Preprints of the 19th World Congress
The International Federation of Automatic Control
Cape Town, South Africa. August 24-29, 2014

Development of an Autonomous Formula SAE Car with Laser Scanner and GPS

Thomas H. Drage*, Jordan Kalinowski*,
Thomas Bräunl*

*The University of Western Australia, EECE,
Crawley, Perth, WA 6009. (email: rev@theREVproject.com)

Abstract: This paper describes the development of a high level control system for an autonomous Formula SAE race car featuring an automotive LIDAR and fusion of a 6-DOF IMU with a consumer grade GPS. The control system is implemented as a multi-threaded C++ program with asynchronous IO, remote web interface and a comprehensive safety system. Autonomous driving is based upon recording and driving mapped waypoints with a Kalman filter used for predictive filtering of the vehicle's motion parameters while the LIDAR is used for road edge detection.

1. INTRODUCTION

Formula SAE is a long-running annual competition organised by the Society of Automotive Engineers with competition events in the U.S., Europe, and Australia (SAE, 2013). In former years Formula SAE has been a design competition only for petrol cars, but recently the new class SAE-Electric has been introduced.

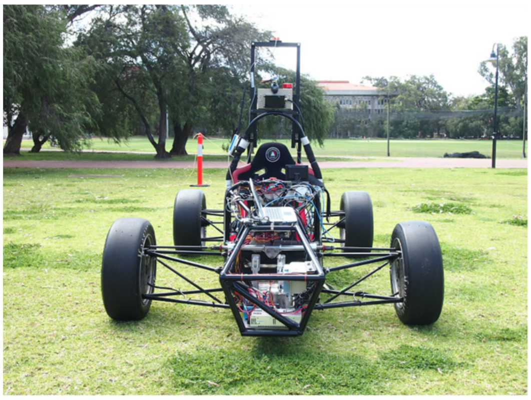

Figure 1 Autonomous Formula-SAE Electric Vehicle

In addition to two road-registered electric vehicles (EVs), a Hyundai Getz and a Lotus Elise, UWA's Renewable Energy Vehicle Project (REV) has built two electric SAE cars. The vehicle discussed in this paper features electric motors driving each of the two rear wheels via independent controllers and has full drive-by-wire control of the throttle, steering and (hydraulic) braking system.

This is outside the scope of the SAE competition which allows neither drive-by-wire or autonomous drive systems, however, the Autonomous SAE car will provide a platform for research into driverless vehicles at UWA with the ultimate goal of improving the efficiency and safety of modern transportation systems. In this particular project the Autonomous SAE Car can drive following a map consisting of "waypoints" and "fence posts". Maps can be recorded by either driving the course manually or through a GoogleMaps driven web-interface. Mapped driving is augmented by the use of a laser scanner for detection of obstacles as well as road edges.

Safety systems are essential in such a system, as the car weighs in excess of 250kg and is capable of driving at a speed of 80km/h. Safety systems are implemented as part of the high level controller as well as through independent hardware. Facilities for remote intervention and emergency stopping are provided through a wireless link to a base station as well as wired on the car itself.

2. REVIEW OF AUTONOMOUS AUTOMOBILES

Research into autonomous vehicles began in the 1980s with projects such as the EUREKA Prometheus Project in Europe and the United States' Autonomous Land Vehicle Project (Dickmanns, 2002). The DARPA Grand Challenges (Miller et al., 2006; Thrun et al., 2007) in 2004 and 2005 saw teams of autonomous vehicles competing to navigate a desert environment whilst the 2007 Urban Challenge (Bacha et al., 2008; Bohren et al., 2008; Montemerlo et al., 2008) required navigation of a road based course and adherence to traffic protocols. In Europe, the VisLab Intercontinental Autonomous Challenge in 2010 (Broggi et al., 2012) required an autonomous drive following a leader car from Italy to China. These competitions saw massive development of the field, with advanced technologies already becoming available for automotive use. Autonomous driving technology is evolving rapidly and is well on its way to finding commercial use in years to come. Google recently revealed that their fleet of autonomous cars had travelled 140,000 miles on US public roads without human intervention (Markoff, 2010). Locally, Rio Tinto plans to have 150 autonomous trucks supplied by Komatsu working in their Pilbara mining operations by 2015 (ENP, 2011).

Detailed research on control systems for autonomous driving has also been carried out at Stanford (Talvala et al., 2011) as well as at the University of Parma (Broggi et al., 2012) and through a collaboration of Spanish universities (Milanés et al., 2010). At UWA Lochlan Brown has conducted research into stability control using the same car and has provided evidence that computer control of only the throttle has resulted in significant advantages in increasing the vehicles stability (Brown, 2013). A substantial body of work exists concerning sensory techniques for use in autonomous vehicles which are built upon in this project. Examples include the implementation of sensors such as GPS on both road vehicles (Brown, 2010; Li, 2006; Qi and Moore, 2002) and in agricultural applications (Gomez-Gil, 2011), computer vision (Seigemund et al., 2011) and LIDAR (Fardi et al., 2003; Wijesoma et al., 2004; Zhang, 2010).

Recently technologies have matured and research into the potential of autonomous cars in racing has begun, with projects such as Stanford University's autonomous Audi TTS, which aims perform as to well as seasoned racing drivers. This project is of particular interest as its aims in using electronic control systems to drive “at the limits” of the car's mechanical abilities are similar to our project. A sophisticated suite of navigation sensors are used and have seen the car drive complex, long (20km) race courses (Brown, 2010).

3. SENSOR SYSTEMS

3.1 GPS

A variety of commercial products are available which provide built-in sensor fusion as well as functionality including D-GPS correction and RTK measurements. The cost is proportional to the accuracy and ranges from around $1000 for sub-meter accuracy (e.g. D-GPS), to around $5000 for decimetre accuracy (dual channel commercial SBAS) to tens of thousands for advanced survey-grade RTK systems (USU/NASA, 2010). Products such as the Applanix POS LV which include these high-level features have commonly been used in autonomous vehicles (Brown, 2010), however, with no public SBAS available in Australia and systems such as D-GPS and RTK still extremely expensive, a standard GPS device was selected for this project. It was found that the jitter of the GPS device used (QStarz BT-Q818X) resulted in a mean deviation from the average position of 1.5m over a 40 minute test. As a result of this, and the sensors limited resolution and update rate, it was found to be necessary to improve the performance via sensor fusion.

3.2 Fusion of IMU and GPS Data

An Xsens MTi IMU is used in this project in order to improve the accuracy and frequency of the data provided by the GPS. Early investigations showed that the accelerometer data is too noisy to integrate directly, with the calculated position diverging in just seconds. The internally filtered heading data is reasonably good but tests have shown that it is prone to errors introduced by magnetic field disturbances, particularly when used on the SAE car.

Fusion of the IMU heading and GPS track angle (heading) measurements is required due to the speed dependent accuracy of the GPS track angle and the limited absolute and dynamic accuracy of the IMU. If the car is not moving at an appreciable rate the reported GPS track angle will initially fail to change and jitters substantially at low speeds making it impossible to tell which way the car is facing when starting the autonomous drive. In order to determine the extent of this issue data was collected by driving in a straight line (directly south) and varying the cars speed and starting/stopping on multiple occasions. The resultant deviations from the average angle over the drive were then binned according to speed and standard deviation plotted.

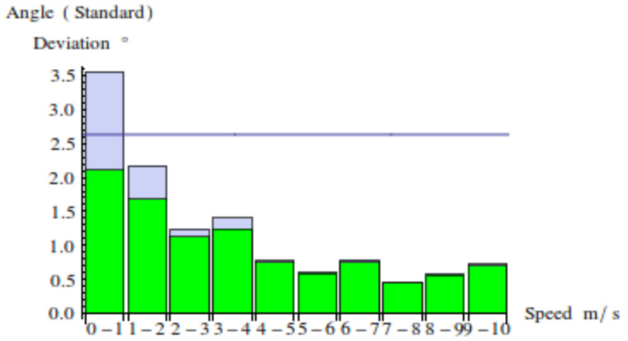

Figure 2 Unmodified GPS (blue) and fused (green) angle deviation vs. speed over a trip with IMU deviation (line)

It was observed as expected that the GPS track angle measurement is more reliable at greater speeds. The IMU mean deviation over the same trip was found to be 2.6° however, a considerable risk exists in relying on IMU data as magnetic disturbances can cause result in a loss of absolute accuracy and a “wandering” heading.

A speed-dependent weighted average with characteristics based on these observations was therefore implemented in order to ensure that the most accurate heading is always available. The weightings were calculated by applying the method developed by Elmenreich (2007) which seeks to minimise the variance of a linear combination (Z) of Gaussian random variables (X_G and X_I) which form the fused estimate:

$$Z = w_G X_G + w_I X_I \tag{1}$$

The weighting factors of the GPS and IMU (w_G and w_I) as well as the combined variance (σ_Z^2) were thus calculated as:

$$w_G = \frac{1}{1+\frac{\sigma_G^2}{\sigma_I^2}} \quad , \quad w_I = \frac{1}{1+\frac{\sigma_I^2}{\sigma_G^2}} \quad \Rightarrow \quad \sigma_Z^2 = \frac{1}{\frac{1}{\sigma_G^2}+\frac{1}{\sigma_I^2}} \tag{2}$$

These weightings were calculated for each speed bin and a piece-wise linear relationship between speed and the required weighting ratio was developed and implemented in the control software (see Figure 5). In order to ensure a reliable heading estimate at stationary, the linearised implementation applies zero weighting to the GPS from zero to 0.25m/s. The most significant result of this approach is the ability to obtain a stable heading throughout the process of acceleration from stationary whilst maintaining absolute accuracy at speed. The effect of reducing the variance of the heading data through this technique can be noted in the overlay in Figure 2.

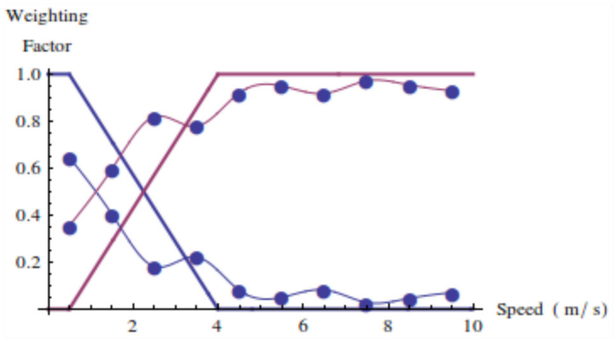

Figure 3 Calculated and linearised weighting factors

Note that it is necessary to apply a correction due to steering as the IMU and GPS angle measurements are no longer physically identical during turns. There is also a desire for faster heading updates and so the high-rate IMU data is used to interpolate the calculated heading by adding the difference between the IMU derived heading at the last GPS fix and an intermediate time to the heading calculated by the weighting formula at the last GPS fix.

A simple filtering algorithm was been implemented in order to smooth and improve the accuracy of the car's positioning data. This system can be considered to be a variation upon a loosely coupled GPS-INS (Inertial Navigation System) fusion with the acceleration data used directly (Qi and Moore, 2002) and was chosen over more complex sensor fusion algorithms such as (Li, 2006) due to the low computational loading, simplicity of implementation and applicability to the low cost hardware employed in this project. Firstly the acceleration data is transformed into a North-East-Down coordinate system using the orientation data (read as a cosine matrix from the IMU) and then combined in a Kalman Filter with the GPS velocity information. The estimated velocity from this filter is then combined in another Kalman Filter with the position data, which is then used to compute the cars trajectory. Both filters currently use a linear free body model for the system. Initial tests have shown that this method is reasonably effective however future work on this project is expected to focus on improving the model used for the vehicle's dynamics and comparing the use more advanced GPS-INS fusion techniques.

3.3 Laser Scanner

The LIDAR system used in this project is an IBEO Lux automotive LIDAR. The IBEO sensor features measurement in four layer layers and has internal data processing functionality including object detection and classification. The range of the sensor is 200m with an 85° horizontal view and can provide data at up to 12.5Hz. In this project the sensor was mounted on a specially constructed bracket above the car's roll cage which provides a stiff mounting and is adjustable in vertical angle. LIDAR scan data is used to detect the road edges by determining the horizontal extent of the points which make up the roads surface. The road points on a bitumen surface tend to be arranged collinearly with a small deviation whereas points belonging to uneven surfaces (such as grass and curbs) tend to be scattered and their arrangement depends on the contour of the surface away from the road edge.

The algorithm implemented in this project operates by first identifying a candidate group of points close to the centre of the scan data which meet the slope condition (i.e. the slope of a line through these points is less than a pre-set value). This group is then expanded iteratively, with a least-squares linear regression performed at each step. By minimising the square residuals between the fit line and the data, this technique obtains the most appropriate line (y) for the given data set (x_i,y_i), the success of which can be measured by means of the product-moment correlation coefficient (r). In this case, the slope (b) and r^2 values are of relevance and are tabulated.

Specifically:

$$y = (\bar{y} - b\bar{x}) + bx \quad , \quad r = \frac{s_{xy}}{s_x s_y} \text{ where,}$$

$$b = \frac{s_{xy}}{s_x^2} \quad , \quad s_{xy} = \frac{\sum x_i y_i}{n} - \bar{x}\bar{y} \quad , \qquad (5)$$

$$s_x^2 = \frac{\sum x_i^2}{n} - \bar{x}^2 \quad , \quad s_x^2 = \frac{\sum y_i^2}{n} - \bar{y}^2$$

This technique is performed independently for the left and right hand sides of the candidate point group to allow for the fact that roads are often sloped about the centre (e.g. to let water run off) which results in improved accuracy. In the initial implementation, the road-edges were then determined to be the outer edges of the candidate groups which maximised the respective correlation coefficients and met the slope condition. However, it was found that this approach was not reliable across various different road scenarios.

Thus, an improved algorithm was developed which adds heuristics to recognise features which result in extremely poor correlation (e.g. curbs) and uses predictive filtering to smooth data as well as to provide a priori knowledge of the edge location to the algorithm at each step. If the maximum correlation point is found to be in an unreasonable position compared to the previous position, either a poor correlation point or a *local* maximum correlation point will be selected closer to the expected value. This filter takes the form of a Kalman filter with state variables x (the measured quantity), the lateral position and v (unmeasured), the lateral velocity of the road edge relative to the car. The state transition matrix and observation matrix are thus:

$$\begin{bmatrix} x_n \\ v_n \end{bmatrix} = \begin{bmatrix} 1 & \Delta t \\ 0 & 1 \end{bmatrix} \begin{bmatrix} x_{n-1} \\ v_{n-1} \end{bmatrix} , \quad H = \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix} \qquad (6)$$

The plots below show scan data from the bottom layer collected by the IBEO sensor and projected onto an XY plane. The red line indicates the section of the plot identified as road after applying the algorithm.

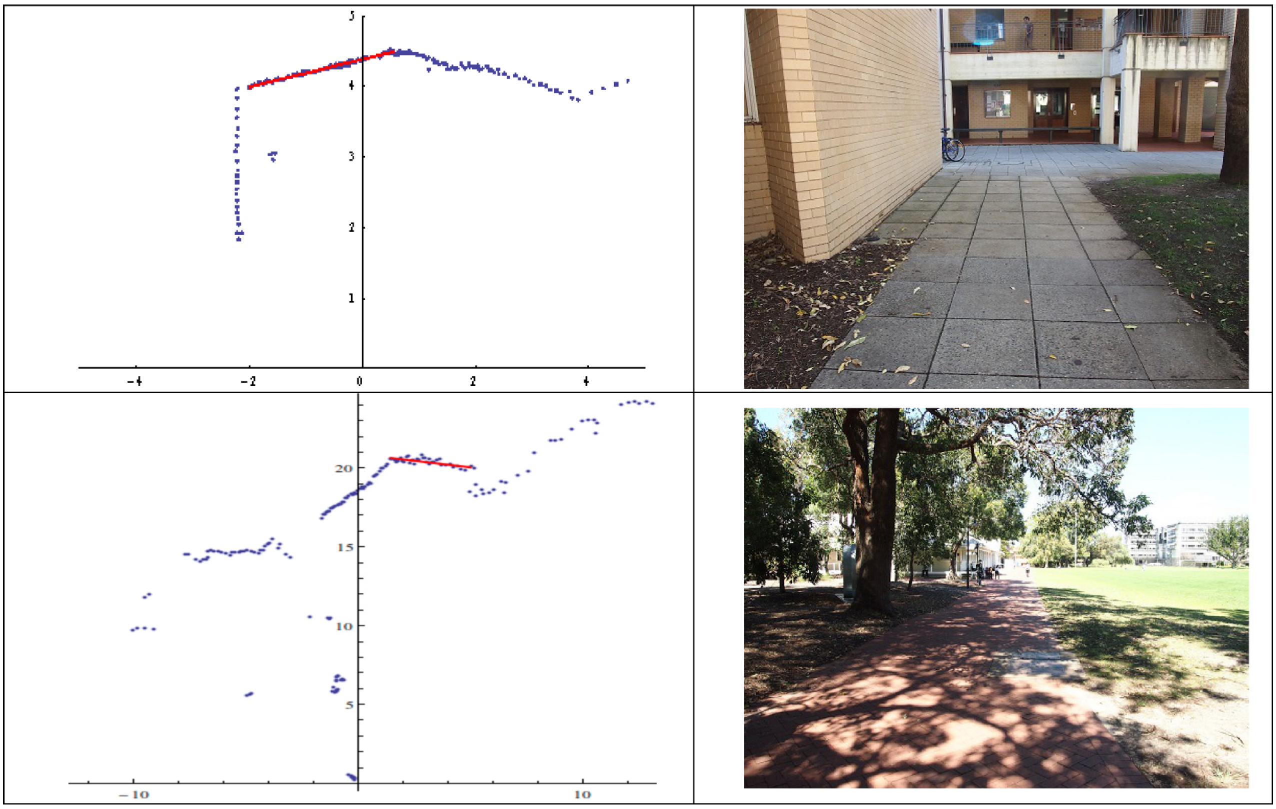

Figure 7 Samples of road-finding algorithm output (LIDAR scan data looking forward, distances in metres).

The algorithm was able to successfully identify the road edges in these two static test cases as well as in three dynamic tests. Identification of the edges whilst in motion along a paved road with grassed edges revealed an accuracy of 0.06±0.31m whilst driving along a curbed bitumen road gave 0.10±0.16m. In a further test, the algorithm was able to identify the edges of a curbed road when an oscillatory pattern was driven to simulate dynamic conditions.

The IBEO sensor also provides object data. This data is projected onto the cars running map data as “fence posts” in order to mark obstructions in the path. The graph below shows the range of obstructions (red circles) detected while parked in the laboratory overlaid on the raw scan data.

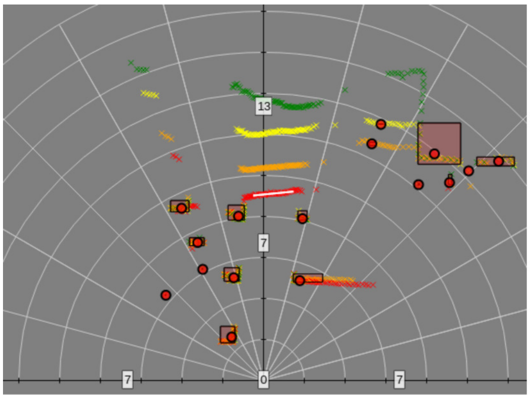

Figure 8 Identification of objects outside the laboratory

4. SYSTEM DESIGN

The Autonomous SAE Car control system consists of two segments – a “low level system” which handles physical outputs (e.g. signal to the motor controllers, operation of the brake servomotor and steering control system) and the “high level system” which encompasses processing sensor data, communication with the human operator, calculation of control parameters and the subsequent instruction of the low level system.

The high level control system is centred around a small PC running Linux located on the car – the control program itself is implemented in C++ and operates asynchronously. Ethernet connectivity is present on the car and a Wifi link is used to communicate with the base station, all other sensors and outputs communicating using simple serial interfaces. Software on a “base station” PC generates a safety “heartbeat” signal which is able to be interrupted with a large physical stop button. The autonomous driving features such as map collection and selection as well as information outputs are provided in a web interface which communicates asynchronously with the main control system.

Alongside the two control sub-systems, a “safety supervisor” hardware device was implemented which comprises a PIC microcontroller and several interface circuits. It provides interlocks which ensure that the car will not accelerate immediately after drive power is applied or if base station communications are present. It is able to disconnect the drive power or apply the emergency brake if it, or either of the other two sub-systems, detects an emergency condition.

5. TRAJECTORY CALCULATION

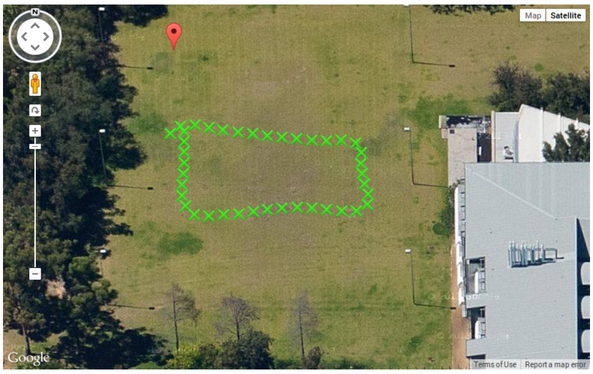

Figure 11 Waypoint data recorded automatically displayed on the GoogleMaps interface.

Maps (captured either by manual input or by driving the course) are stored in text files and contain latitude and longitude coordinates for the maps "datum" followed by waypoints and "fence posts" in Cartesian coordinates with distances in meters referenced to the datum. During recording and driving the map is loaded into RAM and can be edited dynamically within the control program. Fence posts represent obstructions or boundaries which the car is not allowed to travel near to and an emergency stop will be initiated should the car come within a fixed distance of them. The desired speed is set based on the current turn radius, giving slower speeds when cornering.

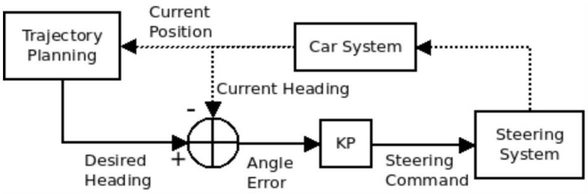

Figure 12 Steering controller overview

The car's trajectory is calculated each time an updated position measurement is received and a cubic spline is calculated from the current position through the next several waypoints. This allows for generation of a realistic, smooth, trajectory and solves stability issues that would occur should the car not arrive at each waypoint facing a direction reasonable for continuing the required path. The polynomial splines used are defined piece-wise, parametrically, between the cars current position, the first waypoint and each successive waypoint. They are fitted against time as in the case of a general map the y coordinate is not a function of x and because the motion of the vehicle is required to be smooth in time as well as in space. In order to achieve this, a "pseudo-time" scale is developed based upon the assumption of constant velocity between the waypoints. The required bearing is then calculated and used as the set point for a PID loop. This PID loop compares measurements of the cars current heading (from the fusion algorithm) to the heading required by the trajectory and manipulates the steering angle in order to bring the cars heading in line. By tuning the PID controller parameters factors such as the angular velocity of turning and response to changing trajectories are easily adjusted.

The value of the brake travel and throttle level required are computed by two separate PID controllers. This enables tuning of each action independently and is required given the nature of operation of the two functions. Braking occurs in response to a high-speed condition and the brake is operated only until the desired and actual speeds match. Manipulation of the throttle is incremental since the PID controller is required to converge on a non-zero throttle value, thus the PID output is summed with the current throttle set position to give the new set position at each step.

6. EXPERIMENTS

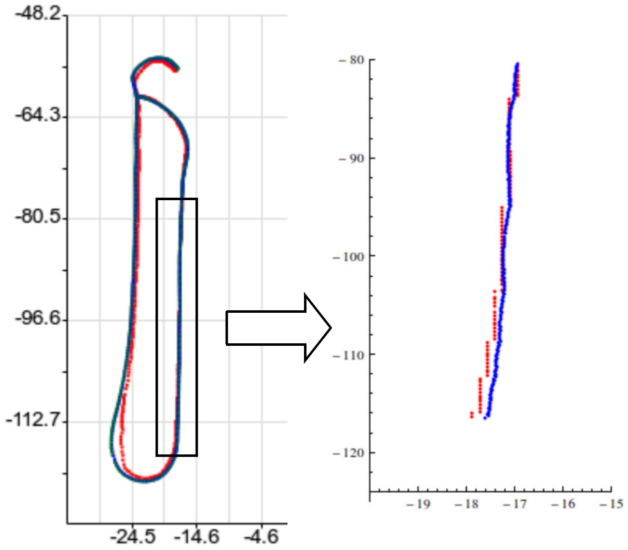

Figure 13 Manually driven path – GPS (red) and Fused (blue) position measurements (metres)

In order to verify the operation of the sensor fusion algorithms, a series of tests were conducted where a loop consisting of two straight sections aligned North-South (marked on the ground) was driven. This experiment reveals the effectiveness of the position filtering algorithm – the GPS position resolution and update rate are somewhat limited, whilst the fused position is substantially smoother, masking the GPS modules shortcomings. The GPS data around the South-North turn in this test drive shows a deviation of approximately two metres as well as increased noise, however, the filtered position in this region is smooth and more importantly physically consistent with the pattern being driven in the turn. It is therefore evident from this data that the filtering algorithm is of benefit.

A test was conducted by autonomously driving an automatically recorded map with a waypoint radius of 2.5m. The path taken is shown in figure 12 with a mean accuracy (distance from path to waypoint) of 80cm and a closest approach of 10cm. Note that the "cutting" of the corner at the top right of the path is due to the combination of the radii of the waypoints in that section.

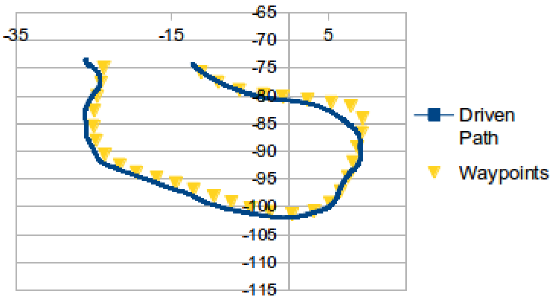

Figure 15 Autonomous driving test trajectory (metres)

7. CONCLUSIONS

This project shows significant promise as a platform for the development of autonomous driving technology and it is anticipated that with development of the underlying systems completed future work will focus on refining and improving the methods presented here. The safety systems and UI implemented have provided safe operation without impact upon flexibility during testing. Future work on this project will encompass optimization of driving behaviours so that race grade performance can be achieved.

REFERENCES

Bacha, A., C. Bauman, R. Faruque et. al. (2008). "Odin: Team VictorTango's entry in the DARPA Urban Challenge", *Journal of Field Robotics*, **25**(8), pp. 467-492.

Bohren, J., T. Foote, J. Keller et al. (2008). "Little Ben: The Ben Franklin Racing Team's entry in the 2007 DARPA Urban Challenge", *Journal of Field Robotics*, **25**(9), pp. 598-614.

Broggi, A., P. Medici, P. Zani, A. Coati and M. Panciroli (2012). Autonomous vehicles control in the VisLab Intercontinental Autonomous Challenge, *Annu. Reviews in Control*, **36**(1), pp. 161-171

Brown, L. (2013). "Improving Performance Using Torque Vectoring on an Electric All-Wheel-Drive Formula SAE Race Car." BE thesis, UWA.

Brown, T. (2010). "Handling at the limits", *GPS World*, vol. 21, no. 8, pp. 38-41.

Dickmanns, E 2002. "The development of machine vision for road vehicles in the last decade", *Intelligent Vehicle Symposium*, pp. 268-281.

Electronic News Publishing (2011). "Komatsu and Rio Tinto Enter into Agreement for 150 Autonomous Truck Deployment into Pilbara Iron Ore operations by 2015", *ENP Newswire*, 4 Nov. 2011.

Elmenreich, W. (2007). "Fusion of Continous-valued Sensor Measurements using Confidence-weighted Averaging", *Journal of Vibration and Control*, **13**(9-10), pp. 1303-1312.

Fardi, B, U. Scheunert, H. Cramer, G. Wanielik (2003). "Multi Model Detection and Parameter-based Tracking of Road Borders with a Laser Scanner", *IEEE IV2003 Intelligent Vehicles Symposium,* pp. 95-99.

Gomez-Gil, G., S. Alonso-Garcia, F. J. Gomez-Gil, T. Stombaugh (2011). "A Simple Method to Improve Autonomous GPS Positioning for Tractors", *Sensors*, **11**(6), pp. 5630-5644.

Li, Y., J. Wang, C. Rizos, P. Mumford, W. Ding (2006). "Low-cost Tightly Coupled GPS/INS Integration Based on a Nonlinear Kalman Filtering Design". *Proceedings of the 2006 National Technical Meeting of The Institute of Navigation,* Monterey, CA, pp. 958-966.

Markoff, J. (2010). "Google Cars Drive Themselves, in Traffic", *The New York Times*, Oct. 9, 2010 [Online]. Available: http://www.nytimes.com/2010/10/10/science/10google.html. [Accessed: Aug. 17, 2013]

Milanés, V., D. F. Llorca, B. M. Vinagre, C. González, M. A. Sotelo (2010). "Clavileño: Evolution of an Autonomous Car", *Annu. Conf. Intelligent Transportation Systems*, Madeira Island, Portugal, pp. 1129-1134.

Miller, I., S. Lupashin, N. Zych et al. (2006) "Cornell university's 2005 DARPA grand challenge entry", *Journal of Field Robotics*, **23**(8), pp 625-652.

Montemerlo, M., J. Becker, H. Dahikamp, et al. (2008) "Junior: The Stanford entry in the Urban Challenge", *Journal Of Field Robotics*, **25**(9) pp. 569-597.

Qi, H. and J. B. Moore (2002). "Direct Kalman filtering approach for GPS/INS integration". *IEEE Transactions on Aerospace and Electronic Systems,* **38**(2), pp. 687-693.

SAE International (2013). "About Formula SAE Series" [Online]. Available: http://students.sae.org/cds/formulaseries/about.htm. [Accessed: Oct. 20, 2013]

Seigemund, J., U. Franke, W. Forstner. "A temporal filter approach for detection and reconstruction of curbs and road surfaces based on Conditional Random Fields" (2011). *2011 IEEE Intelligent Vehicles Symposium*, pp 637-642.

Talvala, K., K. Kritayakirana, J. C. Gerdes (2011). "Pushing the limits: From lanekeeping to autonomous racing", *Annu. Reviews in Control*, **35**(1), pp. 137-148.

Thrun, S., M. Montemerlo, H. Dahllamp et. al. (2007). "Stanley: The Robot That Won the DARPA Grand Challenge" in *The 2005 DARPA Grand Challenge*, M. Buelher, K. Iagnemma, S. Signh, Eds., Berlin: Springer, pp. 1-43.

USU/NASA Geospatial Extension Program (2010). "High-End DGPS and RTK systems" [Online]. Available: http://extension.usu.edu/nasa/files/uploads/gtk-tuts/rtk_dgps.pdf. [Accessed: Sep. 3, 2013].

Wijesoma, W.S., K. R. S. Kodagoda, A. P. Balasuriya (2004). "Road-Boundary Detection and Tracking Using Ladar Sensing", *IEEE Transactions on Robotics and Automation*, **20**(3), pp. 456-464.

Zhang, W. (2010). "LIDAR-Based Road and Road-Edge Detection". *2010 IEEE Intelligent Vehicles Symposium*, pp 845-848.