

Autonomous Ground Vehicle Technologies Applied to the DARPA Grand Challenge

Carl D. Crane III*, David G. Armstrong Jr.*
Mel W. Torrie**, and Sarah A. Gray**

* Center for Intelligent Machines and Robotics
Dept. of Mechanical and Aerospace Engineering
University of Florida, Gainesville, Florida, USA
(Tel : 352-392-9461; E-mail: ccrane@ufl.edu)

**Autonomous Solutions, Inc., Young Ward, Utah, USA
(Tel 435-757-5483; E-mail: sarah@autonomoussolutions.com)

Abstract: This paper describes the design, development, and performance testing of an autonomous ground vehicle that was developed to participate in the DARPA Grand Challenge that was held in March 2004. The authors of this paper are members of Team CIMAR which was one of twenty five teams selected by DARPA to participate in a competition to develop an autonomous vehicle that can navigate from near Los Angeles to near Las Vegas at speeds averaging twenty miles per hour. Most of the event was held on open terrain and trails in a rocky desert environment. This paper describes the overall system design and the performance of the system at the event.

Keywords: autonomous vehicles, navigation, sensor fusion, obstacle avoidance

1. INTRODUCTION

The Grand Challenge was established by the Defense Advanced Research Projects Agency (DARPA) in order to encourage researchers to accelerate the development of autonomous vehicle technologies that can be applied to military requirements [1]. The event consisted of three parts, i.e. (1) application and acceptance into the event, (2) qualification, inspection, and demonstration (QID), and (3) the actual race from Barstow, CA to Primm, NV. The team that was able to complete the course first within a ten hour time frame would be awarded a prize of one million dollars.

Teams were selected for the event based on DARPA's evaluation of the team's submitted technical report that described their vehicle and approach. Twenty five teams were invited to participate in the event from a total of approximately eighty technical reports that were submitted. Team CIMAR was notified of acceptance into the event in mid December 2003.

Twenty three of the accepted teams arrived for the QID event at the California Speedway in Ontario, CA on 8 March 2004. A test course was set up on which vehicles would have to demonstrate autonomous path following and obstacle avoidance. Only those vehicles that demonstrated an ability to navigate the course would be allowed to participate in the actual race on 13 March 2004. Each team was guaranteed at least two opportunities to run their vehicle on the QID track.

Fifteen of the teams, including Team CIMAR, were selected to participate in the actual race event. The vehicles were transported from the California Speedway to the starting area in Barstow, CA on the day before the race. Two hours before the start of the race, each team was given a data file that contained approximately three thousand waypoints along with a corridor width for each pair of waypoints. Teams could use the two hour period to plan a path for the vehicle based upon any a priori data such as trails. The first vehicle to start on the course departed at 6:30 am. Other vehicles were started at five minute intervals. The subsequent sections of this paper will describe the vehicle platform, sensing system, and integration architecture that were used by Team CIMAR during the event.

2. SYSTEM DESIGN

The problem statement as presented by DARPA required that a vehicle be able to travel across desert terrain for a distance of up to 250 miles within a ten hour time period. Underpasses, bridges, railroad crossings, cattle grates, and fences were possible items that could be encountered on the course. The course was defined by a series of waypoints with a corridor width between waypoints (see Figure 1) to keep the vehicles bounded to a limited region. Exiting the defined corridor would be cause to stop the vehicle.




Fig. 1 Corridor definition

The following subsections present an overview of the mobility platform, the sensors used to detect obstacles and smooth terrain, and the architecture that was used to integrate the varied components of the system.

2.1 Vehicle platform

Cost was the limiting factor in selecting a vehicle platform to use for the event. A 1993 Isuzu Trooper that had been in an accident was chosen to become what is now referred to as the NaviGATOR. Vehicle modifications included adding sensors, actuators, and controllers to perform closed loop control of steering, throttle, and transmission together with a driving and emergency hydraulic brake system and backup fuel system for long distances and redundancy. A 32-Bit PhyCore MPC565 PowerPC microcontroller, a typical microcontroller found on most automobiles, was selected to perform the job of the Vehicle Control Unit (VCU). The

VCU is responsible for running many processes to include vehicle localization by fusing onboard motion sensors (WAAS GPS, quadrature shaft encoder, and an inertial/magnetic orientation sensor), closed loop path tracking and velocity control, and low-level hardware control. The I/O capabilities (10 Mbit/s Ethernet CS8900A controller, four UARTs, three on-chip CAN controllers, and 40 10-bit A/D channels) and the integrated 64-bit Floating Point Unit (FPU) make the MPC565 an ideal processor for the VCU function. Figure 2 shows the vehicle as originally purchased and Figure 3 shows the final NaviGATOR platform.



Fig. 2 Vehicle in original state



Fig. 3 NaviGATOR

2.2 System architecture

The system architecture is based on the Joint Architecture for Unmanned Systems (JAUS) Reference Architecture, Version 3.0 [2]. JAUS defines a set of reusable components and their interfaces. In order to ensure that the architecture will be applicable to the entire domain of mobile systems, the following four characteristics have been considered:

1. Vehicle platform independence. In order for JAUS components to be interoperable, no assumptions about the underlying vehicle or its means of propulsion are made.
2. Mission isolation. The JAUS components can typically be assembled such that a variety of missions can be supported.
3. Computer hardware independence. No assumption of or requirement of particular computer hardware is made. This allows for future adaptability and enhancement as new computer hardware becomes available in the future.
4. Technology independence. This is similar to the computer hardware independence, but focuses more on the technical approach rather than the computer hardware. For example, there are many approaches that could be used to determine vehicle position and

orientation. No one approach, such as for example GPS, inertial dead reckoning, or landmark based navigation is specified.

Figure 4 presents a simplified representation of the system architecture that was used in the vehicle. The purpose and functionality of each of the components shown in the figure is presented in the following sections.




Fig. 4 Architecture schematic

2.3 Primitive driver

The Primitive Driver component is the interface to the vehicle actuators. The component does not imply any particular platform type such as tracked or wheeled, but describes the mobility in six degrees of freedom using a percentage of available effort to translate the vehicle in each axis direction and to rotate the vehicle about each axis.

A coordinate system is attached to the vehicle where the x axis points forward and the z axis points down. The primitive driver receives two primary commands; a propulsive wrench and a resistive wrench that are defined in terms of this coordinate system. Each wrench is comprised of six scalar values and can be written as

$$\hat{\mathbf{w}} = [f_x, f_y, f_z; m_x, m_y, m_z]^T. \quad (1)$$

Due to the non-holonomic constraints associated with the vehicle motion, the system can only respond to propulsive efforts to translate in the x direction (f_x) and rotate about the z direction (m_z). For this case these two values are mapped directly to the throttle and steering actuators. For the resistive wrench, the vehicle can only respond to the f_x component which requests that translation along the x direction be restricted. This value is thus mapped to the vehicle brakes. The other components of the propulsive and resistive wrench are ignored.

It is important to note that the propulsive and resistive wrench commands are open-loop. By this it is meant that no velocity (linear or angular) of the vehicle is implied. Other components which perform feedback sensing and control are incorporated to accomplish closed-loop control of vehicle motion.

2.4 Sensor systems

Three categories of sensors are employed on the vehicle, i.e. (1) position and orientation sensors, (2) obstacle detection sensors, and (3) terrain evaluation sensors. The positioning sensors consisted of a NavCom StarFire GPS together with a Smith's Aerospace North Finding Module (NFM) and Inertial Navigation System (INS). The StarFire GPS uses a satellite based augmentation system to obtain position information

within 10 cm of truth. The NFM does not have the limitations of magnetic systems and in effect has the performance comparable to high accuracy ring laser gyros. Together the systems provided accurate position and orientation information for the system.

Three sensors were used for obstacle detection. The first was a 2D SICK LMS200 Laser Range Finder attached to a spinning shaft with an encoder and a slip ring that transmits data at 500 kbaud to a Digital Signal Processor (DSP) where the information is processed for obstacle detection and avoidance. It creates a 3D representation of the world that is orthogonally projected into a 2D space. Positive and negative obstacles as well as drivable and non-drivable slopes are estimated from the data. Figure 5 shows the rotating SICK lidar and Figure 6 shows a typical range data image.




Fig. 5 Sensor suite



Fig. 6 Output of 3D lidar

The second obstacle detection sensor was a stereo vision system that was manufactured by Videre Design. The system utilized 12.5 mm focal length lenses which provided a horizontal and vertical field of view of 50 degrees and 38 degrees, respectively. Image data are transferred via an IEEE1394 interface to a single board computer. The single board computer utilizes SRI International's Small Vision System to handle image rectification, correlation, and ultimately extraction of three-dimensional data.

The third obstacle detection sensor consisted of three short-range Preview Radar Systems from Preco that were mounted on the front of the vehicle. These sensors were used primarily to modify velocity when objects were detected. The closer the object, the slower the commanded velocity giving the obstacle detection and avoidance systems more time to make corrective decisions. A prototype long-range radar unit from Preco provided additional information on free space.

Two sensor systems were utilized to evaluate the smoothness of terrain and to identify regions of good traversability. A stationary SICK lidar (shown in Figure 5) was mounted to provide range information at a location on a theoretical flat ground-level plane twenty meters in front of

the vehicle. As the vehicle moved range data would be grouped as belonging to $0.5m \times 0.5m$ grid cells in front of the vehicle. For each grid cell, the best fitting plane would be calculated. The standard deviation of the data points from this plane would give an indication of the smoothness of the grid cell area. The slope of the best fitting plane would also be used to evaluate the traversability.

Monocular cameras and vision processing algorithms were used as the second sensor system to identify traversable terrain. Here, parts of images that closely resemble the region of the image directly in front of the vehicle were identified. Figure 7 shows a sample image. On the left is an original image and on the right is shown all regions that resemble the training region that is located directly in front of the vehicle. Periodically reclassifying based on the training image allowed for a robust determination of the traversable area.




Fig 7 Image processing to evaluate terrain traversability

2.5 Sensor arbitration

An important aspect of the project was to integrate all the sensor information into a format that could be acted upon to guide the vehicle. A sensor fusion approach was developed whereby the output of all sensors would be in a common grid based format. The environment around the vehicle was modeled by a 120×120 grid where each grid cell was $0.5m \times 0.5m$ in size and where the orientation of the grid lines was always maintained parallel to the north-south and east-west lines, no matter what the current orientation of the vehicle was. The vehicle was situated at the center of the grid and the grid data would be appropriately shifted as the vehicle moved to an adjacent grid cell in order to keep the vehicle located at one of the center cells.

Every sensor output an estimate of the traversability of each grid cell. For the three dimensional obstacle avoidance sensors (rotating lidar and stereo vision) the three dimensional point data was projected onto the grid plane. The terrain traversability sensors, i.e. fixed lidar and monocular vision, estimated traversability based on smoothness of the spatial plane fitted to the range data or the commonality in appearance of pixels in the grid cell to those directly in front of the vehicle, respectively.

Figure 8 illustrates the sensor arbitration approach. The top rendering shows the vehicle moving through a three dimensional environment. The bottom-most grid shows the output of a process that monitors the lane corridor to make sure that the vehicle does not stray outside of it. The red cells indicate regions that cannot be traversed while the green cells identify regions that can be traversed at high speed. The next grid above shows output from the sensors that are looking to identify smooth terrain. Here yellow grid cells indicate regions that are traversable, but not at high speed. White grid cells identify regions where the traversability is unknown. The next grid above shows similar output from the obstacle detection sensors and the top grid illustration shows the result of integrating all the grid information into one grid.

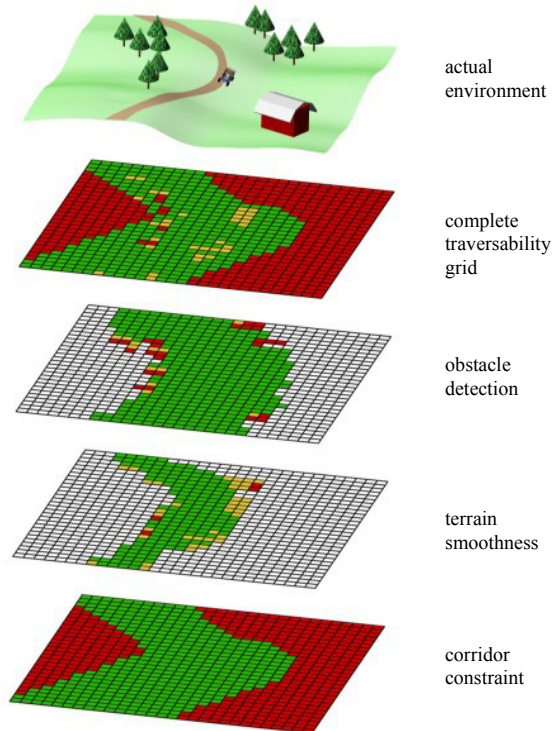


Fig 8 Sensor integration

2.6 Path following

The path segment driver depicted in Figure 4 acts to control the vehicle to follow a given path segment at a given speed. Many control techniques have been developed over the years by many researchers to accomplish this type of task. The implemented component performs closed-loop velocity control by comparing the sensed speed of the vehicle to the desired speed. The f_x component of the propulsive wrench is then adjusted to impact the speed of the vehicle.

Closed-loop steering control is accomplished by comparing the vehicle's sensed position and orientation to a point on the path that is ahead of the vehicle. The current steering wheel value is adjusted via the m_z component of the propulsive wrench in order to minimize the errors in position and orientation.

2.7 World model

The World Model component depicted in Figure 4 was developed in order to be a repository of information, i.e. a priori road data that was obtained from a USGS database as well as latitude and longitude data of trails and roads that were driven months before the competition. The World Model component also was given the course corridor that was defined by the given waypoints and corridor widths.

Currently, the World Model is not updated with sensed environment data. Current efforts are focused on accomplishing this task.

2.8 Planning

A Global Mission Planner and an onboard real-time Reactive Planner were utilized for planning operations during the Grand Challenge. The Global Mission Planner took as input the corridor data stored by the World Model and output an initial global path for the NaviGATOR to follow from start to finish. Previously logged GPS road data and GIS road

data was utilized in path generation in an attempt to keep the vehicle on roads throughout the majority of the event.

The onboard Reactive Planner utilizes the fused sensor data in the form of a drivability grid to check the current path for intersections with non-drivable regions of the grid. If an intersection is detected, a new path is planned to avoid the obstruction and return to the original path plan.

3. TESTING AND RESULTS

The integration schedule for the project was very aggressive. Personnel at Autonomous Solutions, Inc. in Young Ward, Utah installed actuators on the throttle, steering, brakes, and transmission and implemented computer control of the vehicle via the Primitive Driver component. The vehicle was then transported to Gainesville, Florida where it arrived on 21 January 2004. Sensor systems were installed and system integration and testing were conducted at the facilities of the Gainesville Raceway until the vehicle and team began traveling west for the competition on 21 February 2004.

After additional testing was performed in Utah, tests in the desert were held at the Stoddard Valley off-highway vehicle area during the period 4-8 March 2004. Figure 3 shows the test environment.

On 8 March 2004 the NaviGATOR arrived at the California Speedway for Qualification, Inspection, and Demonstration (QID) test runs. Each team was guaranteed a minimum of two runs on an obstacle course that is depicted in Figure 9. The NaviGATOR vehicle successfully navigated the course until reaching the underpass. During this run the NFM was not integrated into the system and GPS was lost at the underpass requiring that the vehicle be stopped. The NFM was subsequently added to the system to address this problem, but other technical issues prevented the team from completing the course on the final day of the QID event.




Fig 9 QID obstacle course

Team CIMAR was one of fifteen teams that were invited to participate in the actual desert race that was held on 13 March 2004. The course waypoints were given to the teams at approximately 4 am on the day of the race and each team had two hours to plan a route through the corridor. Figure 10 shows the NaviGATOR at the starting gate and Figure 11 shows the given course.



Fig 10 NaviGATOR at starting gate



Fig 11 DARPA Challenge course

No team completed the course which was approximately 150 miles in length. The furthest that any vehicle traveled was approximately 7 miles. The NaviGATOR vehicle moved well from the starting line, however after approximately 0.5 miles, while still within the waypoint corridor, it traveled parallel to a trail and became entangled in barbed wire causing it to stop.

4. CONCLUSION

The NaviGATOR vehicle performed well at the QID event. Obstacle avoidance and path tracking were successfully

demonstrated. None of the fifteen competitors completed the off road course. The furthest that any vehicle moved along the course was 7 miles of the total 150 mile distance, but the event was an astounding success from the competitor's vantage point.

Overall the Grand Challenge did succeed in accelerating the development of autonomous vehicle technologies. The lessons learned from the event are being applied towards next year's entry. Continued advancements in the areas of mobility, sensing, data interpretation, and planning will ultimately make the vision of autonomously navigating vehicles a reality.

ACKNOWLEDGMENTS

The authors gratefully acknowledge the contribution of the team sponsors: Autonomous Solutions, Inc., the University of Florida College of Engineering, Smith's Aerospace, The Eigenpoint Company, Gainesville Raceway, Videre Design, NavCom Technology, Inc., Preco, Phytex, and Firestorm Signs and Graphics.

REFERENCES

- [1] DARPA Grand Challenge,
<http://www.darpa.mil/grandchallenge>.
- [2] Joint Architecture for Unmanned Systems (JAUS)
Reference Architecture, Version 3.0,
<http://www.jauswg.org/>.