DARPA Grand Challenge 2005

Technical Paper

PV Road Warriors

Team: A122

Principal Author: Dr. Arthur Schwartz, artsdarpa@cox.net

Abstract

Autonomous operation of the 2005 PV High School Road Warrior's Honda MDX is carried out by a unique distributed processing architecture based on a series of six single board Pentium D processors. The navigation suite features a new dead reckoning sensor that measures vehicle displacement every two msec to a precision of one mm. Combined with a DGPS receiver and a military grade INS, the car's state vector is determined to a high degree of precision. Environmental sensors consist of two SICK Lidars and a pair of Redlake cameras operating as a stereo pair all mounted on an inertial platform.

Introduction

The Road Warriors of Palos Verdes High School are pleased to return to the 2005 Grand Challenge. Most of the team has returned this year with a strengthened set of adult mentors and participants from individuals formerly members of Team Designated Driver. A strengthened financial structure and business management organization has allowed the team to progress in both technology and sponsorship. As a result, the team has been able to develop a far more aggressive technological approach than in the 2004 Grand Challenge. The project is divided into three groups namely business, software, and hardware. Many members have extensive computing experience, and include experienced C/C++ programmers, the language chosen for the software development. We have had additional help from Industry representatives most notably Lambda Tech International who are vision systems experts.

The Navigation sensor suite adds to the usual INS and DGPS configuration a new sensor called the ground mouse. This sensor whose sampling time is 2 msec provides displacement data both longitudinally and transversally at each measurement point. A prototype was instrumental in the team's performance at the site visit last May. With the full design, sub mm resolution may be obtainable. Thus, if we know our initial position, with a complete failure in the DGPS receiver, the car would still be able to find its way from waypoint to waypoint.

The environmental sensor suite consists of two CMOS based cameras operating in conjunction with two laser range finders providing for both the path and object detection functions. The cameras provide long-range detection in a narrow path whose direction can be commanded. One of the lasers provide mid range object detection and the second provides rut and furrow detection. The Optical sensors are mounted on an isolation platform that also houses the ground mouse sensor. The optical sensors have a three-degree of freedom gimbal and each of the lasers has a two-degree of freedom gimbal.

Vehicle path control is implemented by a three-tier priority structure, our three laws of robotics. When an object of large enough size is detected the route planning software changes the route so that the highest priority is to avoid the object. At the second level, if there is a well-defined path as detected by the vision system that parallels and has an area collocated with the planned corridor, the path is followed. When neither objects nor well-defined paths are present, the preset corridor will define the planned route. The corridor is calculated in real time from the RDDF data to keep a planned route of approximately 1 km distance ahead of the vehicle.

1. Vehicle

1.1 Vehicle Overview

The PV Road Warriors vehicle is a 2005 Acura MDX SUV. The vehicle rides on 4 tires made with Kevlar side walls for improved off road endurance and a standard gasoline powered internal combustion engine. Automation of the vehicle is accomplished using servo motors designed and installed by EMC but with our own servo amplifier system replacing their electronics. The gearshift, engine cutoff, remote start and other vehicle controls all use actuators from EMC as well. The system updates the vehicle state vector every 100 msec, which is adequate for vehicle control loops.

PV Road Warriors "Doom Buggy"

Length: 187.7 "
Width: 76.3"
Height: 71.3"
Fuel Capacity: 30.4 gallons
Weight: 4500 lbs. Gross
Characteristics: 4WD, Automatic Transmission

Figure 1- Doom Buggy

1.2 Power Subsystem

The electrical power for the system is comprised of 4 isolated 12 VDC deep cycle storage cells. The batteries are stored in the trunk area as shown in figure 2 below from where power is distributed to the front of the car. The batteries are connected to the vehicle alternator via an isolator circuit that allows them to be charged without compromising each circuit with noise feedback through the power supply. A voltage converter is installed in the car to provide 24 VDC for the servomotors. Military standard Amphenol power connectors are used to ensure the integrity of the circuit connections in the presence of a high vibration environment. Some redundancy has been built into the power circuit by utilizing a secondary power bus for critical equipment in case one of the batteries fails. A UPS operating from one of the 12 V batteries provides 600 Watts of AC power for the processor boards.

The peak power budget is calculated to be approximately 1200-1500 Watts. Quiescent power for computers and other ancillary equipment is expected to be about 600-800 watts. As a result, the car's 1200-watt alternator should be sufficient for maintaining system power. A new pulley for the alternator has been installed to increase the quiescent power at engine idle. An additional fuel tank with a 10-gallon capacity has been added to ensure that pause periods of as long as 10 hours may be tolerated in case of delays in the race. The vehicle's air conditioning system is utilized to provide needed cooling for the electronics.

2. Autonomous Operations

2.1 Processing

All processors in the Doom Buggy are Xeon and Pentium-4 processors. Two of the processors are dual processors to provide enhanced processing speed for data intensive operations such as imaging processing. Interprocessor communication is through small data packs over a 100 Mbit Ethernet shown in green in the figure. Inter processor image movement is performed over a 1Gbit Ethernet which is reserved solely for that function and is shown in blue. Between the speed of the net and the design of the task transfer messages there is no need for any sophisticated real time message control.

Figure 3 below shows the system architecture. The larger size boards represent the 4 Xeon processors. The smaller boards are PentiumD processors. Servo controls are provided to point all environmental sensors and to control the car. One processor, to the left of the figure, provides main system control as well as execution of the car's steering and speed control. Digital controls are provided to allow gear shifting and remote start capability. The Estop is connected to this computer as well as to the car external to the computer. Thus the computers can read the pause/run mode to allow the car to enter correct resting states during prolonged pause modes.

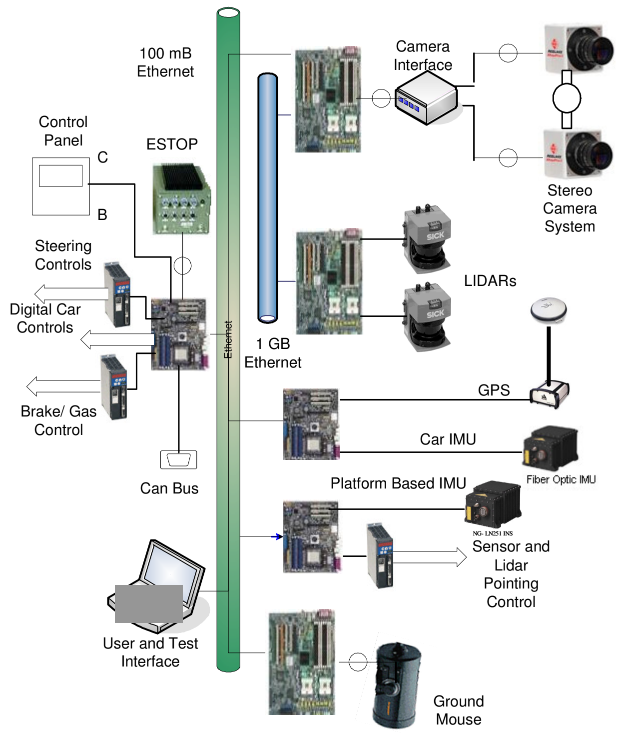

Figure 3- System Architecture

Both cameras are read into a single board Pentium dual Xeon card. The rate at which image capture can be accomplished is at 20 Hz. limited by I/O. Processing to determine clear path is performed and the result is passed to the LIDAR processor. Here sensor fusion is performed between the two types of environmental sensors and a real time 3D model is generated. The last of the dual processor boards is devoted to the ground mouse that has its own intense processing requirement. The last two processors are devoted to navigation. They keep track of the state vectors of both the isolation platform and the car. The system actually navigates to the platform state vector, which has a strongly suppressed vibration power spectral density. The car state needs to be tracked so that sensors located off of the platform can be commanded to a required line of sight direction.

The top-level software layout is shown in figure 4 below. The three main modules are State Control, Driver and Shotgun. The State Control, which is on the left side processor card in figure 3, provides Health and Safety monitoring, vehicle state control and master clock control. This module controls autonomous startup and shutdown procedures. The Driver module implements all of the servo commands for steering and speed control as well as navigation sensor processing and state vector update and determination. The Shotgun is responsible for sensing the environment and warning the driver to change planned paths due to obstructions in the planned path. The modules reside on multiple processors and due to the optical processing, are the most processor intense function of the vehicle.

2.1.1 State Control

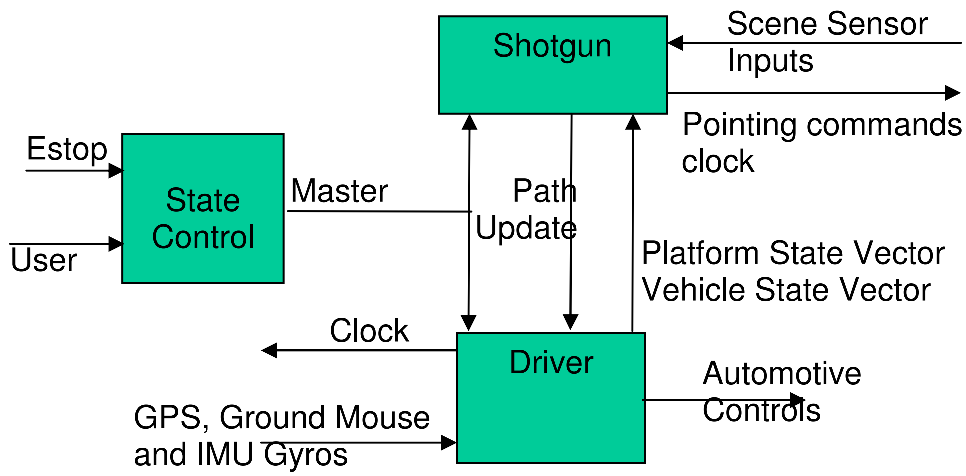

Figure 4 –Top Level Software

The State Control performs a variety of functions related to defining the operating state of the car.. Input is accepted from the Estop to determine when the vehicle shall be allowed to switch in and out of operational states that are under software control. User input is also accepted to allow manual override of the robot for either manual driving or under certain test conditions. The CAN bus is also utilized to determine various state parameters of the vehicle such as engine rpm, wheel rpm and speedometer readings which can be used to monitor the safe operation of the vehicle and take corrective actions when called for.

2.1.2 The Driver

The driver module is provided with measurements from the ground mouse, two sets of gyros and GPS. One set of gyros measures the attitude of the inertial platform that houses the ground mouse and stereo viewing system. The second set of gyros measures the attitude of the car. From this data an appropriate set of filters are applied to provide a state vectors for the automobile and the platform. The state vector data is used with the planned path information to provide both steering and velocity commands for implementation by the servo systems. Appropriate checks are made to insure that the commands fall within a range that the car is able to perform considering both design constraints and vehicle dynamics.

The velocity command takes many dynamic and static conditions into account in determining what speed to command the vehicle to go. Looking ahead along the planned path, the safe speed to execute the planned turn radii ahead are calculated. Their positions are propagated back in time to the current time, and the change in velocity needed to reach these speeds is calculated. The critical speed as controlled by a maximum acceptable acceleration command is determined and the car is ordered to slow down to meet these speed criteria as the vehicle reaches the planned position. If the path ahead is determined to be excessively rough due to the presence of a field of rocks, a safe speed is calculated based on the measured roughness. A measurement of the current road roughness is also done utilizing the accelerometers and a safe speed is calculated from these measurements as well. Finally, all of these speeds are compared with the government provided speed guideline for the current segment of the course. The speeds are combined to provide the speed command for the car. In this fashion the car can react to road curvature, road roughness and traffic control in a manner quite similar to a human driver.

The driver module also calculates the updates to the platform and car state vectors. Multiple measurements of the data is provided by the ground mouse, the GPS, the two IMUs. The ground mouse provides both transverse and lateral velocity components as well. The data is combined using a Kalman filter to provide the required state vectors. The GPS receiver may be in one of several modes depending on signal strength and topographic features. Thus when the system detects a change from HP mode to WAAS mode for instance, the appropriate covariance matrix is used to take into account the signal degradation.

Finally, the driver module is provided map updates from the Shotgun. The map takes the form of a horizontal plane projection of the sensed path ahead of the vehicle. The new information provided on the map represents the newly revised "clear path" data as sensed by the stereo camera system and the lidars. A search for the cleared route is performed utilizing varying weights for corridor edges, objects in the path and varying boundary conditions at the corridor edge. Thus the corridor edge where there is vegetation or just more dirt path found at the corridor edge will get a smaller avoidance weight then a rock ledge or a cliff edge. The route now picked diverges from the position of the robot outwards away from the previously planned path. The new steering and velocity commands are derived from this planned route. A parameter is also passed from the shotgun to the driver, which deals with road smoothness for use in calculating the velocity command that the car is to execute.

2.1.3 The Shotgun

The shotgun is the most hardware intense module in the car. There is one dual processor card devoted to the cameras, and one to the Lidars and sensor fusion function. The stereo cameras are the primary vision system. They are tasked with providing a measurement of the clear path at long distance. The cameras operate in a tasked mode so that the IFOV is limited to a small angular window in the direction that the car is planning to traverse in the immediate future and a window large enough to extend to the horizon line. This is the technique used by the brain in processing the retinal data. The IFOV of the cameras is not more than 10 ^{o} to 15^{o} azimuth by 5^{o} elevation. This is compared to the human eye whose central detection envelope is 5^{o} azimuth by 5^{o} elevation, albeit the eye maintains 30 million pixels in this region. Like the human eye, the sensor is pointed in the direction where the "brain" tells it that high quality detection is needed. At 100 meters range where the cameras are needed to provide enhanced safe speed limits, the field of view is not wider than approximately 40-60 feet.

At the start of the race, the lasers sense the clear path for a distance up to 30 meters. The cameras are then commanded to clear path for the rest of the visible field. This allows us to set up the initial environmental model and determine the initial path for whatever distance is visible to the sensors. The velocity filters set the maximum speed the vehicle will be allowed to go by use of the initial clear range data. The corridor has been calculated and the direction that the camera will have to look at in order to sense the next clear range sector is known and commanded. Additional scans are taken and the map updated by the shotgun with the new sensor data. The lidars provide additional information as to potential objects in the FOV. Resolving the difference between misidentified shadows and objects is accomplished by the lower lidar. The upper lidar can measure ruts and ditches but is more primarily used to understand the sides of the road. If there are rock walls or sheer drops, the top mounted lidar will readily detect. The weights of the side boundary for the object avoidance algorithms are adjusted accordingly.

2.2 Localization

The primary sensor used for localization is a proprietary sensor called the ground mouse. The prototype is shown mounted in the doom buggy in figure 5 was used during the site visit. The final unit will be integrated later in August. The ground mouse was developed specifically for the Grand Challenge. The ground mouse consists of a narrow spectral filtered light source, a detector array and a modified Newtonian optic

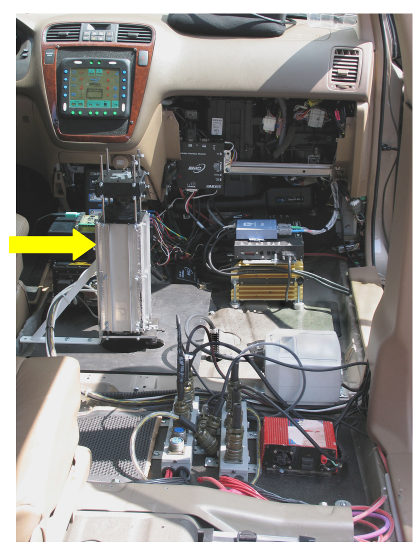

Figure 5 Ground Mouse

system. The distance of the sensor to the ground is not critical due to the optic design so there is no loss of accuracy due to vertical motion of the car on its suspension. The ground mouse takes measurements every 2 msec and integrates the position change until requested by the navigation system. The measurement consists of the longitudinal and transversal displacement of the center of the ground mouse. As the time is accurately known, the data also provides the average velocity in the measurement period as well. Measurements of the system show an accuracy of a few mm with the prototype system. The final system, which has a higher resolution sensor, is expected to have an accuracy of approximately 1 mm. If the ground mouse is correlated with a GPS location at some initial time and position, then the position of the car is well known thereafter from the ground mouse measurement alone. However, to maintain accuracy using a projection of the earth to a locally flat coordinate system, especially in the presence of elevated terrain, the ground mouse is recalibrated to the GPS periodically. The recalibration is performed only if the GPS system is in the HP mode where 0.5 m accuracy is achieved while the vehicle is in motion. The ground mouse works in conjunction with the IMU. The three gyro outputs are used to sense turn rate for each of the three angular axes. The ground mouse uses the yaw rate output in particular to correct for image rotation during the 2 msec sample periods. Looking at the power spectral density of the accelerometer output for a surface similar to open dessert and at speeds in the range of up to 35 MPH suggests that the traditional approach of using a Kalman filter on the navigation sensor suite to provide optimized estimates in a noisy environment does not contribute any substantive improvements for any period longer than at most a second or two. Road vibration completely dominates the measurement process and the low signal to noise ratio does not warrant the use of the traditional approach.

The IMU chosen to track platform position is the Northrop Grumman LN250. The particular unit that we have obtained has a drift rate of $.06^0$ per hour. Thus the angle errors are quite small for substantial periods of and can be used with the ground mouse data to provide high accuracy state vectors independent of GPS. With the drift rate of the IMU, the expected error over the entire course would be expected to be of the order of .01 radians if the course were finished in 10 hours. This translates into a position error of approximately 2.5 km. The contribution of the error due to the ground mouse errors is completely negligible in comparison. In order to maintain a destination expectation error of less than 1 meter, the IMU would need to be recalibrated every

300 feet on average. Practically, there is no need to recalibrate that often as route planning would account for these errors whenever a recalibration was performed. Thus, over the 100 foot GPS tunnel outage, the expected propagation state error of the vehicle will be of the order of 1 cm.

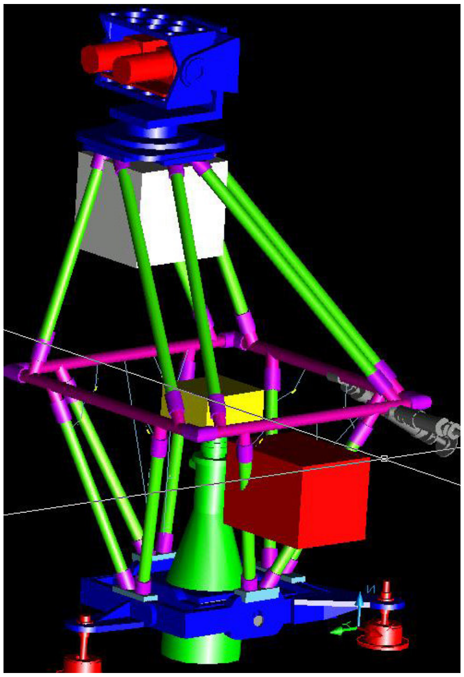

Figure 6- Motion Platform

2.3 Sensor Subsystem

2.3.1 Sensor Suite

The sensor suite consists of two sick lidars mounted on active 2 axis controlled platforms. Additionally there is a 3 axis controlled stereo camera system which is mounted on the motion platform internal to the car. Each sensor may be commanded to point in a desired direction. Figure 6 shows a CAD design drawing of the motion platform internal to the car. The high quality IMU (red box in lower right quadrant) is mounted to the platform and gives precision angle displacement data for the platform with respect to the normal coordinate plane. The ground mouse (green structure mounted to blue gimbal at lower half of drawing) is a two axis controlled system. The stereo cameras (blue and red structure at top of drawing) are mounted on a 3 axis controlled gimbal system.

2.3.1.1 Lidars

Figure 7 shows the location of the two sick lidars on the vehicle. The lidar mounted on top of the vehicle is tasked with sensing road ruts, pits, trenches and any other road depressions. The lidars are shown hard mounted to the vehicle. The active platforms will be integrated onto the lidars during the last week in August. The locations of devices will be approximately as shown but with modifications to support the spatial needs of the motion base. The pitch and roll angle can be commanded over a $\pm 30^0$ interval. The principal sensors for controlling the

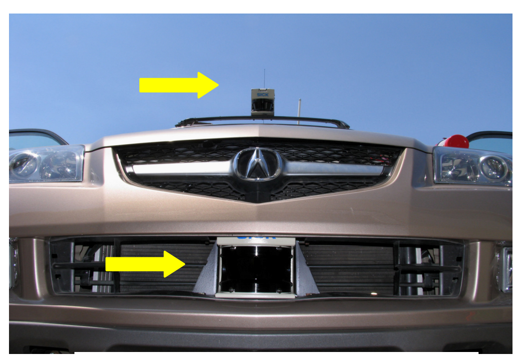

Figure 7- Lidar locations

pitch and roll angle of the laser are the angle resolvers for each axis and the IMU hard mounted to the car. The instantaneous roll angle of the car is cancelled out by commanding the opposing angle into the servo control. The pitch angle is maintained at a fixed negative pitch to provide a slant range of approximately 30 meters to the road. Ruts are seen as discontinuities in the slant range to the road. The laser is kept clear by running a film past the laser detector so that the window is cleaned continuously. The lower laser is tasked with object detection in the short to mid range. We do not expect to see objects further than 30 meters with the lidar and will not track them within 5 meters of the vehicle. The laser is angled downward so that a 6 inch size object is detectable at 30 meters.

2.3.1.2 Optical Sensor

The Redlake video cameras are mounted about half a meter apart on the same motion platform as the ground mouse. They are inside the passenger compartment protected by the windshield of the car as shown in figure 8. Unobstructed visibility through the windshield is provided by the windshield wiper cleaning system, which is activated periodically or on dirty windshield conditions by the main processor.

Figure 8- Stereo Camera on 3 axis gimbal.

The mount is an active 3-axis gimbal. Instantaneous angle data is derived from the IMU mounted to the motion platform. The angle resolvers on the gimbals provide the LOS with respect to the platform. The required rotations are calculated to bring the LOS into the desired inertial direction. The system is commanded to that direction and the camera samples the IFOV.

2.3.2 Environmental Model

The environmental model is built up each scan of the sensors. It is a 3 dimensional representation of the path toward which the vehicle is moving. The transverse axis of the model is the next guideline to be encountered. The longitudinal axis is the normal at the position where the car is planned to cross the next guideline. Local gradients of the road plane are used to detect obstructions in the path and weight them accordingly.

Initially, the model is built from the laser scans and augmented by visual scans based on where the laser information becomes unreliable. A field of obstacles is built up which is proportional to the local gradients with a low number being a high likely clear path. The boundaries have variable weights based on information from the lidars. Thus rough vegetation or other forms of non critical side information will get a small proportionality constant and rock walls or sheer drops will get high proportionality constants. In this way, the route planning software can preferentially allow boundary violations (e.g. let the left wheel drift onto a shoulder to avoid a hazard at the center of the road) to avoid hazards or force the vehicle to avoid an object but strictly enforce the boundary.

2.3.3 Object Avoidance

Pictured on Figure 9 is a sample test screen from the object avoidance algorithm. The 2 dimensional sky view is updated by transforming the cleared path data from the three dimensional environmental model. Objects are shown as white blobs on the plot below. The current vehicle position is indicated by the blue cross hairs. The green bar is the next guideline to be encountered.

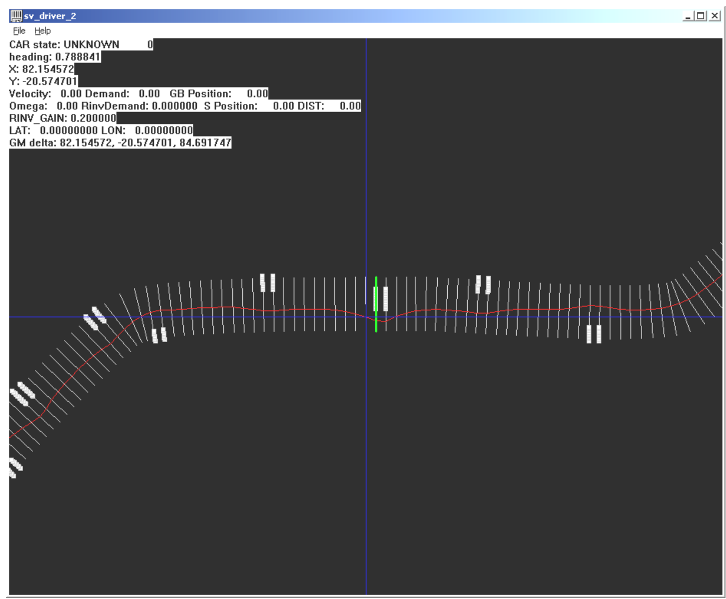

Figure 9- Object Avoidance Result

Planning is done for several 100 meters based on a combination of sensor data and preplanned corridor placement. On each clock cycle, new objects are added to the map and the path is optimized based on current information. There are both smoothing criteria and vehicle constraint information built into the path planning process. The red line shows the planned path in the above example. At each time step, this line may change but it always changes from the current position of the vehicle. Reassessments of the avoidance weights between objects and the corridor affect which choice of path will be selected.

2.4 Vehicle Control

2.4.1 Contingencies

2.4.1.1 Missed Waypoints

There is no need for any special consideration for missed waypoints using the method of navigation adopted by the PVRW. As shown in figure 9 above, the waypoints are used solely to construct a series of guidelines as represented by the red crosshatched line. They are constructed using a Bezier spline to construct a smooth path through the boundaries formed by projecting the waypoints backward and forward at the corridor distance in the direction of the previous and next waypoints respectively. A corridor is constructed by projecting points at the local waypoint accuracy requirement projected in either direction along the local normal. The guidelines are constructed along these lines spaced approximately 20 cm apart. All legal paths are within the bounds of this corridor with some exceptions noted below. The guidelines in conjunction with the shotgun have lain a clear path ahead which is continuously updated by new sensor observations. The robot steers along this planned path and the waypoint has no further meaning.

2.4.1.1 Car Stuck

There are two preplanned cases of the vehicle becoming "stuck". In one case, the vehicle becomes lodged on a feature such as an undetected rock. With the native four wheel control system of the Honda MDX, such an occurrence is a remote possibility. The car has both slip differential and traction control systems. They allow power to be distributed amongst the four wheels so that a free wheeling tire gets little power and more is distributed to the wheels with positive traction. A wheel, which is not turning, will be given more power so as to help it to dislodge. There are however possible geographic features such as certain types of trenches where if the car where to get stuck in one above the critical size, then there would be no recovery. On the other hand, if the vehicle were to come to a situation where the shotgun could not see a clear path in front, then a back up procedure is planned. The system is capable of storing the previous 30 minutes of data. The vehicle can then execute a reverse of its motion. The shotgun would search for an alternative route. When found, the car would move forward upon the newly planned route.

2.4.1.3 Over- Boundary Obstacle

As described above, there is a weight associated with the boundary as well as objects. Objects always have a high weighting to ensure they will not be hit. The boundary on the other hand may be a smaller weight. If there are features such as a rock wall on one side and a sheer drop on the other side, then the weights are as high as objects and if a path cannot be cleared, the vehicle will be “stuck”. If on the other hand, there is shoulder, vegetation or other non critical features at the side of a trail, the vehicle can choose to violate the corridor in order to plan a path around objects that would otherwise prevent it from moving forward. The larger the distance from the boundary, the larger will be the weight thus if a boundary needs to be violated, it will do so at the minimal distance possible.

2.4.2 Maneuvers braking, starting on hill, sharp turns without boundary violation

2.4.2.1 Speed Control

The speed control algorithm takes many different parameters into account in order to determine the velocity command that the servo should execute. One of the parameters that control the speed is turn profile in the path ahead. The vehicle is programmed with an acceleration limit on turns. In order to meet this criterion, the maximum acceptable velocity along any point of the planned path can be calculated. Based on a combination of the current speed, the range to the point and the maximum acceptable speed at that future point, a braking command may be executed causing the vehicle to slow down so as to achieve the desired velocity at that future turn point. Other conditions such as road roughness or dynamic stability issues produce their own velocity command. The actual command is the minimum of the speed commands calculated from a list of constraints.

2.4.2.2 Starting on Hills

The only potentially stressing condition for the Honda in hill starting is if the vehicle is facing uphill. The native sensors of the vehicle detect only speed; they do not detect the direction of the velocity. When the vehicle is commanded to start, it is put in gear and the break is held. If at this time the vehicle were to start sliding down hill, the combination of the IMU indicating that the vehicle is pitched upwards and a finite velocity while the brake is on would indicate sliding down hill. If we are not ready to start as yet, the brake will be depressed further to hold the vehicle. Corroboration of the slide downhill will come from the ground mouse by a negative velocity. Once the velocity command is issued, the PID loop will enforce the proper motion.

2.4.2.2 Enforcing Boundary on Sharp Turns.

The best technique for avoiding boundary violations on sharp turns is to avoid sharp turns as much as possible. The first issue is to plan a path which minimizes such turns. The Bezier spline adopted for the planned corridor removes as much of the "planned sharp turns" as possible. There were no waypoints in the 2004 Grand Challenge course that provided sharp turns using the algorithms implemented for 2005. The most likely occurrence of a sharp turn is thus likely to be due to object avoidance. The minimum turn radius of 5.8 meters is a constraint used in the route-planning algorithm. The speed control algorithm will automatically cause the vehicle to slow down as the turn radius decreases until the minimum of acceptable acceleration and roll stability is achieved.

2.4.3 Integration of Navigation with Sensing

The interface between navigation and sensing is readily understood from the information flow as shown in figure 5 above. The state vector determined in the driver module is passed to the shotgun that uses position and orientation data to construct and update the 3D environmental model. The newly sensed data is added into the model and cleared regions or newly discovered blocked regions are marked. This information is transformed into the 2D sky view model centered about the vehicle as shown in figure 9 above. In this fashion the task is divisible between the driver/navigation function and the shotgun/environmental sensor tasks.

2.4.4 Manual Operations

The vehicle has two modes of operation both of which are controlled by the computer. The computer must be powered up in order for the car to operate. Ignition is controlled by the state of a switch on the vehicle control panel. If the ignition switch is on, the computer will start up the engine as part of its initialization sequence. Otherwise, the computer will wait to sense the change of state of the ignition switch to determine of the engine should be on or off. If the switch is turned off while in operation, the brake is depressed to stop the vehicle, and the engine is shut down. The H&S subsystem ensures that the vehicle is brought back on line safely. On the same control panel is a switch labeled manual/autonomous. In manual position, the car is drivable by a human driver in precisely the same manner as an unmodified vehicle.

2.5 System Test

Testing philosophy for the Doom Buggy required testing as each subsystem is developed, built, and incorporated into the vehicle. Both performs as designed and interface verification with the other subsystems are measured. Tests are performed in real-time to verify the performance of each new feature and to ensure that no existing capability is compromised. For example, many of the inherent capabilities of the vehicle have been verified during initial autonomous testing under controlled conditions at PVHS. Velocity control for example was performed in steps from 1 to 35 kph where both servo response time and long term accuracy were measured. Step function responses over large command changes were also made. All safety equipment, including flashers, a siren, four manual external emergency stop buttons, and the E-Stop supplied by DARPA have also been installed and tested.

A set of formal operational requirements have also been generated for each subsystem and for the vehicle as a whole, which together comprise the set of necessary vehicle capabilities to ensure success at the DGC. Each requirement will be verified during the integration and test phase. Many of the requirements will be verified by analysis. Whenever possible, specific requirements will also be verified by testing. A set of tests has been generated to verify each subsystem and vehicle requirement. Many of the subsystem tests will be carried out under controlled conditions to verify the operational requirements for each subsystem. The vehicle tests will be carried out over a test course that simulates the environment anticipated for the DGC. Reliability will be determined by extrapolation of vehicle and subsystem performance changes during repetitive long-term testing. Once all specific requirements have been verified, the vehicle will continue testing under adverse conditions. This additional testing will provide a high degree of confidence that the PVHS entry in the DGC will demonstrate safe and effective autonomous operation and will be fully capable of meeting the many challenges that are anticipated during the DGC.