TEAM BANZAI TECHNICAL PAPER
DARPA GRAND CHALLENGE 2005

JOHN CHOI. PH.D.
CHIEF SCIENTIST
BANZAI RESEARCH INSTITUTE
JCHOI@GOBANZAI.COM

Disclaimer

The information contained in this paper does not represent the official policies, either expressed or implied, of the Defense Advanced Research Projects Agency (DARPA) or the Department of Defense. DARPA does not guarantee the accuracy or reliability of the information in this paper.

Table of Contents

Abstract

Team Banzai is a privately funded effort by Banzai Research Institute to develop a fully autonomous vehicle capable of traversing long distances over extreme terrain. Our vehicle is a luxury SUV with off-road racing capabilities. Using object-oriented technologies, we developed custom software for three Mac mini computers to process data from various sensors and drive the vehicle-control actuators. In addition to operating autonomously, our vehicle has a 'cooperative driving' mode that allows human and computer to share in controlling the vehicle as co-drivers. We feel this represents a significant application for consumer use.

Vehicle Description

Model Selection

Our race vehicle is a minimally modified 2004 Volkswagen Touareg with a V6 engine, electronic air suspension, all time four-wheel drive, and electronic center differential locking. In our initial assessments we determined that the best way to ensure reliable performance of sensitive electronic equipment and computers was to provide a stable, climate controlled environment. Therefore, rather than building on a highly custom, open air chassis platform (e.g. sand dune buggie, modified golf cart, or ATV), we choose to start with a commercial luxury SUV. We also wanted to find a vehicle that already had highly intelligent vehicle control systems such as automatic traction control, anti-lock braking, and electronic stability control. We wanted to minimize our efforts in controlling the car under normal driving conditions.

The VW Touareg satisfies these requirements spectacularly, and additionally delivers excellent race-proven off-road capabilities. Modified versions of the Touareg are routinely raced by the VW factory team in premiere desert rallies throughout the world. Furthermore, the vehicle is controlled by numerous advanced sensors and processors communicating through an in-vehicle digital CAN bus network. We were able to easily tap into this network to exploit the inertial and mechanical telemetry data for our autonomous control system.

Autonomous Control

The vehicle was modified for autonomous driving in several areas. In place of the stock plastic covers, the underbody was reinforced with a custom fabricated metal skid plate. Additionally, some of the plastic body flaring was cut away to avoid inadvertent breakage and to allow access for mounting sensors. A custom equipment rack was fabricated for shock mounting the computers in the cargo area and mounting of sensors and antennas on the roof.

The steering control system consists of opposing wire steel cables connected to a linear actuator. Similarly, the brake and accelerator actuation is accomplished with an opposing pair of wire cables feeding into a single linear actuator controlling both pedals. We are in the process of outfitting a similar system for transmission shifting. The design philosophy behind the wire drive system was to develop an unobtrusive solution that would retain the native driving characteristics while operating in multiple modes: human driver, remote control, fully autonomous, and cooperative driving.


2004 VW Touareg


Linear Drive Actuators

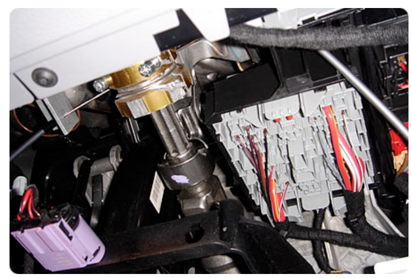
Wire cable steer system


Equipment Rack


Co-operative Drive System


Compact Form Factor

Information Processing

Systems Selection and Design

The foundation of our computing system is a trio of Apple Mac mini computers running custom developed software. We choose the Mac mini for its compact form factor, low energy consumption, and high computing power. Having identical units allowed us to easily swap out and exchange individual units during the development phase. The compact form factor facilitated placement and secure mounting, and has proven reliable in a mobile environment. The computers are connected by ethernet and communicate with each other through a custom developed UDP synchronizing protocol.

We decided to use multiple computers so that we could distribute the computational load onto specialized pre-processing units. The 'Navigator' computer is primarily responsible for processing the serial stream from GPS and LIDAR sensors and transmitting summary information to the other units. The 'Tactician' computer is responsible for determining the best course of travel using the sensory information and the pre-programmed route map. Once a course of action is determined, it is transmitted to the 'Driver' computer that is responsible for managing the vehicle control actuators and monitoring the internal dynamics.

We selected a Serial to Ethernet server device to minimize the additional overhead involved in listening on dedicated serial ports. By converting the serial data stream into ethernet packets, we were able to flexibly route of the information to any of the computers. This allowed us to implement a robust fail-over system that can detect when a primary computer unit goes off-line. In such cases, the serial information can be automatically re-routed in real-time to secondary units.

Differentiating Innovations

Throughout the development process we employed an extension of the iterative development methodology which emphasizes the importance of building on a scalable architecture. We call this the "concentric development process" whereby each iteration is not only rapid and incremental, it is strongly gated upon proving that the previous iteration, i.e. inner-core, is solidly capable of supporting the successive layers. Our software modeling process is likewise an extension of object-oriented design called "reality-based modeling" where object classes are tightly modeled after real-world analogues.

Architecture Diagram

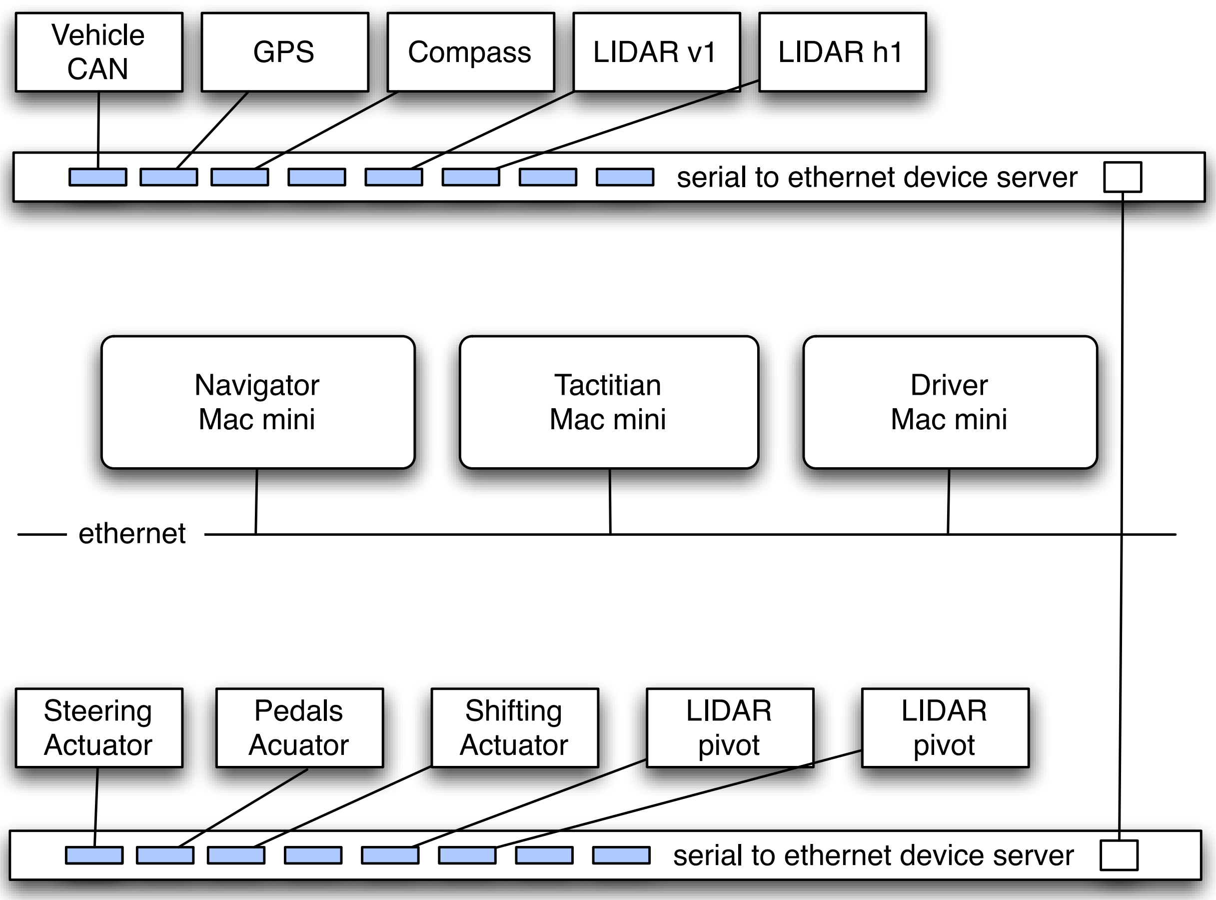

Sensory Subsystem

Inertial and GPS Localization

Primary positioning information is provided by a commercial GPS unit, Trimble AG 132. This unit is capable of receiving public DGPS signals and commercial subscription services such as Global Star. The system can reliably provide sub-meter resolution at a 10 Hz update rate. Secondary positioning information is provided by calculating the change in position based on vehicle inertial measurements (i.e. "dead reckoning"). For inertial measurements we rely on the in-vehicle sensors reporting independent wheel rotational rates and steering angle. This information, when combined with reliable time stamps, provides reliable updates when the car has not moved significant distances. To compensate for the additive error produced by dead-reckoning calculations, we re-sync with GPS readings whenever that signal is deemed reliable.

When the vehicle is traveling at higher speeds, derivative information such as acceleration, heading, and drift can be reliably calculated from GPS readings. At slower speeds, GPS fluctuations can often yield spurious measurements, and in these cases, calculations from inertial measurements are compared to projections based on previous time point using custom developed filters to reject improbable data.

Additionally, static posture information is provided by a digital magnetometer that utilizes a dual-axis linear tilt sensor to provide pitch, roll, yaw data, as well as magnetic heading information.

LIDAR and Sonar Obstacle Detection

The vehicle uses a course map that provides the turn parameter and general direction of travel. Our initial intention was to combine this information with maps from the USGS 1-meter topographical series and digital road information from the US Census Bureau's

TIGER project and USGS DEM 3 meter series. All this information was to be pre-computed and pre-processed to provide the vehicle with a composite value map containing desirable routes of travel (e.g. roads) and areas to avoid (e.g. ravines, canyons). However, after preliminary investigation, we determined that the resolution of publicly available maps was insufficient to provide meaningful additional information. We have chosen instead to rely on the real-time dynamic assessment from our vehicle sensors as the primary means of route determination.

Primary sensing is provided by four LIDAR scanning units (SICK LMS-291). Two are mounted horizontally on the roof angling forward at 7 degrees and 9 degrees respectively. Two are vertically mounted in front of vehicle one over each wheel. We chose this arrangement to provide the rapidly accessible information about terrain quality and presence of obstacles. Rather than attempting to construct a full 3-D model based on 2-D scan information, we use information from each laser to quickly assess go / no-go areas.

The vertically mounted units scan a 100 degree arc at 72 Hz. Even in a fast moving car bouncing up and down, each scan can reliably determine the slope of the terrain immediately ahead of each forward wheel and distinguish a impassible step or ditch from a steep, but drivable, hill. As these vertical scanners alternatively sweep slowly inward at 1 Hz, they can detect obstacles in front of the vehicle.

The two horizontally mounted sensors are responsible for determining the approximate center of road travel. They do this by examining the signal variability and lateral slope (i.e. flatness). The double offset placement provides redundancy and partial immunity from vehicle pitch.

In addition we anticipate the mounting of multiple ultra-sonic range finders on the front and rear bumpers to provide short- range obstacle detection during backing up and close-quarter maneuvering.

The LIDAR units are housed in a dust and moisture-proof NMEA-grade casing. We mounted them onto rigid brackets bolted directly onto the vehicle chassis frame or roof cargo rails using aluminum angles and tubes. Vibration mitigation is accommodated by a layer of shock-absorbent foam padding sandwiched between the unit and bracket. This arrangement, combined with the vehicle's main air-suspension system, appears to be sufficient to avoid damaging shocks and has provided useful data output in several hundreds of testing miles. Additionally, the lower LIDAR units are mounted close inside car body to prevent damage from accidental collisions.

Forward Scanning LIDAR Array

SICK LMS-291

Aositilt EZCompass

Vehicle Control

Lane Management and Maneuver Algorithms

In keeping with our reality-based modeling philosophy, we processed the sensory to extract the type of information a human being would normally be interested in. Instead of trying to develop a full 3-D model based on partial and noisy data, we chose to focus on immediately interpretable metrics to extract salient information. For example, because the vehicle is capable of surmounting a 32-degree slope, a ground slope of greater than 32 degrees and higher than the ground clearance of the vehicle denotes an impassable obstacle. This slope and height information is quickly determined by the vertically mounted front sweeping LIDAR units.

As mentioned before, on-board vehicle operations and dynamics information, such as engine speed and gear selection, wheel rotations, wheel angle, etc. are transmitted via the in-vehicle CAN bus network. By connecting into this network, we were able to access this information and monitor the condition of the vehicle. This information is processed by the 'Navigator' computer to determine inertia-based positioning (i.e. 'dead reckoning') and by the 'Driver' computer to provide feedback control of the accelerator and brake pedals.

The 'Tactician' computer is responsible for using the sensed route information to direct the vehicle via a series of drive maneuver commands. This is model is based on the real-world rally racing practice of pre-running a course to develop a course guide. Analogously, during the initial loading of the route file, the route map is analyzed and pre-calculated into a series of maneuvers such as keep steady, arc at an angle, brake, etc. Based on the vehicle position-velocity and the route map, we can readily determine the target steering angle. The target speed is calculated by cascading through a series of filters. First we determine the maximum or "safe" travel speed based on turning requirements. Then we adjust this speed for terrain conditions and for anticipated future speed such as entering a lower speed limited segment. Lastly, we compare this speed with the route speed limit to truncate if necessary. Using this maneuver-based driver command model, we can interject in between pre-calculated maneuvers any number of new maneuvers as conditions change and obstacles are detected.

Contingency Strategies

For safety and rule compliance we have programmed our vehicle to stay between lateral boundaries at all times. If the vehicle strays outside the boundary, it is programmed to reduce its speed and return at the earliest opportunity. If the vehicle falls completely outside the course and there is no direct route to safely return, the vehicle is programmed to stop and self terminate.

A primary constraint of the driving algorithm is to stay within the lateral boundaries and stay in the direction of travel as determined by the course file. Because of this constraint, it is anticipated that waypoints cannot be "missed". Once the vehicle has passed the perpendicular extension of the waypoint, it is designated as acquired and obtained regardless of how far or close it is to the waypoint. Additionally, rather than moving from waypoint to waypoint we have decided to focus the vehicle to keep moving in the direction of travel.

Detected obstacles are avoided by switching lanes of travel if possible or by stopping if necessary. In most cases, it is anticipated that obstacles will be detected early and avoid in fluid travel. If the vehicle has to stop, either because the obstacle suddenly appeared or the vehicle failed an avoidance maneuver, the vehicle will assess the situation determine if there is a theoretical possibility to move around the obstacle. If such a possibility exists the vehicle will attempt to reverse sufficiently to acquire the available passing lane. If the vehicle becomes stuck, i.e. it cannot move forward regardless of maximum application of pedal, the vehicle places a virtual obstacle at that location and will attempt to reverse out and apply the same avoidance logic as before.

Basic and advanced vehicle control is greatly simplified because of our selection of the Touareg luxury SUV. Its automatic transmission, hill-lock, and various electronic vehicle stability and control features allow us to apply a simple point wheel and press pedal algorithm. By using our reality-based driving maneuvers model, we slow down before a turn, adjust the steering angle based on speed of travel and available lane width, and apply different correction factors to translate desired course change into steering wheel angle.

Route Finding

Intelligent Traction Control

Electronic Center-locking Differential

Non-Autonomous Operation

Our selection of the wire cable actuation system was significant influenced by our desire to retain the original appearance and function of the vehicle under normal driving conditions. Doing so allowed us to develop a multi-model vehicle that could easily transition between several control modes.

Even with the car fully connected, the car can still be driven in 'normal mode' because all of the actuators have a zero holding current setting which allows for free spin operation. In 'remote control mode', each actuator can be selected operated by a human using a computer interface. This mode is useful for testing and debugging, and it facilitates the evaluation of modification and alternative algorithms or strategies. In 'fully autonomous mode', all vehicle control is handled by the on-board computer modules. This mode is designed to mirror the race condition and faithfully adhere to the predetermined route parameters with necessary adjustments for obstacles along the way.

By retaining the normal characteristics of the vehicle, we have been able to combine the all of the control modalities into a special mode called the 'cooperative driving mode'. This mode will allow for an interactive driving experience where human and computer can operate as "co-drivers".

Conclusions

System Tests

Using our concentric development model, we have tested our vehicle and our systems in small stages and by building on previous stages. Initial stages involved testing the computer architecture and rudimentary control of the vehicle on level vacant lots and abandoned roads. Second phase tests involved preparing the vehicle for off-road travel by shock mounting the electronics equipment and adding underbody protection. The completed off-road system was then stress tested in actual desert environments by driving the vehicle many hundreds of miles around the trails in the Ocotillo Wells California Desert Vehicle Park.

The third phase of testing involves close simulation of the DGC race course in the desert by using an example route file and ensuring full compliance with all DGC rules. There the vehicle has been tested against all standard required maneuvers such has waypoint following, obstacle avoidance and stopped vehicle passing.

Findings

As we approach the DGC after 9 months of development we have concluded that many of our initial design directions proved correct. For example the selection of the luxury SUV with extreme off-road capable suspension has yielded an extremely comfortable and stable development environment for both electronic equipment and human developers. We have suffered no computer failures or detected any vibration related wear. In our desert testing during this summer, we have routinely operated the vehicle over hundreds of rugged terrain in temperatures in excess of 115 F.

The selection of LIDAR as our primary mode of terrain sensing has proven to be both versatile and practical. Our strategy in simplifying the LIDAR data interpretation to more directly yield obstacle detection and best-terrain road-centering information appears to be promising. However, although we can determine crudely which portion of the terrain ahead seems most like a "road", we do not have the route-finding ability to determine in which "road" or "trail" will lead us most efficiently to our next route point. This is a difficult problem that remains un-addressed in our current testing and design.

A novel enhancement that we are developing here at the Banzai Research Institute is a special control mode called the 'cooperative driving mode'. This mode will allow for an interactive driving experience where the human operator can issue voice commands to the computer such as "turn right ahead" or "slow down here". The computer will then execute the command or respond appropriately with information such as "lane not available for turning" or "cannot slow down here because of approaching traffic". While the 'fully autonomous mode' will be very useful for industrial and military operations, we feel that the 'cooperative driving mode' is the most significant application consumer use.