

The TerraMax Autonomous Vehicle

Deborah Braid

Rockwell Collins
Cedar Rapids, Iowa 52498
e-mail: dmbraid@rockwellcollins.com

Alberto Broggi

VisLab
University of Parma
43100 Parma, Italy
e-mail: broggi@ce.unipr.it

Gary Schmiedel

Oshkosh Truck Corporation
Oshkosh, Wisconsin 54902
e-mail: gschmiedel@oshtruck.com

Received 22 March 2006; accepted 20 June 2006

The TerraMax vehicle is based on Oshkosh Truck's Medium Tactical Vehicle Replacement truck platform and was one of the five vehicles able to successfully reach the finish line of the 132 miles DARPA Grand Challenge desert race. Due to its size (30 000 pounds, 27' 0" long, 8' 4" wide, and 8' 2" high) and the narrow passages, TerraMax had to travel slowly, but its capabilities demonstrated the maturity of the overall system. Rockwell Collins developed, integrated, and installed the intelligent Vehicle Management System, which includes vehicle sensor management, navigation, and vehicle control systems. The University of Parma provided the vehicle's vision system, while Oshkosh Truck Corp. provided project management, system integration, low level controls hardware, modeling and simulation support, and the vehicle. © 2006 Wiley Periodicals, Inc.

1. INTRODUCTION

TerraMax™, a completely autonomous vehicle, was developed by Oshkosh Truck Corporation in cooperation with its partners Rockwell Collins and Uni-

versity of Parma in response to Congress' goal that one third of military vehicles be unmanned by 2015. The Oshkosh TerraMax™ was one of only five vehicles that successfully completed the 132-mile DARPA Grand Challenge course in October 2005 (5th



Figure 1. The TerraMax vehicle.

place), and it was the only vehicle whose mission is to provide medium- to heavy-payload logistic support to the battlefield. During the race, the fully autonomous vehicle was successful in demonstrating obstacle avoidance, negotiating tunnels, narrow roads and cliffs, GPS waypoint following and 28 h of nonstop continuous operation—all applicable to military missions.

2. THE VEHICLE

The TerraMax vehicle shown in Figure 1 is based on Oshkosh's Medium Tactical Vehicle Replacement (MTVR) MK23 truck platform. The MTVR was designed with a 70% off-road mission profile. It can carry a 7-ton payload off-road or a 15-ton payload on-road. All-wheel drive, TAK-4™ independent suspension, and central tire inflation make rocks, dips, holes, and crevasses easier to handle. And the truck can handle 60% grades and 30% side slopes. A 425-hp Cat C-12 engine powers the truck. This kind of vehicle was chosen for the DARPA Grand Challenge (DGC) because of its proven off-road mobility, as well as for its direct applicability to potential future autonomous missions. Team TerraMax also participated in the 2004 DGC (Ozguner, Redmill & Broggi, 2004) with the same vehicle. Two significant vehicle upgrades for the 2005 DGC were the addition of rear-wheel steering and integrated sensor structure/roll cage. Rear steer has been added to TerraMax to give it a tighter 29-ft turning radius. Although this allows the vehicle to negotiate tighter turns without needing frequent backups, the backup maneuver is required to

align the vehicle with narrow passages. The sensor mounting structure/roll cage provided added protection to the sensors as well as key vehicle components.

2.1. Autonomous System Integration

The autonomous system consists of computers, communication network, sensors, vehicle control interface, and the supporting mounting and protection structures. The autonomous system utilized in the 2004 DGC was completely removed and upgraded for the 2005 DGC.

2.2. Computers and Communication Network Integration

The computers and communication network hardware was packaged in a modular shock absorbing rack located inside the base of the passenger seat as shown in Figure 2. The video monitor, keyboard, and mouse were securely mounted on the dashboard. This arrangement allowed the sensitive computer equipment to survive the high-G shock and vibration experienced on the trail.

2.3. Sensor Installation

The sensors were mounted to modular adjustable mounts that were integrated into the roll cage. The roll cage also serves as a protective conduit, through which, the vital sensor communication and power cables are routed. The adjustable mounts used were selected for their ability to retain the set position regardless of the pounding taken from the trail. The location of each sensor was optimized for functionality while maintaining a high level of protection from the environment, i.e., rain, sun, engine heat, brush, and, yes, even bridge supports. A sensor cleaning system was also developed to keep the lenses of the TerraMax sensors free of debris such as dust, water, and mud. The main components of this system are: cleaning controller, valve array, and washer tank. The cleaning controller controls the sequence and duration the sensors are dusted, washed, and dried. The valve array has electrically controlled valves that pass pressurized water and air through pattern nozzles to the sensor lenses.




Figure 2. Computers mounted under passenger seat: (a) Computer aided design (CAD) simulation; (b) real installation.

2.4. Vehicle Control Actuator Integration

The vehicle control integration was comprised of four key areas: brake, throttle, gear selection, and steering. The brake was controlled via Ethernet by a proportional voltage to pressure valve. This method was utilized, in-part, because of safety concerns of mechanical systems interfering with a driver while on the road. The throttle was controlled via an I/O card in the vehicle manager computer. A pulse width modulated (PWM) signal allowed precise control of engine throttle level. The gear selection was controlled via a relay card in the vehicle manager computer. A binary pattern applied to the transmission control harness allowed the ability to select the de-



Figure 3. ADAMS model of the MTRV negotiating a simulated obstacle.

sired gear required by the trail conditions. The steering was controlled via serial communications to a servodrive and motor. The servomotor is connected in parallel with the steering wheel shaft through a gearbox. The servomotor has an integrated high-resolution encoder that allows precise control of wheel angle.

3. TERRAMAX MODELING AND SIMULATION

Modeling and simulation efforts supported the controls development by providing information such as underbody clearance, steer angles, and lateral stability. A full vehicle model of the truck was created in Advanced Dynamic Analysis of Mechanical Systems (ADAMS) by assembling subsystem models of suspensions, steering, chassis, and tires. A typical NATO Reference Mobility Model (NRMM) obstacle course with over 70 different obstacles of different sizes and shapes was used to evaluate the underbody clearance (see Figure 3). The results of this simulation gave an idea about the truck's capability to maneuver through different obstacles at low speeds.

The steering model was used to predict the front and rear steer angles (see Figure 4) for a given steering wheel input. The rear steer model included a dwell and had different gear ratios than the front.

The lateral stability of the truck was evaluated through constant-radius tests. Tire forces were monitored to detect tire lift-offs. The results of these simulations as shown in Figure 5 were used to evaluate the capability of the truck to take a particular turn at different speeds without rolling over.

4. VEHICLE MANAGEMENT SYSTEM

The Rockwell Collins intelligent Vehicle Management System (iVMS) (Braid & Johnson, 2006) consists of hardware and software components that together

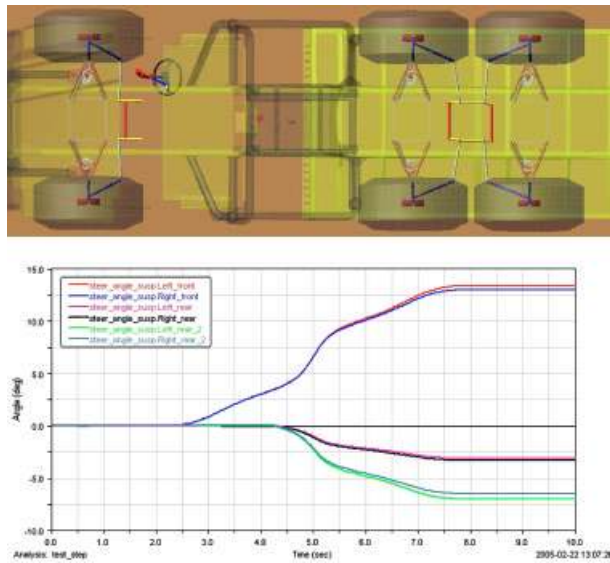


Figure 4. Steering angle behavior of each wheel end throughout a 360° rotation of the steering wheel originating from a straight ahead position.

provide an extensive set of autonomous capabilities. In order to accomplish this, the iVMS interfaces with the vehicle systems and all onboard sensors. The primary commands to the vehicle interface are throttle, brake, steering, and transmission.

The general architecture for the iVMS software is a set of applications that communicate to each other over a 100BaseT Ethernet network utilizing transmission control protocol (TCP) and user datagram protocol protocols and a commercial Ethernet switch. The iVMS software has the key role of performing all autonomous behavior and interfacing to numerous line replaceable units and the key vehicle systems. The software applications are as follows:

- Vehicle control—controls and receives feedback from the throttle, brakes, and steering in order to control the vehicle while in autonomous mode.
- Real time path planner—computes the real time path utilizing the desired path while avoiding the obstacles along the desired path.
- Obstacle detection—uses LIDAR and Vision to detect positive and negative obstacles. Obstacle data coming from the various sensors are merged into a single obstacle database used by the real-time path planner.
- Behavior management—decides what mode the vehicle should be in based on the current conditions of the other functions
- Navigation—computes present position and provides a dead reckoning function.




Figure 5. Lateral stability simulation of a MTVR traveling on a constant radius path with increasing speed. Graph depicts weight transfer on each of the six wheel ends. Simulation was conducted at the maximum gross vehicle weight rating of 58 000 lb.



Figure 6. The driving cabin, with the monitor showing the graphical user interface.

A graphical user interface (GUI) provides multiple functions to the user including data visualization, recording, and playback. The GUI is primarily a development tool and is not considered to be an integral part of the real-time iVMS system. Figure 6

shows the GUI displayed on a monitor in the cab.

A system management function is also implemented that provides a user interface for execution control and status display for the iVMS applications. Once the system has been initialized, the system manager performs a health management function that continuously monitors the status of the application and automatically stops and restarts applications as necessary to maintain normal functionality. The iVMS can continue to operate normally without the system manager once initialized so it is not included as one of the iVMS applications. The system architecture can be viewed in Figure 7.

The following sections of this paper will go into further detail on each of the iVMS functions.

4.1. Vehicle Control

The vehicle control function of the iVMS provides the TerraMax control actions that emulate the actions a human would perform when driving the truck. The controls provided by the iVMS are steer-

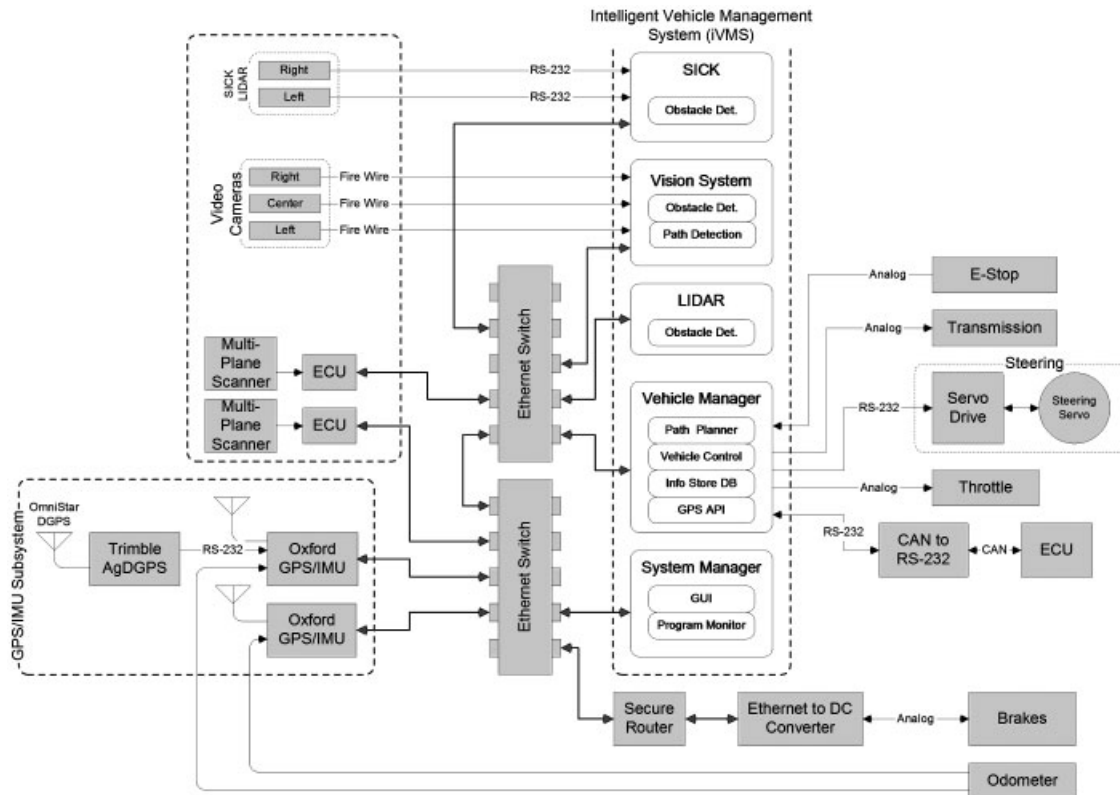


Figure 7. TerraMax iVMS system architecture.

ing, throttle, brake, and transmission control. Steering control is provided through an electronic servo connected directly to the MTVR steering gearbox. The standard MTVR steering gearbox has dual inputs so the steering servo for autonomous operation and hand wheel are both connected to the steering gear allowing the steering control to be switched between manual and autonomous operations without changing mechanical linkages. The steering control function is responsible for providing wheel angle commands that guide the vehicle to the path defined by the real-time path planner. This is accomplished by computed deviations from the desired path and converting the deviations to steer angle commands that are sent to the steering servo. The steering control uses capture and track steering control modes. Capture steering control is used for initial course capture and track steering control is used during normal operation. Capture and track control modes are automatically selected based on current conditions.

The capture controller uses course error as the control parameter. The controller creates a steer angle command that aligns the ground track of the vehicle with the direct bearing to the active (TO—"next") waypoint. This type of control is sometimes referred to as homing control since the path followed is uncontrolled and the path to the TO waypoint may not be a straight line. Capture conditions occur during initial course capture so the capture controller is only active if certain conditions exist at the time the autonomous mode is first activated.

The track controller uses linear cross track deviation and cross track deviation rate to align the vehicle's path along the ground with the active TO waypoint course. Track angle error and steer angle command limiters are used to limit the commanded steer angle to values that are achievable by the vehicle. The command limiters incorporate vehicle dynamic limits with margins built in to ensure the vehicle does not get into an unsafe condition. This also means that the vehicle operates at levels below its maximum dynamic capability when in autonomous mode. Turn anticipation for waypoint sequences is also used so the transition onto the new course is accomplished without overshoots.

The throttle controller interfaces directly to the electronic engine control unit through a digital PWM interface. The throttle controller is responsible for controlling the vehicle's speed to the desired speed specified by the path planner. This is accomplished

primarily through throttle position control but engine and service brakes are also used in certain situations to manage the speed.

The throttle position control uses proportional and integral control. Reset conditions to the throttle position are provided for transmission up shift and down shift and to activate the engine brake. Engine brakes are activated during engine idle so throttle position overrides are used when engine brakes are required. Throttle position faders are used to reactivate the throttle position control when the engine brake is disabled. Engine and service brakes are used primarily to control speed on steep grades and for speed management during deceleration.

The brake controller provides an analog signal to a pressure actuator connected to the air brake system (service brakes). The throttle and behavior control functions provide brake actuation parameters to the brake controller and the brake controller determines the pressure actuator signal. The brake control parameter provided by the throttle control function is speed deviation which is used by the brake controller to provide a brake application that is proportional to the speed deviation. Behavior control provides brake override signals for emergency stop (e-stop), e-stop pause, and other special situations requiring speed control or position holding. The emergency stop condition results in a full brake command. Brake modulation to limit slipping in full brake conditions are provided by the antilock brake system that is part of the basic MTVR.

The MTVR has a seven speed automatic transmission. The transmission control function provides forward, neutral, and reverse gear control for the automatic transmission. The selection of the transmission gear is through a digital signal to the transmission control unit. The transmission controller receives a desired gear signal from the behavior control function and converts the desired gear into the digital interface to the transmission. Behavior control uses the actual gear position to determine allowable state transitions and to prevent transmission faults due to incorrect gear selection sequences.

4.2. Real-Time Path Planner

The real-time path planner (see Figure 8) is responsible for deriving the desired trajectory of the vehicle and providing that trajectory to the vehicle control function. The trajectory includes a desired path along the ground as well as the desired speeds and

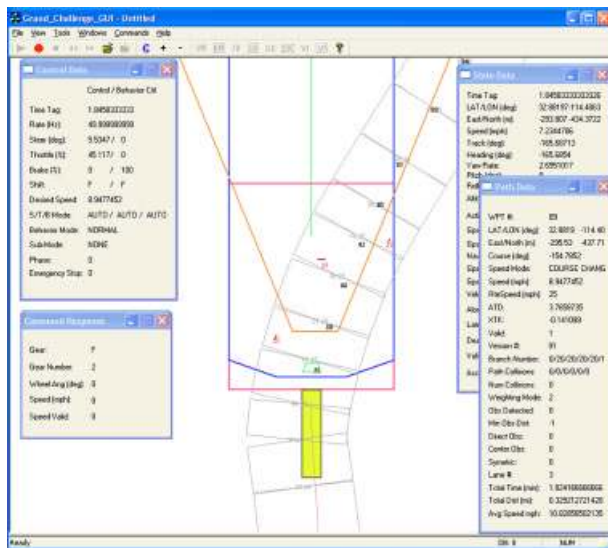


Figure 8. iVMS graphical user interfaces depicting real-time path planning along a defined route.

boundary area. The desired trajectory is derived using the path and speed constraints contained in the DARPA route data definition file (RDDF), which contains a list of waypoints that define a path along the ground, a path boundary, and maximum speed for each leg of the path. The real-time path planner provides reactive path corrections to this nominal path to account for current conditions, such as vehicle dynamic limits, obstacles, road edges, terrain grade, etc.

The path planner implements a tree algorithm that branches from the base at the current TO waypoint. Constraints for path boundary and speed are applied to the tree build function so the tree size is bounded by the constraints. Branches of the tree are computed using a model of the steering system and vehicle dynamics to insure that the candidate paths are drivable. The tree algorithm was derived from the rapidly-exploring random tree path planner (Kuffner & LaValle, 2000) where the growth of the tree was limited to a fixed number of branches (levels).

Once built, the tree represents a series of candidate paths, one of which is selected as the path to be used by the vehicle control. Selection of the best path from the candidate paths is based on a scoring algorithm that considers distance from the route centerline, path curvature, obstacle avoidance, boundary

area constraints, and other factors. Over 2000 candidate paths are evaluated each planning cycle to determine the best path.

The real-time path planner also contains a speed management function that adjusts the speeds as necessary to account for path geometry and current conditions. The initial desired speed is set to the RDDF speed constraint for the leg and the speed management function reduces the speed as necessary.

Figure 8 shows a graphical representation of the RDDF data and the resulting real-time path generated by the path planner. In the main window of the illustration, the central box represents TerraMax vehicle, light grey boxes represent the desired path defined in the RDDF file, and the black numbers near the center of the boxes represent the real-time path generated by the path planner. The width of the grey boxes defines the lateral boundary of the path and the length of the boxes is defined by the waypoints in the file. The short red lines in the diagram are representations of obstacles that have been detected by the perception sensors. Dialog boxes along the sides of the main window show data associated with the path, vehicle state, and control state. As shown in the example diagram, the real-time path is adjusted to the right of the RDDF path center in order to avoid the obstacles in the turn.

4.3. Obstacle Detection

LIDAR and vision sensors are used to detect obstacles in front of the vehicle. Obstacles detected by the sensors are registered to the vehicle navigation position and stored in an obstacle database. The real-time path planner queries the database to determine if obstacle collisions occur on the proposed paths.

Several different types of obstacle clearance information are provided to the path planner to aid in path selection. Obstacle collision information is reported by the database in terms of the closeness of the object collision to the proposed path. Buffer regions of various sizes are used to determine the collision proximity relative to the path.

Bearing and distance to the nearest collision is provided by the obstacle database that is an indication of the proximity of the obstacles to the proposed path. Obstacle distance is used primarily in the speed manager function to lower the speed if an obstacle is in close proximity to the vehicle's planned path.

Road and cliff edges are handled as special cases

by the obstacle database. Since the consequences to the vehicle of breaching a cliff edge are very severe, additional weight to negative road/cliff edges are used. The database also reports if any negative road/cliff edges are in the immediate area that is used by the speed manager to reduce speeds accordingly.

4.4. Behavior Management

The behavior management module is the central “brain” of the system. Its purpose is to monitor and react to dynamically changing conditions. This module receives input from the real-time path planner, obstacle database, navigation sensors, and the vehicle interface module.

Several behaviors have been designed into the behavior module, using a state transition architecture. When a specific event or a change from normal operating conditions is detected, one of the behaviors is activated to handle the situation at hand. Each behavior executes an ordered list of instructions, providing a set of commands to the vehicle controller.

Some of the conditions the behavior module will react to are as follows:

- Transition in e-stop state: When the e-stop is in pause mode, a behavior will command the vehicle to come to a stop. When e-stop transitions to /Run, another behavior is initiated to begin normal operation.
- No valid path ahead: The behavior initiated in this condition commands the vehicle to come to a stop and wait for a valid path. If no valid path is found, it will command the vehicle to back up and try again.
- Obstacle detected behind the vehicle while backing up: Another behavior will stop the vehicle and command it back into normal operation to try to find a valid path ahead.
- A large course change requiring a backup maneuver: The switchback behavior guides the vehicle around a three-point turn.
- Narrow tunnel condition: The tunnel behavior will guide the vehicle through a narrow tunnel, using the LIDAR scan data.
- Stuck between obstacles: If the vehicle cannot make progress along the route because it continues to go back and forth, getting stuck be-

tween obstacles, the stuck behavior will take over. It will first try to position the vehicle at different angles to search for a valid path. If no valid path is found, it then commands the system to ignore low confidence obstacles, in an attempt to eliminate false obstacles. The last resort is to go forward toward the DARPA route, ignoring all obstacles.

The real-time path planner, behavior management, and vehicle control functions work together to determine the actual path the vehicle follows. Nominally, the vehicle follows the path generated by the real-time path planner but that real-time path can be overwritten by the behavior manager based on the current conditions. This design approach is similar to the distributed architecture for mobile navigation (DAMN) (Rosenblatt, 1997) developed by Carnegie Mellon University where the behavior management functions as the DAMN arbiter and behaviors. Unlike the DAMN architecture, the behavior manager uses rules-based decision logic to determine control behavior modes rather than a voting scheme to select the control mode. This approach was chosen over the more complex voting scheme since it is deterministic and more robust, which were considered to be important attributes given the nature of the competition. This approach lends itself to fleet applications, which was also an important consideration in the iVMS design.

4.5. Navigation

Two Oxford Technical Solutions RT3100's (www.oxts.co.uk) supply GPS position information to the iVMS system. The RT3100 is a combined GPS/IMU sensor that provides real-time data even in the absence of GPS signal. The high 100 Hz update rate has a very low latency to insure that the system is using the most accurate position possible. One RT3100 is configured to use DGPS corrections transmitted via RS-232 from an external GPS receiver subscribed to the Omnistar correction service. The other RT3100 is configured to use WAAS corrections.

In the case of loss of GPS signal, such as driving through a tunnel, the IMU portion of the RT3100 takes over and begins dead reckoning. In order to aid the INS solution in dead reckoning mode, a wheel speed sensor on the vehicle provides input to the RT3100. Tests have shown that the wheel speed

input helps to keep the IMU solution stable and extends the time the RT3100 is able to dead reckon.

In the case of a failure or short-term loss of the RT3100's, a second dead reckoner is implemented using sensed wheel speed and wheel angle. This represents an independent backup navigation function. Because of the potentially large errors that can buildup when it is in a dead-reckoning mode, the RDDF boundary area checks in the path planner are disabled so the vehicle can continue navigation relative to the terrain and terrain obstacles for short periods of time.

The RT3100's were capable of dead reckoning after loss of GPS using the IMU but, during field testing, it was found that the error characteristics of the wheel speed/wheel angle dead reckoner were more desirable than the IMU error characteristics. The conditions where the dead reckoner was most likely to be used was while traveling through railroad tunnels and highway overpasses where the GPS satellite signal would be masked. This meant that the vehicle would be moving in a straight line and only in the tunnel for a short time (typically less than 30 s). Under these conditions, the wheel speed/wheel angle dead reckoner was able to maintain a predictable accuracy of less than 1 m. The IMU was also capable of similar accuracies but the error characteristics were less predictable. The predictability of the error was important since the obstacles while in the tunnel (i.e., the tunnel walls) were very close to the vehicle and small, abrupt changes in position error caused the collision detection function to stop the vehicle.

The wheel speed/wheel angle dead reckoner relied on an analytical model to relate wheel angle and speed to heading (yaw) rate so the accuracy deteriorated significantly during turns. The large heading error build up during turns makes this implementation practical only under a very narrowly defined set of conditions.

5. SENSORS

The sensors were carefully selected to provide the required navigation and perception capability. Figure 9 shows the sensor locations on TerraMax. The sensors selected for the DARPA Challenge 2005 are as follows:

- Oxford GPS/INS;




Figure 9. The TerraMax sensor suite: the picture shows the GPS position, the three front looking cameras, the two SICKs, and the two multiplane laserscanners.

- Trimble GPS;
- single-plane LIDAR;
- multiplane LIDAR; and
- forward-looking vision and system.

5.1. Oxford GPS/INS

The OXTS RT3100's are mounted on the floor of the cab on the approximate centerline of the vehicle. In order to obtain a more accurate position solution and eliminate any errors over time, the position solutions from the two RT3100's were averaged together. In the case of a failure of one of the RT3100's, the system will switch to using the remaining RT3100 as the sole GPS source.

5.2. Trimble GPS

The Trimble GPS (www.trimble.com/aggpps132.shtml) is an agriculture GPS unit used to receive differential corrections used by the GPS receivers embedded in the Oxford RT3100's. The Trimble receiver outputs differential corrections at 1 Hz through RS232. In order to output the differential corrections the Trimble receiver is placed in base station mode and must also have a subscription.

5.3. Single-Plane LIDAR

There are two SICK LMS-291 LIDARs used for positive and negative obstacle detection (see Figures 10 and 11). They are mounted on the outermost edges



Figure 10. The SICK LIDARs and the cameras are placed on a rigid bar onto the vehicle hood.

of the front rollbar. They are pointed 10 deg down and 25 deg outward from the truck so that there is good coverage on extreme turns. The two LIDARs are configured to scan a 100-deg scan area with a 1-deg resolution.

The orientation of the SICK LIDARs was chosen to gain visibility to positive obstacles near the vehicle and detect negative road edges. Positive obstacle detection was accomplished by translating the range returns into local level coordinates and com-




Figure 11. iVMS graphical user interface depicting SICK obstacles as green and yellow polygons. In the figure the vehicle is located at waypoint 27 with iVMS calculated micro-waypoints extending ahead of the vehicle.



Figure 12. The front of the TerraMax vehicle. Two four-plane laserscanners are visible: the one inside the bumper is the one used during the race, the other (over the bumper) is a backup.

paring the relative heights of neighboring scan points. A history of scan returns were maintained that effectively mapped the surface directly in front of the vehicle. Detection thresholds were set so obstacles below a specific height would not be detected as an obstacle. The minimum obstacle height was set based on the capability of the vehicle. A convex hull algorithm was used to define the outermost edge of the obstacle.

Negative road edge detection followed a similar approach as for the positive obstacle detection. A specific search algorithm was used to find any negative height discontinuities. Each discontinuity that was detected was further evaluated to determine if the true edge and if that discontinuity was a continuation of the previously detected edge.

5.4. Multiplane LIDAR

The IBEO ALASCA LIDAR (Lages, 2004) is a four-plane scanner that is used for positive obstacle detection. The LIDAR is mounted level in the front bumper (see Figure 12) and has two planes that scan toward the ground and two planes that scan toward the sky. With a range of 80 m and a resolution of 0.25 deg it can detect obstacles accurately at long and close range. The 170-deg scan area allows seeing obstacles around upcoming turns.

The LIDAR sends scan data via Ethernet to the LIDAR PC via a TCP connection. An algorithm then transforms the raw scan data into obstacles by looking for large positive slopes in the scan data (see Figure 13).




Figure 13. iVMS graphical user interface depicting IBEO obstacles as blue lines. In the figure the vehicle is located at waypoint 162 with iVMS calculated microwaypoints extending ahead of the vehicle.

The IBEO obstacle detection algorithm followed a similar approach as the SICK obstacle detection algorithms. The scan returns were translated from the sensor coordinate frame to a local level coordinate frame. The scan returns were then compared to neighboring returns and discontinuities were detected. Since the IBEO LIDAR was a multiplanes, the obstacle detection algorithm was able to compare the scans from each plane to aid in the obstacle detection.

The SICK and IBEO LIDARs were configured so the data provided by the sensors were complementary. The SICK LIDARs were configured to detect obstacles from 0 to 20 m in front of the vehicle. The IBEO LIDAR was configured to detect obstacles from 20 to 60 m in front of the truck. Obstacles detected from both sensors were put into the obstacle database and used for path planning.

Field testing indicated that, although the two LIDAR sensors provided similar data, the characteristics of the data were somewhat different. For example, the SICK LIDAR data were very consistent

but more susceptible to reflections than the IBEO LIDAR. The IBEO LIDAR provided more accurate data but would occasionally return spurious data spikes. Putting both sets of data into the database without prior filter resulted in multiple copies of some obstacles database. A fusion of the LIDAR sensors to eliminate the multiple copies reduced database utilization and sped up the execution of the collision detection algorithms.

5.5. Trinocular Vision System

The vision system is based on multistereoscopic vision (forward looking trinocular system). It consists of three identical cameras mounted on a rigid bar on top of the hood. The two lateral cameras lay at a distance, which is about 1.5 m, while the central one is placed asymmetrical at about 0.5 m from the right one. Thanks to a precise calibration of the cameras—performed on a graduated grid—the three degrees of freedom specifying cameras orientation are fixed to known values, and in particular—in order to ease and speed-up the subsequent processing—the yaw and roll angles are fixed to zero for all cameras. The pitch angle is chosen so that the cameras frame a small portion over the horizon (to limit direct sunlight) and frames the terrain at about 4 m from the vehicle.

The trinocular system sends three video streams at 10 Hz (640×480 , color with Bayer pattern) to the vision personal computer (PC) via a firewire connection. The PC selects which stereo pair to use depending on the speed of the vehicle. Since the base line of the stereo vision system influences the depth of view, the large base line is used at high vehicle speeds so that a deeper field of view is obtained, the medium one at medium speeds, and the short base line is used at low speeds. This is one of the very few examples of very large base line stereo systems (1.5 m) used on rough off-road terrain and delivering a robust environmental perception at more than 50 m, regardless of terrain slope.

The rationale behind the design is mainly the need for mechanical robustness: three nonmoving cameras have been preferred with respect to a pan-tilt solution such as the one used by other teams, for example the Red Team (Whittaker, 2005). A few other considerations were the basis for this choice: vision must be able to sense obstacles at large distances (more than 50 m away on rough terrain), therefore a stereo vision system was the only choice.

Furthermore, the base line (distance between the stereo cameras) had to be large enough to guarantee depth perception at large distances. Systems based on moving cameras when both stereo and a large base line are used, are subject to a number of mechanical problems such as—for example—the non-negligible momentum caused by vehicle vibrations which has to be compensated for. As a result, the experience with multiple cameras providing different video streams to choose from turned out to be a winning solution, capable of removing most of the mechanical problems of a gazing system.

Vision provides sensing for both obstacle detection and path detection (see Figures 14 and 15).

1. Image disparity is first used to estimate the average terrain slope in front of the vehicle (Labayrade, Aubert & Tarel, 2002). Slope information is then used for both obstacle detection and path detection. Any significant deviation from the average smooth slope detected previously is then identified as an obstacle. The exact location of obstacles is then obtained via stereo triangulation between the two views of the object. A fairly precise localization is obtained, but nonetheless it can be further refined via sensor fusion with raw data coming from the multiplane LIDAR. In this way it is possible to detect thin vertical posts and fence poles. The system is able to detect even small obstacles (Broggi, Caraffi, Fedriga & Grisleri, 2005), but—due to both the size and capabilities of the vehicle and to the team strategy—it was tuned with very high thresholds, so that the number of false positives was reduced to a minimum. In other words, the capability of detecting small obstacle was traded for a higher robustness of the detection. Nevertheless, the system was demonstrated to be able to detect small construction cones used during both the tests and the qualification runs. Anyway, since the vehicle is able to negotiate 60 cm steps, obstacles smaller than 60 cm needs to be detected primarily for speed management issues.
2. Image disparity is also used to compute the area in front of the vehicle which features a smooth slope, the so-called free-space. The free-space is one of the features that concur to construct a representation of the path to be



Figure 14. Images showing left and right images of different situations; on the right images, colors show the presence of detected obstacles: different colors mean different distances. The additional horizontal lines represent the 5 m, 50 m, and horizon position. Posts and thin poles are correctly detected.

followed by the vehicle: also similarity in texture, similarity in color, and shape information are taken into account, fused together, and delivered to the following path planning module. Free space is obtained using a standard image warping (Bertozzi, Broggi & Fascioli, 1998) in order to localize deviations from a smooth road surface: Figure 15 shows the right image (a), the warped left image (b), and—in green—the matching cluster repre-




Figure 15. Image showing different steps of path detection: first the free-space is determined then the path is localized. Right image (a), warped left image (b), free space (c), the final result of path detection (d).

senting free space (c). Figure 15(d) shows the final result of path detection. This algorithm also signals the presence of a straight segment of road, in order to increase vehicle speed. When a curved path is present (the red dot at the top right of Figure 15 shows the presence of a nonstraight road), vehicle speed is reduced.

Vibrations are automatically filtered out since the slope detection algorithm (Broggi et al., 2005), which is the first to be performed, also extracts information that is used to electronically stabilize the oncoming images. Different light levels are compensated for by an automatic gain control scheme, which allows to sense the environment even with direct sunlight into the camera. Figure 16 shows some examples of the custom gain control mechanism.

The camera boxes have a sun shade aimed at reducing to a minimum the quantity of direct sunlight hitting the cover glass, in order to avoid over saturation and reflections due to dirty glass.

The MTRV is a production vehicle with thousands of units produced and in service with the US




Figure 16. Images captured at sunrise, representing the view with (a) and without (b) the developed camera gain control scheme; (c) and (d) show the result of obstacle detection in bad illumination conditions: although a part of the images are oversaturated, the terrain is visible and the algorithm can be run; obstacles can be detected until they enter the oversaturated area.

Marine Corps, US Navy, and other services throughout the world. Performance of the MTRV was not specifically tracked as part of the TerraMax development efforts, rather the dynamic capabilities of the



Figure 17. Some phases of the qualification runs.

vehicle were identified through ADAMS simulations and used to define the limits within which the real time path planner developed alternative paths.

6. VEHICLE PERFORMANCE

The vehicle was able to conclude the qualification runs with excellent results, avoiding obstacles, passing into tunnels, and maneuvering (with backups) in order to align itself with narrow barriers and gates. Figure 17 shows some pictures taken during the qualification phase.

The iVMS development philosophy was to create an autonomous system that could, in the future, be

utilized in military operations. This allowed for a more rugged implementation of the iVMS for real time navigation across unknown terrain. As a result, Team TerraMax was one of only five teams to traverse the 132-mile course and the only vehicle to overcome an overnight “pause” of the autonomous system. During the race, TerraMax reached a maximum speed of 42 mph. This is impressive not only due to the size and weight of the TerraMax, but due to the fact that true obstacle avoidance was achieved at these speeds. Figure 18 shows a few pictures taken during the last part of the race.

During the race the TerraMax was paused 13 times by DARPA officials to maintain a minimum distance between the competing vehicles or for passing



Figure 18. Two pictures taken during the race (courtesy of DARPA).

stopped vehicles. The TerraMax automatically stopped and realigned its path approximately 52 times during the race. The majority of the path resets occurred while traversing beer bottle pass where the road was very narrow and the turns were very tight compared to the size of the vehicle (see Figure 18).

An automatic reversion mechanism was implemented to manage redundant sensors. The reversionary logic was activated due to sensor abnormality twice during the race. The reversionary logic correctly selected the operational sensor and continued to operate normally after the reversion.

A software application health monitor was implemented to monitor the health of the system and start and stop applications as necessary to keep the system executing normally. During the race the health monitor function was activated and correctly reset applications to keep the system operating normally with only minor interruptions in service. A pe-

riodic database integrity check was also performed to prevent fatal errors from corrupting the database and to recover data if a data error was found.

During the race the TerraMax struck the edge of one of the concrete underpass barriers. The impact of the tunnel caused the IBEO LIDAR and passenger vision camera to be severely misaligned. The misalignment caused a georegistration error of the detected obstacles. This caused the path planner to offset the path to compensate for obstacles, slowing the progress of the vehicle for the last half of the race. TerraMax completed the 132-mile race with an official time of 12 h, 51 min and over 28 h of continuous operation.

7. CONCLUSIONS

This technology demonstrated by TerraMax has the potential of improving soldier survivability in the battlefield by removing soldiers from harms way, especially during convoy operations—the ultimate outcome of the Congressional goal. The development of fully autonomous systems has allowed Oshkosh and its partners to fully understand the requirements related to both leader-follower and autonomous operation. The opportunity exists to develop and deploy this technology to allow for robotic replacement of convoy personnel, allowing the personnel to be refocused on more pressing duties, and ultimately reduce the convoy personnel exposure to enemy threats.

On the technical side, the choice of the (i) sensors suite delivered the sufficient amount of information for the successful conclusion of the race and demonstrated to be robust enough to deal with the extreme conditions of a desert environment in summer. The experience with multiple cameras providing different video streams to choose from turned out to be a winning solution, capable of removing most of the mechanical problems of a gazing system. The 28 h of uninterrupted service provided by the (ii) processing systems and software architecture demonstrated the stability and robustness. Finally, the (iii) algorithmic solutions proved to be fast and reliable, reducing the number of wrong detections to a minimum.

The TerraMax partners of Oshkosh Truck, Rockwell Collins, and University of Parma have demonstrated the scalability and portability of the autonomous technology by installing and operating the system on an Oshkosh palletized loading system (PLS). The PLS is a 10 × 10 vehicle with a gross vehicle



Figure 19. PLS operating autonomously near Barstow, CA.

weight of 84 000 lb and is capable of delivering a 33 000 lb payload. The vehicle was successfully demonstrated at the Yuma Test Center in January 2006, exhibiting the same autonomous capabilities as the TerraMax. The project was completed in approximately 75 days. Figure 19 shows the PLS vehicle during an autonomous run.

REFERENCES

Bertozzi, M., Broggi, A., & Fascioli, A. (1998). Stereo inverse perspective mapping: theory and applications.

- Image and Vision Computing Journal, 16(8), 585–590.
- Braid, D., & Johnson, C.W. (2006, June). DARPA Grand Challenge intelligent vehicle management system (iVMS) and the transition to leader/follower. 6th Annual Intelligent Vehicle Systems Symposium & Exhibition, Traverse City, MI (pp. 104–113).
- Broggi, A., Caraffi, C., Fedriga, R.I., & Grisleri, P. (2005, June). Obstacle detection with stereo vision for off-road vehicle navigation. 2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition, San Diego, CA (Vol. 3, pp. 65–65).
- Kuffner, J., & LaValle, S. (2000, April). RRT-connect: An efficient approach to single-query path planning. Proceedings 2000 IEEE International Conference on Robotics and Automation, San Francisco (Vol. 2, pp. 995–1001).
- Labayrade, R., Aubert, D., & Tarel, J.-P. (2002, June). Real time obstacle detection in stereovision on non flat road geometry through “v-disparity” representation. Intelligent Vehicle Symposium, 2002, IEEE, Paris, France (Vol. 2, pp. 646–651).
- Lages, U. (2004, June). Laser sensor technologies for prevent safety functions, ATA EL 2004, Parma, Italy.
- Ozguner, U., Redmill, K.A., & Broggi, A. (2004, June). Team TerraMax and the DARPA Grand Challenge: A general overview. Intelligent Vehicles Symposium, 2004 IEEE, Parma, Italy (pp. 232–237).
- Rosenblatt, J. (1997). DAMN: A distributed architecture for mobile navigation. Doctoral dissertation (Tech. Rep. CMU-RI-TR-97-01). Pittsburgh, PA: Carnegie Mellon University, Robotics Institute.
- Whittaker, W. (2005, August). The Red Team technical paper. DARPA for the Grand Challenge 2005.