to force the vehicle off the course if ignored [5]. In tests carried out by the Team, referenced terrain was erroneously classified

as not navigable 12.6% of the time with a height threshold parameter of $\delta=15$ cm [6].




Figure 5: An illustration of fused sensor maps

Planning and Vehicle Control

With reliable data from the POS LV integrated into the drive-by-wire systems of both Red Team Robots, pure pursuit tracking was made possible. However a method to maximize the performance of both vehicles was needed.





Figure 7: Topography data overlaid on imagery

Human drivers adjust to changing terrain / weather conditions in addition to interpreting a curves apex to maximize the efficiency of a turn rather than following a straight line denoting the curve and ‘jerk’ through it. This is not efficient and providing the robot with apex entry and exit information, in addition to terrain condition, are two ways in which



(a)

performance was maximized. Terrain data can be from maps or aerial imagery, but high 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 [7].

Detailed terrain topography can be acquired by collecting range scanner and vehicle position measurements while driving. This was done with an H1 Hummer called ‘Topographer’ which utilized a POS LV and laser scanner to derive drive by topography typically with .25m resolution and 1.5m accuracy. 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. An example of the detail of topography is shown in Figure 7.

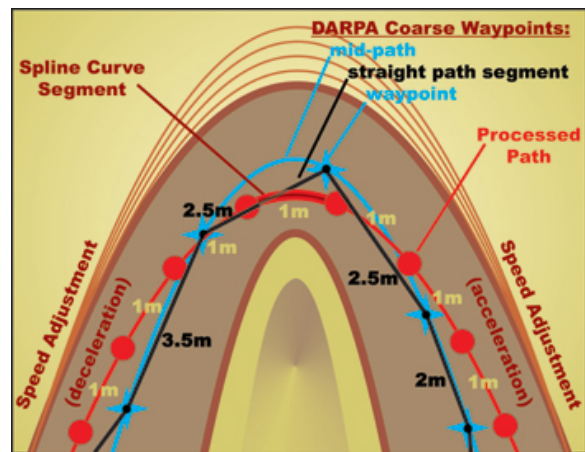
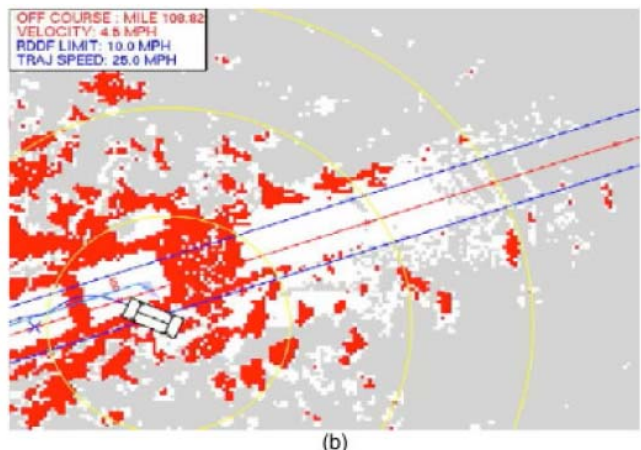


Figure 8: Result of pre-planning process

The entire robot preplanning process relies on accurate terrain and known parameters of vehicle performance to detail safe driving parameters while minimizing the time it takes to complete a section of the course [8]. The result of pre planning is illustrated in Figure 8. The black lines denote raw RDDF file waypoints and speed limits provided by DARPA. The red dotted path illustrates the route as edited by human planners heavily interpolating the original set of waypoints. These smoothed splines form the basis of navigating in and out of curves. Should obstacles be encountered, the robot generates its



(b)

Figure 6: Illustration of how small errors in position and orientation can provide erroneous terrain characterization

own viable path candidates using the interpolated path as a reference. Utilizing this approach, control points, and spline angle vectors that adjust to specify the location and orientation of a path.

4. THE 2007 DARPA URBAN GRAND CHALLENGE – THE CHALLENGE OF A DYNAMIC ENVIRONMENT

The 2005 Grand Challenge pushed participants to develop solutions for terrain perception and obstacle avoidance which required obstacle detection and avoidance at ranges of 40m directly in front of the vehicle. This next iteration of the Grand Challenge in an urban environment pushed the state of the art not only in perception, but object prediction and autonomous vehicle interaction in dynamic environments. Teams needed to contend with spurious GPS conditions in some areas of the course, ensuring the robots followed the rules of the road and interaction with 50 manned vehicles simulating traffic along with the other robots attempting to finish DARPA's three mandated mission tests. Each of these missions tested different core skills (parking, traffic collision avoidance, driving precision) and needed to be completed in the shortest amount of time. The intention of each of the missions was to test how the vehicles interacted with situations human drivers encounter on a regular basis. For example, there were several intersections where four vehicles were waiting at a stop light. The autonomous vehicle needed to know not only when it was safe to proceed, but deal with other vehicles which may malfunction and need to get out of the way while taking into consideration other traffic and the rules of the road. This required a new breed of sensors which provided the vehicle with a 360 representation of its surroundings. Such capability is shown in Figure 9.




Figure 9: 360 degree scanning LIDAR for situational awareness

The Teams demanded very high performance from their positioning and orientation systems. Pose estimation was critical to perception, planning, control and providing key data to the drive-by-wire systems of autonomous vehicles. As demonstrated in the following architecture from Tartan Racing, accurate position and orientation estimation was essential to the perception and world modeling routines constructed by the robot. As discussed previously, data fusion is a key factor in determining how successful the robot is in characterizing and interacting within its environment to achieve its mission. Pose estimation is provided to the behavior generation and motion planning routines which are bounded by the mission planning parameters programmed into the vehicle. In each phase of the

mission, the robot needed to integrate the composite representation of the world and understand what were safe and unsafe maneuvers given the changing targets around it impeding its route.

Given the missions and skills needed to be demonstrated by the robots, the benefits of pre-planning were not as profound as with the 2005 Challenge. Teams did require a substantial amount of data up-sampling from the sparse points provided in the DARPA RNDF file. However, instead of having 2 hours to prepare the vehicle, Teams only had 15 minutes between missions to prepare for the next portion of the race. Route planning was absolutely critical to finish the missions in the least amount of time, however the missions required much more processing of real time obstacle avoidance rather than following exact waypoints. This required high bandwidth, low latency data to be constantly available to the system especially for dynamic data fusion. Detecting a static obstacle is a simple process of determining where the target is located, what the lane corridor as defined by the Robot's sensors as compared to the RNDF is and what is the safest speed and steering angle around the obstacle to avoid it.

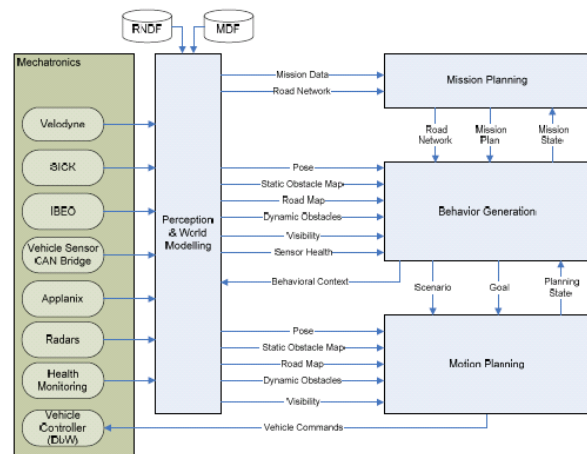


Figure 10: Tartan Racing Architecture [9]




Figure 10: Vehicle Tracking and Prediction

When the obstacle is dynamic, there are three fundamental challenges. The first is reliable position tracking relative to where the vehicle is and where it needs to go (in the local coordinate). Second, with accurate range and target bearing the robot can determine what lane the obstacle is in from the route network definition file (RNDF) or if it is off the road. This

provides vital information on what the obstacle is doing. Third, and most importantly, what is the predicted path of the obstacle. All of these functions depend heavily on highly accurate pose estimation.

Estimation of state and feature extraction all depend on accurate pose estimation. As mentioned previously, errors in pitch and roll of only 0.5 degrees can result in false characterization of terrain and obstacles. This is more critical in sensing obstacles far away rather than close to the vehicle. As the course presented obstacles in rapid succession the robots required accurate pose estimation to avoid colliding with them. However, errors in roll and pitch are more pronounced over longer distances and higher speeds. The absolute vertical error increases as the pitch error angle expands over the range of the sensor.

One of the key elements which determined success in this Urban Challenge was real time situational awareness and data fusion. Such a capability required two levels of characterization, that of the robotic vehicle in relation to the road and the dynamic obstacles on it. The challenge is illustrated in figure 10 where the robot is sensing the way to a waypoint, but encounters traffic around it. The vehicle must not only track and predict where it will go, but it must do this while tracking within its lane, sensing the terrain (road radius of curvature, grade / cross fall) to ensure any maneuvers are within the performance envelope and actually predict where the obstacle will move to. In the previous Grand Challenge robots had a choice of path candidates (in the Red Team example given previously, an onboard computer generates 's' splines or multiple path candidates immediately adjacent to the intended path of travel, all which are viable alternate routes taking into consideration the vehicle's dynamic state). Here, the path candidates around an obstacle need to be able to change rapidly and the vehicle will do most of the thinking.

5. RESULTS OF THE RACE

The DARPA Urban Grand Challenge took place in Victorville California at George AFB. The National Qualifying Event (NQE) saw thirty six Teams participate in a number of rounds designed to illustrate the requisite skills required to successfully complete the three DARPA missions. Of the thirty six Teams, eleven were qualified to participate in the final race on November 3rd, 2007. Of the eleven Teams, only six managed to finish all three missions. Applanix Corporation partnered with Tartan Racing, Stanford Racing and MIT to secure first, second and fourth place finishes.

Team Name	ID#	Vehicle	Type	Time Taken (h:m:s)	Result
Tartan Racing	19	Boss	2007 Chevy Tahoe	4:10:20	1st Place: averaged approximately 14 mph (22.53 km/h) throughout the course
Stanford Racing	03	Junior	2006 Volkswagen Passat Wagon	4:29:28	2nd Place: averaged about 13.7 mph (22.05 km/h) throughout the course
VictorTango	32	Odin	2005 Ford Hybrid Escape	4:36:38	3rd Place: averaged 13 mph (20.92 km/h) throughout the course
MIT	79	Talos	Land Rover LR3	6:00:00	4th Place.
The Ben Franklin Racing Team	74	Little Ben	2006 Toyota Prius	No official time.	Finished
Cornell	26	Skynet	2007 Chevy Tahoe	No official time.	Finished

Figure 11: DARPA Urban Grand Challenge Results

6. URBAN GRAND CHALLENGE LESSONS APPLIED TO REAL WORLD SCENARIOS

The goal of the Urban Grand Challenge is to apply the various technologies employed to successfully navigate the course to real world problems. Looking at this competition at its most fundamental level, these autonomous vehicles are mobile mapping platforms. The advances made here have significant implications for how mobile mapping data is used. Consider the automotive industry for example. Currently, GPS is utilized as a convenience feature utilizing GPS, map matching and odometer data to route a driver (albeit not very accurately) through GPS outages. When looking at position and orientation data in terms of driver assistance / active safety systems, the accuracy required changes dramatically. Data needs to be thought of in a layered approach for this application much like the data fusion discussed above. Base maps utilized by onboard computers need to be very accurate for sensors to determine dynamics in relation to a vehicle's current and predicted path so the vehicle can determine if a driver is making turns at unsafe speeds or passing through an intersection without stopping. By having detailed maps along with accurate position and orientation data, vehicles will be able to actively ensure the safety of passengers.

Military applications present another example of how vehicle automation saves lives. The Pentagon is aiming to have one third of its forces automated by 2015. This applies to combat forces as well as re-supply elements. Mobile mapping will become particularly automated in this field and employ several layers of data from different sources to achieve a particular mission. For example, UAVs employing LIDAR and other sensors will provide up to date intelligence for automated ground convoys traveling through hostile terrain. Ground vehicles utilizing their own LIDAR and optical sensors will map their way to an objective relying on accurate base maps and accurate position and orientation data.

7. SUMMARY

Accurate and reliable position and orientation data is a fundamental part of autonomous vehicle guidance and control. What we have shown is that even small errors in pose estimation can lead to erroneous terrain characterization which impacts vehicle performance. The significance of accuracy was highlighted in the Urban Grand Challenge where dynamic obstacles and terrain characterization in adverse GPS environments were key skills that robots demonstrated in order to successfully navigate the course and complete the three DARPA missions. Position and orientation data accuracy was essential to win the race which required sensor fusion and precise vehicle dynamic control to interact with a constant changing environment. These core elements will revolutionize how we think about mobile mapping in general. The precise location of roads, their geometry and roadside features will be essential elements for vehicle guidance and control, not just basic navigation. Accurate geospatial information and the real time interpretation of that information are essential elements for autonomous vehicles to demonstrate before such technology becomes mainstream.

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] Urmson, C., et al, *A Robust Approach to High Speed Navigation for Unrehearsed Desert Terrain*, Pittsburgh, PA 2006.
- [3] Peterson, K, *DARPA Grand Challenge 2005 Technical Paper*, Pittsburgh, PA 24 August 2005.
- [4] Whittaker, W, Nastro, L., *Utilization of Position and Orientation Data for Preplanning and Real Time Autonomous Vehicle Navigation*, GPS World, Sept 1, 2006.
- [5] Thrun, S., et al., *Stanley: The Robot that Won the DARPA Grand Challenge*, Journal of Field Robotics 23(9), 661-692, 2006.
- [6] *ibid.*
- [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] Urmson, Chris et al., *Tartan Racing: A Multi-Modal Approach to the DARPA Urban Challenge* Pittsburgh, PA April 13, 2007.