Austin Robot Technology, Inc.

Technical Paper

DARPA Grand Challenge 2005

Jon Brogdon	brogdon@austinrobot.com
Robert Dunlap	rob@austinrobot.com
Patrick Dyson	patrick@austinrobot.com
Art Martin-de-Nicolas	art@austinrobot.com
Jorge Martin-de-Nicolas	jorge@austinrobot.com
Juan Martin-de-Nicolas	juan@austinrobot.com
Don McCauley	dwmccauley@austinrobot.com
Dan Miner	miner@austinrobot.com
Jack O'Quin	joq@austinrobot.com
Stephen Polkowski	stephen@austinrobot.com
Jens Roever	jens@austinrobot.com

www.austinrobot.com

CONTENTS

1 Abstract

Austin Robot Technology was one of 43 teams invited to compete at the National Qualification Event for the 2005 DARPA Grand Challenge. Our vehicle is a modified 1999 Isuzu VehiCROSS and this paper describes the modifications made to the vehicle to allow autonomous operation. Our vehicle uses many of the techniques employed by the 2004 Grand Challenge teams including the use of SICK laser rangefinders. Our most significant departure from the design of most 2004 vehicles is the integration of a vision system using a highly modified Birchfield Stereo Correspondence algorithm for 3D perception.

2 Introduction

Austin Robot Technology is a team of volunteers from Austin, Texas. We designed and built our vehicle in our spare time during a span of 9 months (January – September 2005). Our project was sponsored by the Suzuki-Isuzu dealership of Austin (www.suzukiofaustin.com), who generously provided our vehicle and most of our funding. A project of this magnitude would not have succeeded without unwavering financial support, and we are very fortunate to have found such a generous sponsor. Thank you Suzuki-Isuzu of Austin! Participating in the DARPA Grand Challenge would have been **impossible** without your support!

Austin Robot Technology did not participate in the first Grand Challenge (March 2004) and as such we are a *"second generation"* team. We derived our inspiration to participate from Team Digital Auto Drive, who demonstrated that a small team could be competitive against larger, better funded teams. We have benefited tremendously from those who came before us and generously shared their knowledge and experiences from the first Grand Challenge. Thank you Ivar Schoenmeyr from Team CyberRider for placing a greater emphasis on the word *collaboration* than on the word *competition*.

3 Vehicle Description

This section describes our vehicle, the rationale for our choice, along with the modifications we made to achieve drive-by-wire and to generate power for our onboard systems.

3.1 Vehicle

Our vehicle is a 1999 Isuzu VehiCROSS *"Ironman"* edition. Even though our primary sponsor is the Suzuki-Isuzu dealership of Austin, the principal owner of the dealership did not limit our choice and he offered to buy any vehicle regardless of manufacturer. Our vehicle selection was therefore based on technical merits, the right price, and the right opportunity.

We wanted a 4WD vehicle that would perform well in the desert, had good ground clearance and stability, was manageable in size and had a limited slip differential. The VehiCROSS met all our criteria.

The VehiCROSS was designed in the early 1990's as a *"concept car"*, a high-performance SUV designed for off-road duty. The frame is box-section ladder-type, a truck-like construction that performs very well off-road. Its front double wishbone and rear 4-link suspension are an excellent balance of traditional 4WD articulation and modern chassis design. The front suspension provides remarkable vertical travel and articulation. In one comparison in low-range four-wheeling, the VehiCROSS delivered record performance for an independent front suspension four wheel drive vehicle.

The shock absorbers are made from aerospace-grade aluminum and feature a monotube design with heat-expansion chambers. The design and level of specification were proven in the Paris-Dakar Rally and in the 1999 Australian Safari Rally where the VehiCROSS won its class.

The Borg-Warner Torque-On-Demand transfer case, with 2.48:1 low range and one of the tightest rear limited slips on the market can anticipate where traction will be needed, based on sensor input from the tires, throttle, engine, and transmission. The transmission delivers power to wheels with positive traction and removes power from wheels before they start slipping. This gives us confidence that our vehicle will not spin its wheels if it gets caught on a berm.

The VehiCROSS features heavy underframe members and a guard in front of the transfer case. Isuzu has a reputation for excellent truck engineering and is known to pay attention to underside engineering and brake line placement. Unlike many similar trucks, the VehiCROSS has four-wheel disc brakes for excellent performance and water shedding capabilities.

Stock, the VehiCROSS has 8.5 inches of ground clearance yet still handles with alacrity. At less than 4000 pounds curb weight, it is strong yet compact and has performance and economy advantages over full-size SUVs. It is rated at 15 EPA MPG city yet still achieves a sub-9 second 0-to-60 time.

We are very satisfied with our vehicle selection. The performance characteristics of the VehiCROSS far surpass the capabilities needed by the current state of the art in autonomous ground vehicles.

3.2 Modifications

This section describes the modifications made to our Isuzu VehiCROSS in order to achieve drive-by-wire and in order to power our onboard systems.

3.2.1 Steering Control

Steering is provided via a Quicksilver 34HC-1 I-Grade Brushless Servo. The servo is capable of operation at 12VDC (up to 48VDC). The motor is programmable via an RS232 interface and has an integrated encoder with 16,000 line/rev resolution. A 10:1 gear reduction was used in order to achieve the right torque and speed combination. The motor and gear head are mounted under the steering column as shown in Figure 1 (Concept 2) below:

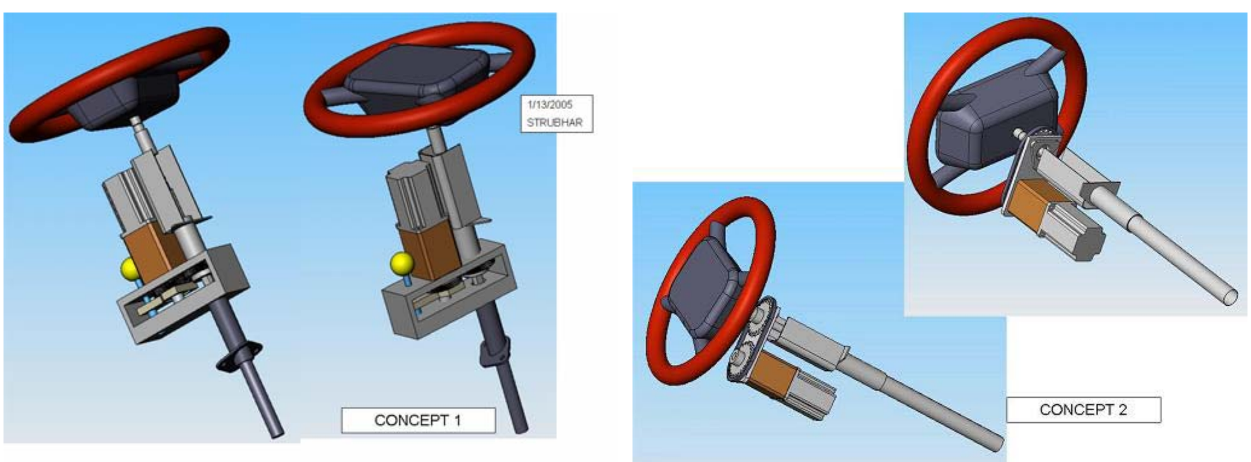

Figure 1: Steering Concepts 1 and 2 (chain in middle or top of steering column)

The ignition switch was removed from the steering column and the chain sprocket was attached in the resulting space. Attaching the chain sprocket on top of the steering column guaranteed that we would not compromise the integrity of the collapsible (safety) steering column. Attaching the sprocket at the middle of the steering column (as shown in Concept 1 above) would have prevented the collapse of the steering column during an accident.

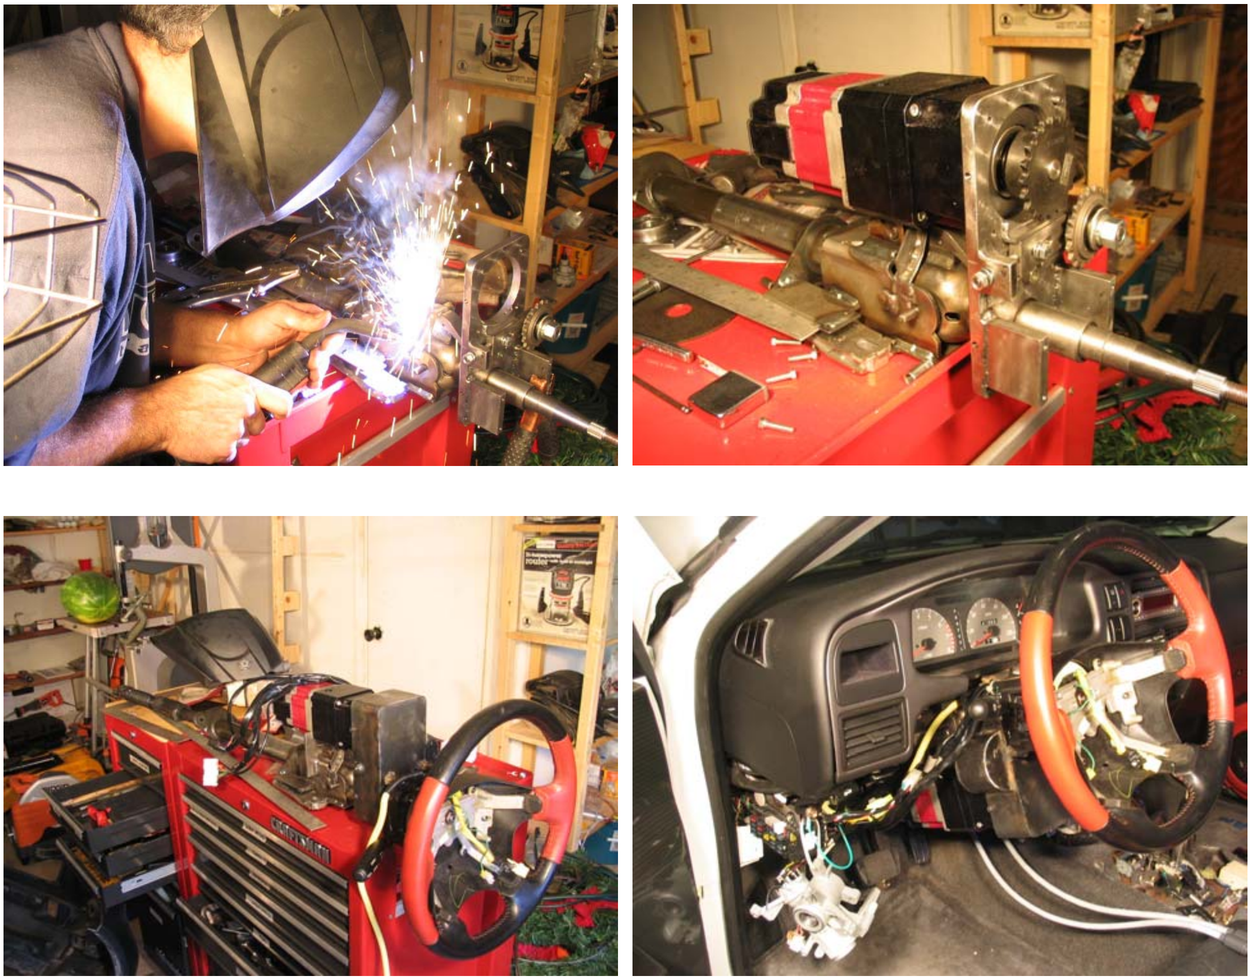

Figure 2: Fabrication of steer-by-wire system

Figure 2 above shows a sequence of photographs from the fabrication of our steering system. Notice the ignition switch temporarily hanging under the instrument cluster on the lower right hand photograph. The ignition switch was later relocated to a custom-made center console.

Our system is capable of providing over 30 foot-lbs of torque and can turn the steering even with the engine off and power steering disabled. The system can turn the wheels lock to lock in close to two seconds with power steering enabled. We have learned that sizing the steering motor is important to prevent possible burnout of the steering actuator. We are satisfied with our choice and with the fact that our system can operate without strain even without the assist of power steering.

3.2.2 Brake Control

Braking is provided via an Ultramotion *Bug* linear actuator with an Animatics SmartMotor. The linear actuator pulls on a cable which is attached to the brake pedal. The entire actuator assembly sits on a sliding cart which is held in place via an electromagnet. Current to the electromagnet can be interrupted via the external emergency stop buttons and via the DARPA-supplied Estop unit. Pressing the Estop button disengages the electromagnet and the entire actuator assembly is pulled back via two 40 lbs constant force springs (80 lbs total). This design guarantees that the Estop mechanism will stop the vehicle in a fail-safe manner even if the vehicle loses power.

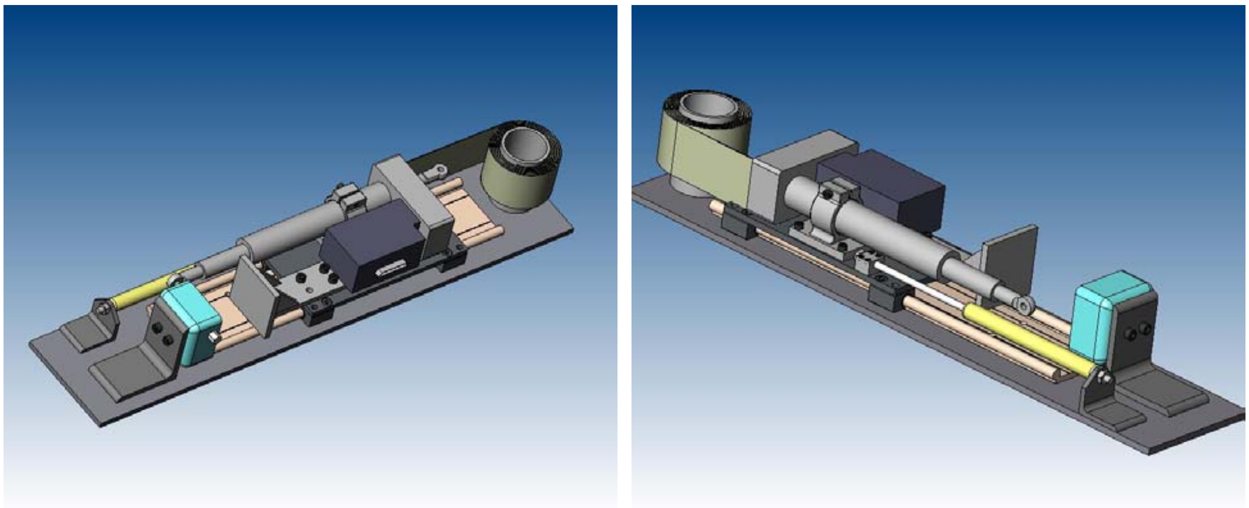

Figure 3: Brake Actuator CAD drawing

Figure 3 above shows the CAD drawings of our brake actuator, and Figure 4 below shows a close-up of the sliding mechanism. The mechanism was bench-tested via a custom test frame also shown Figure 4 below.

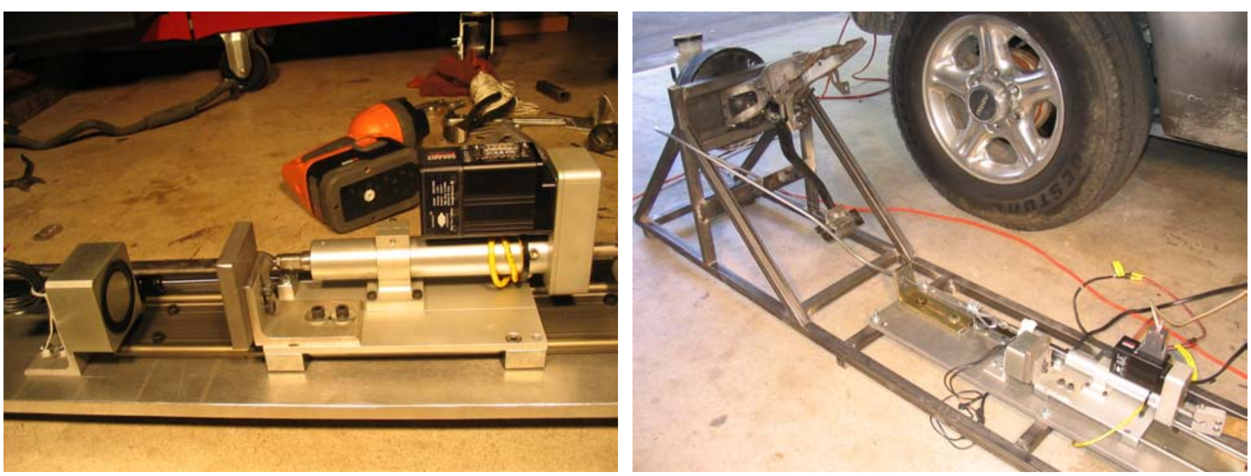

Figure 4: Brake actuator close-up and bench test frame

For our bench tests we installed a spare brake assembly from an Isuzu Trooper (same part as the Isuzu VehiCROSS) to the frame in order to duplicate the resistance from applying the brake. The spare unit was fully functional, including a master cylinder and the brake lines were sealed to provide braking resistance. We tested our unit through several thousand cycles using this setup.

Figure 5: Brake actuator installation

The brake actuator pulls on a Teflon coated steel cable routed from the cargo area of the vehicle through a ¼ inch diameter steel pipe which goes under the vehicle floorboard and through the firewall. Figure 5 above shows the location of the brake actuator behind the driver seat and it shows a close-up of the cable as it exits the firewall and attaches to the brake pedal. We used a cable to pull on the brake pedal from behind in order to allow the vehicle to be operated normally by a human driver.

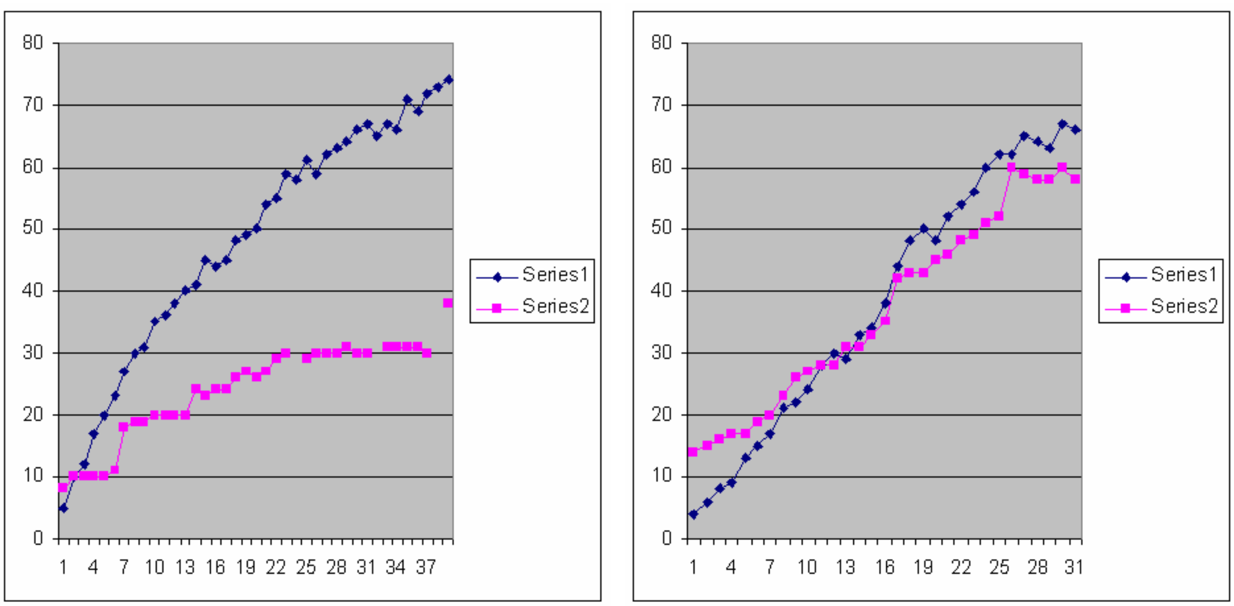

Figure 6: Evaluation of static friction with and without Teflon

The steel cable and the inner wall of the steel pipe were coated with Teflon to minimize losses due to static friction. Figure 6 above shows the force applied at the brake actuator end of the cable (blue) and the resulting force at the brake pedal (fuchsia). The left chart shows the result before coating the cable and pipe with Teflon. Notice a 40 pound loss at the brake due to static friction of the cable. Continued application of force past 60 lbs at the brake actuator did not result in a measurable increase in force at the brake pedal. The chart on the right shows the results after coating with Teflon and shows an almost linear correspondence. Many thanks to TexLoc, Ltd. (www.texloc.com) of Fort Worth, Texas for supplying custom PTFE coated cabling for our brake control system - exceptionally low friction and very good quality.

3.2.3 Throttle control

Throttle control was provided by using the vehicle's cruise control unit. We chose to utilize the standard cruise control actuator (motor) because it is designed to tolerate the harsh environment under the hood and because it is designed to immediately release upon application of the brake. We utilized this same functionality and connected it to our *"Manual Override"* button (panic button) on our custom center console. Pressing this button disengages the steering and throttle and returns control to the human driver (Figure 8 below shows the yellow *"Manual Override"* button mounted to the center console). We replaced the vehicle cruise control electronics that drive the actuator with our own circuit. This circuit accepts commands from the host computer via RS-232 and drives the actuator with a bipolar PWM (pulse width modulated) signal using an H-bridge. Thus the direction and speed of the actuator can be precisely controlled. Figure 7 below shows the proof-of-concept for our cruise control interface before ruggedizing into our vehicle.

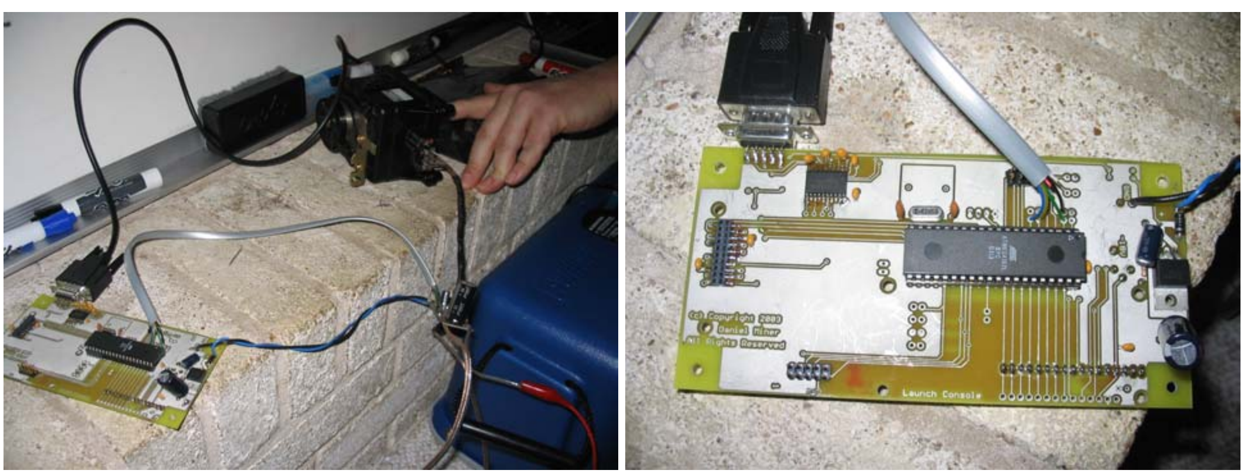

Figure 7: Cruise control interface proof-of-concept

3.2.4 Shift-by-wire system

Our vehicle's shift-by-wire mechanism was provided by Arens Controls Company of Carpentersville, Illinois (www.arenscontrols.com).

Figure 8: Shift-by-wire unit provided by Arens Controls

Arens modified one of their *"Safe-Bus Shift-by-Wire"* systems to fit our vehicle and we are very grateful for their support. Amazingly they were able to fit their system to our vehicle via only the exchange of emails, photographs, and careful measurements. Thank you Arens Controls!

3.2.5 Power

Our vehicle has two power systems, a 12 V system supplied by the stock vehicle alternator and a 24 V system supplied by a 4.5 kw Leece-Neville alternator (24 V 150 Amp), also driven by the vehicle engine. Each of the alternators is attached to a bank of Optima Absorbed Glass Mat (AGM) batteries located in the passenger compartment. AGM batteries were used to prevent emission of hydrogen even under severe overcharge conditions. The output of the batteries then powers the high current systems that take unregulated power, through two standard distribution panels. Regulated power is supplied by buck/boost supplies for equipment that requires regulated, filtered voltages. The system is tuned to supply 2.5 kw at idle to support the two hour pause requirement. It also can run off of the battery bank for extended periods of time if required.

Figure 9: Installation of Leece-Neville alternator

The alternator was mounted inside the engine compartment. In order to provide space for the alternator, we eliminated the vehicle's OEM battery and routed the vehicle's 12 V electrical system to our custom 12 V AGM battery bank inside the passenger compartment. The alternator was mounted on a custom-made bracket attached to the battery tray. The generator is driven by a belt, attached to the main crankshaft pulley. A secondary pulley was welded to the harmonic balancer in order to drive the belt. Routing of the custom belt system required precise fabrication to match the angle of the engine and to route the belt around the vehicle's OEM belts and pulleys and around the main engine fan. The radiator return hose had to be re-routed via a custom steel pipe in order to attain the necessary clearances. The steel pipe was connected to the engine return via a flexible radiator hose to allow for engine rotation and vibration on the engine mounts. The alternator is mounted to the vehicle's chassis therefore consideration had to be made for maintaining proper belt tension as the engine rotates and vibrates during normal operation. Tension in the pulley is maintained via a custom belt tensioner.

4 Autonomous Operation

This section describes the hardware and software used for the autonomous operation of our vehicle.

4.1 Processing

The primary processing system within our autonomous vehicle consists of a pair of computers used for Vision and a third computer used for Navigation.

Each Vision system contains:

- A Tyan K8WE (S2895) dual-socket Opteron mainboard with two nForce Professional bridge chips.
- Two, low-power, 2.0 GHz, dual-core AMD Opteron processors – 4 cores total.
- A PCI-Express based NVIDIA 6600 GT Graphics Processing Unit (GPU).
- Two 1 GB/s Ethernet interfaces; a dual-IEEE-1394a PCI card.
- Two compact flash based disk drives: one 2GB drive for the O/S and application run-time. One 1GB drive for logging.
- 2GB of Registered, ECC-protected Double Data Rate (DDR) DRAM.
- A shock-mounted 160GB Serial-ATA disk drive, used for development.
- A total of 12 CPU, GPU and case fans with filters on the intake fans.
- JASUNY 24VDC 460W power supply.

The Navigation system is similar except for:

- Two, low-power, 2.2 GHz, single-core AMD Opteron processors – 2 cores total.
- The dual-IEEE-1394a PCI card is replaced by a combination IEEE-1394a + USB 2.0 PCI card.
- 4-port RS232 PCI card is added

The Navigation system contains a total of 5 RS232 ports and 11 USB 2.0 ports for sensors and actuators. Many of the USB 2.0 ports are, in turn, connected to USB to RS232 / RS422 / RS485 adapters depending on the sensor or actuator used.

All components selected for use in the computer systems were rated for at least 50,000 hours MTBF and, in many cases, much more. Other than the highly redundant cooling fans, there are no moving parts within the computer systems during normal operation. The symmetric multi-processor computer systems used are highly redundant – with multiple concurrent hardware and software resources available - and provide improved real-time performance. ECC memory provides single-bit error correction.

Using the same, simple basic design for both the Navigation and Vision computer systems reduces complexity significantly. The single-core and dual-core Opteron parts used, for example, share the same power rating of 55 Watts. This means the same cooling solution can be used for all 3 systems. Parts – and even entire systems – can be swapped, if necessary.

The two Vision systems are redundant and operate independently; one Vision system, however, is configured as a spare for the Navigation system, just in case the Navigation system fails before the race; switchover is done manually. The Vision system is capable of running both the Vision and Navigation applications concurrently, though at reduced vision resolution and frame rate.

Each of the 3 primary computer systems is shock-mounted – in a custom made computer rack – using a set of 4 Wire Rope Isolators designed for our team by Enidine Incorporated (www.enidine.com). The design incorporated the computer's mass, center of gravity, and the likely off-road shock and vibration inputs. Thank you Enidine!

4.2 Software Architecture

Figure 10 below shows a diagram with the three main components of our software architecture. The subsystems are labeled *the Commander*, *the Navigator* and *the Pilot*.

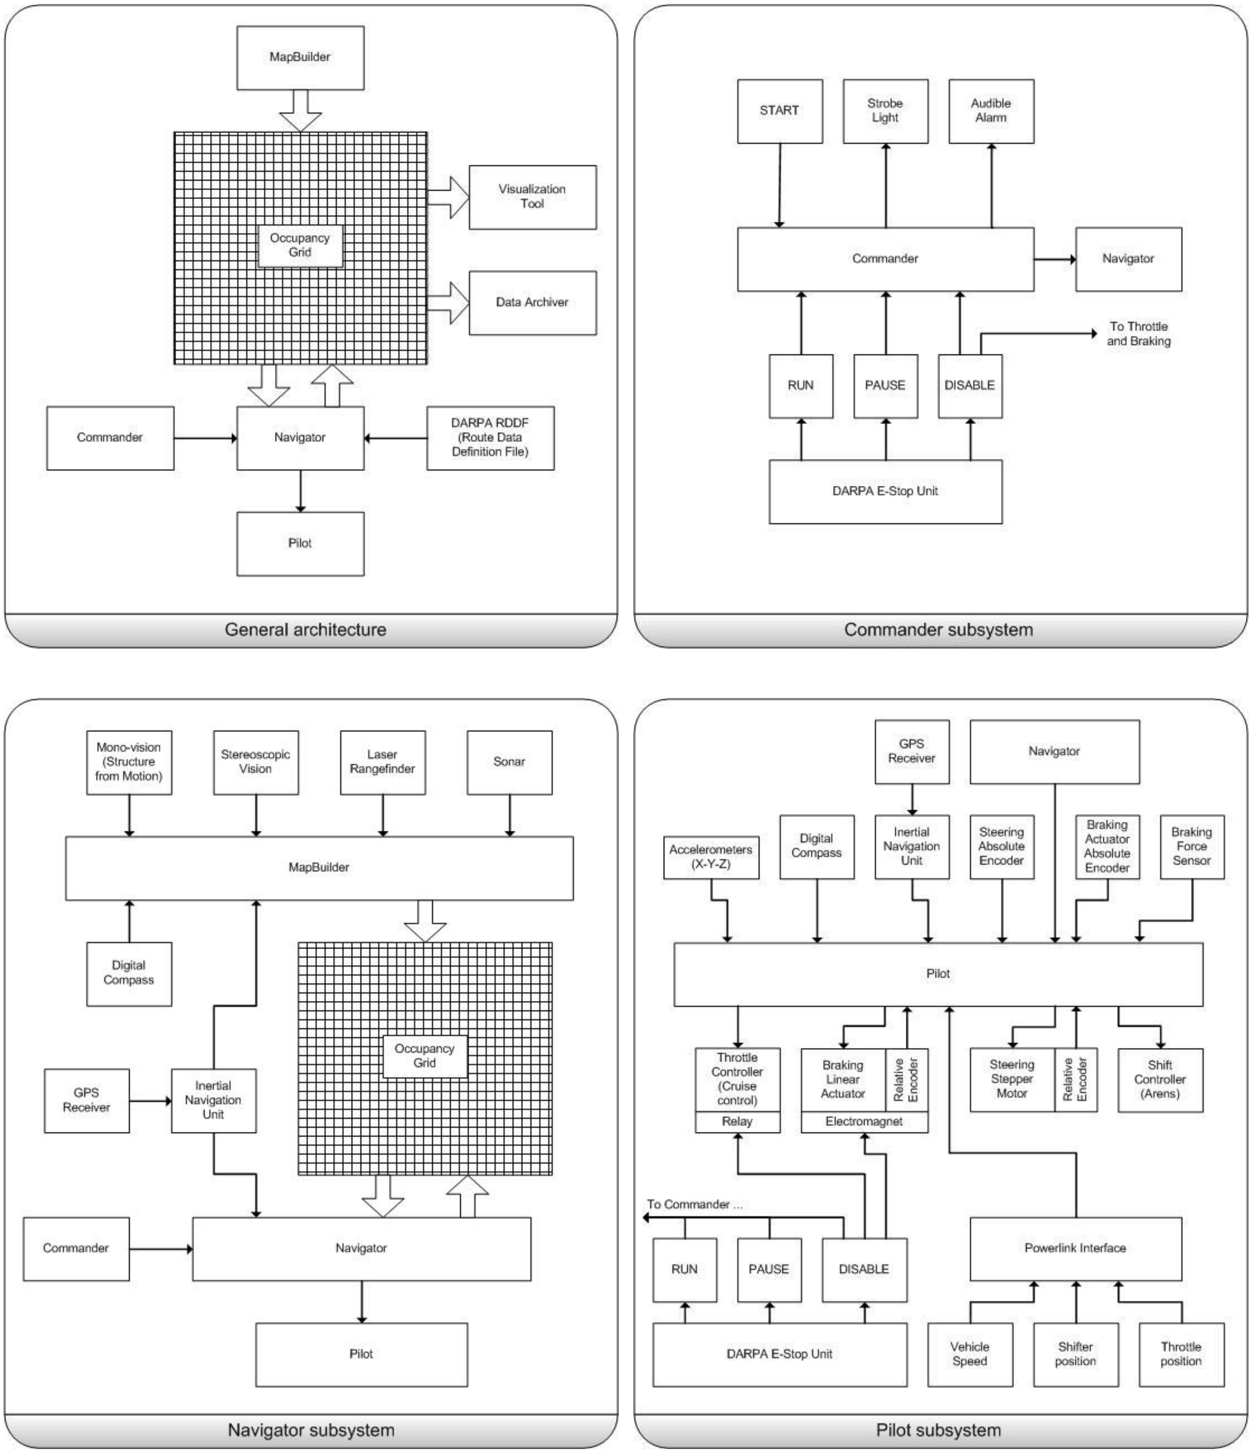

Figure 10: Main components of software architecture

4.3 Localization

This section will explain the GPS system used and inertial navigation systems employed for vehicle localization.

4.3.1 GPS and Inertial Navigation

Our vehicle uses a NavCom SF-2050G receiver (www.navcomtech.com) coupled with a subscription to the StarFire Network Global DGPS Service. The system is capable of decimeter accuracy when in full differential mode.

The GPS is integrated with a Point Research (www.pointresearch.com) Vehicle Navigation Unit (VNU) which integrates the GPS, vehicle odometer, magnetic compass, and accelerometers to provide continuous position data even during GPS outages (such as within a tunnel). The inertial system and the GPS position information are processed with a Kalman filter to give a best-estimate current heading, velocity, and position.

The major drawback of the GPS and inertial system used in our vehicle is the lack of accurate heading information under edge conditions. As is well understood, GPS alone is unable to provide heading information when the vehicle is stationary. Likewise a magnetic compass such as the one provided by our inertial navigation system is unable to provide accurate heading information during periods of sudden acceleration or deceleration (or drastic changes of vehicle pitch or roll). Fiber optic gyroscopes or dual antenna GPS receivers would help in solving this problem but their cost may be prohibitive in certain applications.

Our system solves this problem by combining the output of the GPS and VNU with the output of a Sauer-Danfoss (www.sauer-danfoss.com) Hall Effect rotary position sensor attached to one of the vehicle's steering linkages. The sensor provides a very accurate independent confirmation of the steering angle and is used to generate an independent estimation of position, interpolating when the GPS signal is strong and extrapolating when it is poor. This final position solution is used as the basis for vehicle position.

4.3.2 Map Data

No map data or satellite imaging is used by our software. The RDDF file is parsed and a virtual corridor is built from the waypoints and lateral boundary offsets that the vehicle then navigates within.

4.4 Sensing

This section describes the location and mounting of vehicle sensors and includes a discussion of sensor range and field of view.

4.4.1 Sensor Locations

Our vehicle uses two main sensors for localization: laser rangefinders and stereoscopic vision.

We have two SICK (www.sickusa.com) 291-S14 FAST LMS laser rangefinders currently mounted to our vehicle. One is mounted on the vehicle's front bumper and scans horizontally at 180 degrees in front of the vehicle, 18 inches from the ground. The other rangefinder (set for 90 degree scans) is mounted on the vehicle rooftop at a 6 degree angle and it scans the terrain in front of the vehicle. Both LMS units are set for 80 meter mode.

We have four Sony (www.sony.com) XCD-SX910 cameras forming two stereoscopic pairs using 6mm and 9mm focal length lenses. The cameras are mounted on the rooftop rack with a baseline between 13 and 34 inches, as described later in this document.

4.4.2 Sensing Architecture

The laser and vision system data are fused to build a local occupancy grid by the *Navigator*. This grid also has the corridor boundaries marked as high-confidence obstacles and all navigation must occur within the boundaries. In general, the sensors are merged one at a time into a combined view, then the occupancy grid is populated with that merged data. The occupancy grid has dual coordinate systems - Cartesian for integration with course data and polar for polar sensors and for the local navigation algorithm. This grid is regenerated on demand as the sensor data is received.

4.4.3 Sensing Vehicle State

The main control system is interfaced with the vehicle OEM computer via OBD-II standard protocols. This gives engine RPM, throttle position, odometer readings, and a wealth of other data. This is used to keep the control system within the operational limits of the vehicle. In addition to vehicle operation, the vehicle pose information is read from the navigation sensors and steering position sensor and used to limit the yaw rate of the vehicle.

4.4.4 Waypoint following, path finding, obstacle detection, collision avoidance

Waypoints are processed into a series of goal points, often sub-waypoints are generated to maintain safe theoretical speeds around anticipated course changes. The goals and speed limits are handed to the *Pilot* one at a time and a heading is calculated between the current vehicle position and the goal position.

The *Navigator* uses the occupancy grid and the vehicle dynamics to determine possible routes, then evaluates the routes using a standard costing algorithm. The chosen route is then described as a speed and rate of turn, and this request is handed to the *Pilot*. This solution is updated every sensor cycle.

The *Pilot* calculates the steering angle and hands that request to the steering driver which controls the steering servo. The speed is fed into a PID control loop that uses braking and acceleration to achieve the requested speed. Special cases are made for emergency stops which bypass the smoothing algorithms and bring the vehicle to a stop as quickly as possible.

The vehicle is modeled with a series of constraint tables that describe the operational limits for acceleration and turning, indexed by speed. These in turn are processed into attainable and unattainable regions and are used both by the route planner and the control driver to keep within the vehicle performance envelope.

4.5 Vehicle Control

This section describes methods employed for vehicle control, including common autonomous operation contingencies (such as out-of-boundary conditions), autonomous maneuvers (such as starting on a hill) and manual vehicle control.

4.5.1 Autonomous Operation Contingencies

Contingency planning is reduced to several cases. For out of boundary conditions, the shortest path back to the route is planned and the normal navigation is employed to reacquire the route corridor. The speed is reduced dramatically to approximately walking speed and in the event of an unnavigable solution the vehicle is stopped. This is not a normal contingency as the vehicle cannot consider routes that take it outside the boundaries, but is possible with an emergency maneuver near a boundary or with a major inertial or GPS location change event.

If the vehicle is stuck, and no obstacles are detected it will attempt to extricate itself by increasing throttle (limited by a max RPM value). Obstacles are handled normally, if no free path is available the vehicle will stop and wait for the sensors to clear. If no path is found, a limited distance reversing mechanism may be employed to increase possible paths after a waiting period. If neither succeed, the vehicle will re-approach the obstacle and stop.

For missed waypoints, the vehicle will attempt to reverse course to re-attain the waypoint. If it is unattainable, the vehicle will stop. This is not a normal course as the waypoint is considered achieved if the vehicle is within the lateral boundary offset of the waypoint.

4.5.2 Autonomous maneuvers, starting on a hill, sharp turns, etc.

The vehicle uses dynamic speed limits to set the safe speed and should begin braking well before the safe stopping distance. Slow, smooth approaches to obstacles are the rule. Starting on a hill is not an issue as the vehicle transmission is engaged before the brakes are released. Sharp turns are approached the same way, the vehicle brakes well before the turn, centers on the course, and attempts to keep to the centerline of the turn. Sharp turns with obstacles can be a challenge so the safe turning speed is less than the safe stopping distance for a possible obstacle. The vehicle has a small wheelbase and a tight turning radius which helps minimize course boundary violations.

4.5.3 Manual vehicle control

The vehicle can be driven normally when not in autonomous mode. Other than limited visibility through the rear window, all systems function as expected (steering, brake, signals, accelerator, etc.). Shifting is done via a pushbutton console connected to the Arens Controls electronic shift mechanism.

5 Vision System

This section describes the hardware and the algorithms we used for our vision system.

5.1 Introduction

The Vision computer systems run a custom 3D Stereo Correspondence application based on the Open Computer Vision library [1] originally developed by Intel Corporation. A brief summary: Stereo correspondence finds matching edges in the left and right image pairs and calculates the distance to the corresponding obstacles. The resulting 2D disparity (aka distance) map – where image intensity is inversely proportional to distance – is then converted to a pseudo laser scanner data stream (angle, distance-to-obstacle) and sent to the Navigation system.

5.2 Birchfield Stereo Correspondence algorithm

The stereo correspondence algorithm used is a highly modified version of the original "Depth Discontinuities by Pixel-to-Pixel Stereo" algorithm developed by Stan Birchfield and Carlo Tomasi [2]. The original algorithm is well described in the literature; this paper will describe features specific to our implementation.

5.3 Cameras

The cameras used in our stereo vision system consist of two pairs of high-resolution Sony XCD-SX910 cameras with 6mm and 9mm focal length lenses, respectively. One reason we selected IEEE-1394 cameras is that almost all IEEE-1394 cameras support the "IIDC 1394-based Digital Camera (DCAM) specification" which helps ensure software written for one DCAM-compliant camera will run on other DCAM-compliant cameras. We used considerably less expensive Unibrain Fire-I DCAM-compliant cameras for much of our software development.

5.4 Mounting and alignment

The Sony cameras are mounted in a rigid frame on the roof of the vehicle, pointing downward at a 4 degree angle. Each pair of cameras is carefully mechanically aligned to eliminate variations in roll and pitch between cameras. A large 4' x 8' chessboard calibration target is used to perform the mechanical alignment as shown in Figure 11 below:

Figure 11: Camera mount and chessboard calibration target

5.5 Camera baseline

Each pair of cameras uses a baseline – the separation between cameras – of between 13 and 34 inches in our application. This is a significantly longer baseline than used in most commercial stereo cameras. The longer baseline, however, allows us to better resolve obstacles at greater distances. The tradeoff is that our stereo cameras are noticeably near-sighted; they do not detect small objects near the cameras ('objects outside of the horopter'). However, since it is approximately 7 feet from the cameras to our vehicle's front bumper, this near-sightedness is not a concern; we are only concerned with detecting obstacles in front of the vehicle. Typically the 9mm cameras – with a narrower field of view – are mounted closer together than the 6mm cameras to minimize this near-sightedness effect.

5.6 Image binning

To help reduce our computational workload, while preserving as much camera resolution as possible, we use the 2x2 binning mode of the Sony cameras. Binning combines pairs of pixels horizontally and vertically to reduce our resolution from 1280(H)x960(V) to 640(H)x480(V). Binning also significantly reduces image noise. We also used the auto-shutter and auto-gain features of the Sony cameras.

5.7 Camera calibration

The large camera calibration target described earlier (under *Mounting and alignment*) is also used during camera calibration. Camera calibration, as described in the openCV documentation, calculates the intrinsic parameters (focal length, image center, pixel size, radial distortion) for each individual camera as well as the extrinsic parameters (rotation matrix, translation vector) for the stereo camera pair. We have made a number of minor improvements to simplify the camera calibration step, which is normally only required if the mechanical alignment or the focal length of the stereo camera has changed.

5.8 OpenCV code improvements

We multi-threaded essentially the entire image acquisition and image processing pipeline to distribute the workload over our four Opteron cores. Two threads – one per camera – are used to acquire images from the cameras. The initial image processing steps – 4 threads - include a light Gaussian filter – to smooth or blur the left and right images – followed by undistort and rectification. Undistort applies image transforms – based on the individual camera intrinsic parameters described earlier – to eliminate, for example, barrel distortion. Rectification applies image transforms, based on the stereo camera extrinsic parameters described earlier, to rotate and project the images to align the left and right images, scan-line by scan-line.

5.9 Undistort and rectification transforms

We merged the undistort and rectification transforms into a single image processing step to improve performance. The end result of the undistort and rectify step is a pair of stereo images where objects are aligned scan-line-by scan-line. Because we are using stereo cameras, the left and right images of the same object are offset by some number of pixels. The number of pixels – the disparity - is inversely proportional to distance. We multithreaded the computationally intensive disparity calculation (4 threads).

5.10 Edge detection

Birchfield stereo correspondence uses edges – mostly vertical edges – to delineate objects. We replaced the original Birchfield edge detectors with a Sobel edge detector, which significantly reduced noise in the intermediate edges map.

5.11 Dynamic ground-plane removal

We also added dynamic ground-plane removal and overhead obstacle removal based on the vehicle's pitch relative to the terrain. The vehicle's pitch is determined, dynamically, by comparing the actual disparity of a number of image points in front of the vehicle with the expected disparity – based on known camera height and camera angle - for the corresponding scan lines.

5.12 Pseudo-laser data

Finally, the 2D disparity map is projected to a plane and converted to a pseudo laser-scanner data packet which matches the format used internally by the rest of our navigation software. Packets are sent from the Vision systems to the Navigation systems using socket connections over the vehicle's gigabit Ethernet network.

5.13 Results

Figure 12 and Figure 13 below depict left and right sample images and the resulting disparity map both before and after dynamic ground plane removal.

Figure 12: Left and Right Sample Images

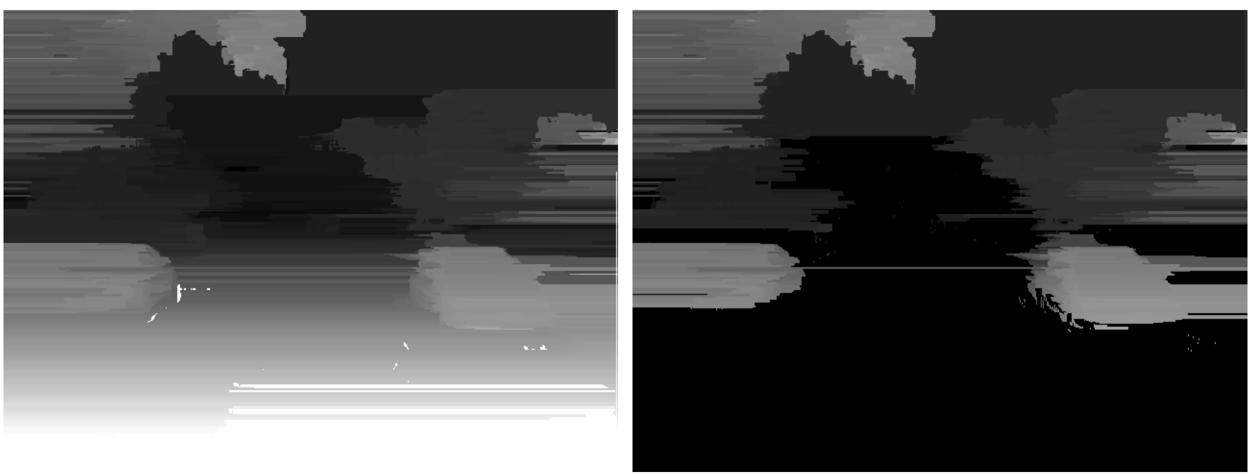

Figure 13: Disparity Map before and after Dynamic Ground Plane Removal

6 Vehicle tests

This section has an overview of the unit and systems testing of our vehicle.

6.1 Unit testing

All physical vehicle systems were unit tested before integration, then again to characterize the behaviors under stress. The vehicle was initially raised on jack stands and a pre-recorded waypoint-following loop was executed to test the steering response. Custom harnesses were made to test the braking system reliability, the throttle control including emergency stop response, and the laser systems. The brake and throttle systems underwent stress tests with literally thousands of repetitions without failure. When issues were found the design was changed and retested until the systems passed with a 2x safety factor. For example our brake actuator was changed from 24 V to 48 V to increase the speed and torque characteristics in order to enhance our brake responsiveness. The software was tested in simulation before and during the vehicle development. A simulation environment enabled parallel development of the software and the vehicle hardware and provides an excellent way to replicate issues that are seen in road tests, solve the issues, and regression test the system.

6.2 Systems testing

The vehicle as gone through four phases of incremental development and test, starting with simple waypoint following, then integrating obstacle avoidance, then increasing the difficulty with steep terrain, sharp corners, and high temperatures as the summer has progressed. Most of the specific hazards of the NQE have been tested though the tunnel tests were only simulated by dropping GPS in a tight corridor. Our vehicle has attained straight-line speeds up to 45 MPH, but we have done only limited high speed testing, concentrating on solid performance under 30 MPH.

6.3 Road testing

Our vehicle has undergone literally hundreds of autonomous runs, including many off-road runs in Texas ranches near Burnet, Blanco and Dead Man's Hole. We have not been able to find and take advantage of desert terrain like the GC course but will be testing the vehicle at the U.S. Army Yuma Proving Ground in Yuma, Arizona prior to the Grand Challenge. Thank you U.S. Army!

6.4 Visualization tool

Each test run produces several log files which are compressed and saved automatically to a log directory. Using our custom visualization tool, we can replay the log from any run as a debugging aid. Our visualization tool is based on OpenGL which provides the capability of using an image as the texture for any rendered surface. We have downloaded Google Earth images corresponding to our various test locations and our visualization tool searches our database for the corresponding image given the vehicle's latitude and longitude for any given run. Figure 14 below shows the base image used for vehicle runs performed at the parking lot of the former Robert Mueller Municipal Airport in Austin, Texas where we hosted our 2nd DARPA site visit.

Figure 14: Base image for the former Robert Mueller Municipal Airport parking lot

The left image above is part of a web page we created for DARPA representatives. It shows the approximate course layout, the area reserved for spectators, and the visitor parking. The right image shows a screen shot from our visualization tool using the same Google Earth photograph to texture the ground. The screen shot corresponds to run number three during our site visit and it shows the vehicle approaching the second obstacle after negotiating the first turn. Figure 15 below shows a front view and a back view of the same scene... For the fun of it, we added the lampposts using Open GL.

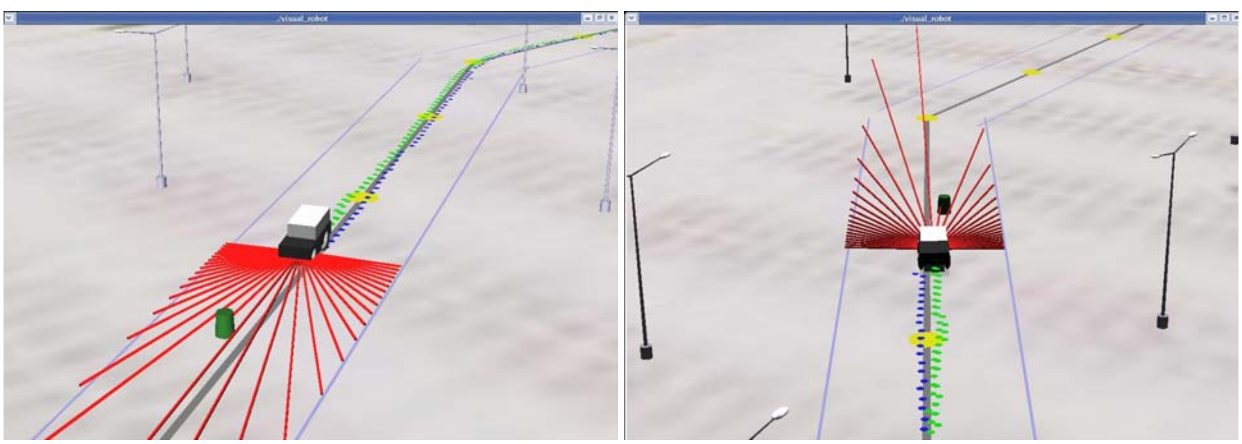

Figure 15: Visualization tool output, front view and back view, DSV run number three

7 References

[1] "Open Computer Vision Library" http://sourceforge.net/projects/opencvlibrary/

[2] "Depth Discontinuities by Pixel-to-Pixel Stereo" Stan Birchfield and Carlo Tomasi http://vision.stanford.edu/~birch/p2p/