CIGR - International Conference of Agricultural Engineering
XXXVII Congresso Brasileiro de Engenharia Agrícola

Brazil, August 31 to September 4, 2008

EXPERIENCES FROM THE 2007 DARPA URBAN CHALLENGE AND IMPLICATIONS FOR HEAVY VEHICLE AUTONOMY

CARL D. CRANE; DAVID ARMSTRONG; ANTONIO ARROYO; ANTOIN BAKER; DOUG DANKEL; GREG GARCIA; NICHOLAS JOHNSON; JAESANG LEE; SHANNON RIDGEWAY; JOHN K. SCHUELLER; ERIC SCHWARTZ; ERIC THORN; STEVE VELAT; JI HYUN YOON[1]

[1] Center for Intelligent Machines and Robotics (CIMAR), University of Florida. corresponding e-mails: ccrane@ufl.edu; schuejk@ufl.edu

Presented at
CIGR INTERNATIONAL CONFERENCE OF AGRICULTURAL ENGINEERING
XXXVII CONGRESSO BRASILEIRO DE ENGENHARIA AGRÍCOLA – CONBEA 2008
Brazil, August 31 to September 4, 2008

ABSTRACT: Team Gator Nation modified a hybrid Toyota Highlander by automating it and adding pose estimation and object detection sensors. A control architecture was developed which integrated planning, perception, decision making, and control elements. The developed Adaptive Planning Framework was utilized for situation assessment, behavior mode evaluation, and behavior selection and execution. The system successfully detected and modeled its environment and then planned and executed appropriate actions in real time based upon an integration of a priori and sensed information. The various technologies and techniques, such as the JAUS architecture, raster local world model, receding horizon controller, perception technologies, and operating behavior modes, which were successful for this vehicle and its situations are applicable to heavy vehicles, such as autonomous agricultural, construction, and mining equipment.

KEYWORDS: autonomous, Urban Challenge, JAUS, perception, control

INTRODUCTION: In order to inspire and focus the creativity of robotics researchers, the Defense Advanced Research Projects Agency (DARPA), the central research and development organization for the United States Department of Defense, has sponsored competitions for totally autonomous vehicles. The first two Grand Challenges required the autonomous vehicles to complete long travels over difficult desert terrain. No vehicle was successful in 2004. Five vehicles were able to complete 2005's 132-mile course with the Stanford Racing Team winning the US$2 million prize with the a time of 6 hours, 53 minutes.

The Grand Challenges were in an off-road environment with limited traffic interactions. For 2007, DARPA proposed an Urban Challenge which "features autonomous ground vehicles maneuvering in a mock city environment, executing simulated military supply missions while merging into moving traffic, navigating traffic circules, negotiating busy intersections, and avoiding obstacles." This added significant structure and complexity, since there was now substantial interaction with the environment beyond simple transversal.

Team Gator Nation had acquitted itself well with vehicles it developed for the two Grand Challenges (Crane, et al., 2006; Touchton, et al, 2006). Although the vehicles did not complete the entire courses, they passed all the screenings and were named finalists. The vehicles could follow planned paths and detect and avoid obstacles.

But the Urban Challenge was more complex. The following technical challenges were identified: pavement and lane detection; detection of static obstacles, detection and classification of dynamic objects; environment data representation and sensor integration with noise in sensor systems; localization; high-level mission planning; determination of appropriate behavior mode and smooth transition between modes; interprocess communication and coordination of multiple threads on multiple computers; and fault tolerance.

VEHICLE AND SYSTEM OVERVIEW: The steering, throttle, transmission, and braking controls on a hybrid Toyota Highlander were automated and vision, ladar, inertial, and GPS sensors were mounted. The computer system architecture is a natural extension of the Joint Architecture for Unmanned Systems (JAUS) Reference Architecture, Version 3.2, which defines a set of reusable components and their interfaces. Figure 1 depicts the four basic elements (Planning, Control, Perception, and Intelligence) of the system.

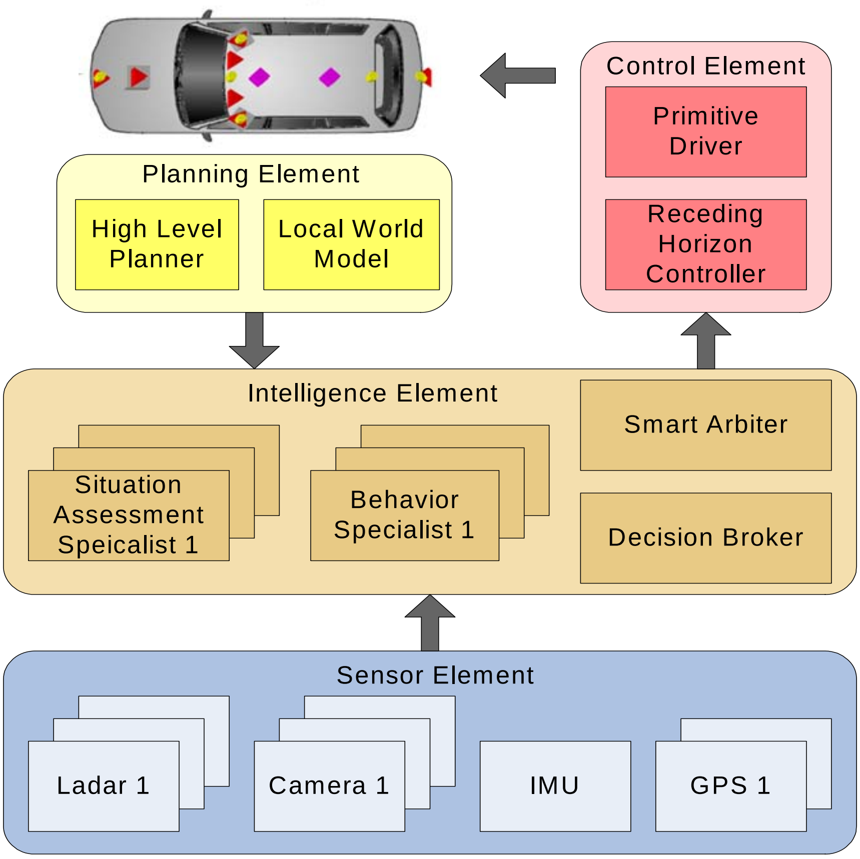

Figure 1: Four basic elements of computer system.

A typical sequence of operations of the system would be to:
- off-line path plan to generate a desired path based upon the route network and the mission,
- generate a 300m x 300m grid local world model (LWM) with 0.5 m resolution,
- integrate data on obstacles, dynamic objects, terrain and road lanes into the LWM,
- situation specialists make specific findings (e.g., a road lane clear of obstacles),
- behavior specialists assess whether behavior modes are appropriate,
- decision broker selects behavior mode,
- smart arbiter generates 60m x 60 m transversability grid, and
- receding horizon controller plans suitable path and generates steering, throttle and braking commands.

This sequence of activities guides the vehicle using both a priori roadway data as well as static and dynamic obstacle and lane information from the perception system. The LWM constantly estimates discrepancies and calculates a net offset that can be applied to to the a priori data in the event that sensed data is momentarily lost. It determines important facts about the vehicle's position in the world in order to help make decisions on the appropriateness of certain behaviors.

SENSOR ELEMENT COMPONENTS: Geolocalization is achieved using a GE Aviation North-Finding-Module (NFM) combined with two GPS units and one odometer. The NFM is an inertial navigation system that maintains Kalman filter estimates of the vehicle's global position and orientation as well as angular and linear velocities. An on-board computer determines whether the Novatel Propak V3-HP with Omnistar subscription or a Garmin WAAS Enabled GPS 16 is providing the best signal and the NFM takes the best signal with a corresponding tuning. Using inertial sensors and encoder signals, the NFM was even able to maintain accuracies of less than five meters when travelling up to three miles without GPS.

The sensor package on the vehicle included six SICK LMS-291 LADARs, two SICK LD-LRS1000 long range LADARs, and six Matrix Vision BlueFox high-speed USB2.0 color cameras. Many of the deployed sensors were also electromechanically articulated on the vehicle with one degree of freedom. Figure 2 shows two views of the sensor package on the vehicle.

The first major sensor thrust was to characterize terrain by using a combination of vision and LADAR to look at slope, relative height, and texture regularity of the terrain around the vehicle. The second thrust was to localize static and moving obstacles. The long-range LADARs mounted on the front fenders provided obstacle data out to 275m. This was supplemented by close-range sensors on articulated mounts on the front and rear bumpers. The last thrust determined road characterization and vehicle pose within the lane high-frame-rate vision images from the center of the bridge and on the wings, supplemented by vertical-fan LADARs.

The various sensor inputs were unified by using a generated format in which a transverability grid was used to allow information to be added and fused by an arbitration component. Figure 3 depicts three example transversability grids and the result of sensor fusion. This common framework allowed the developers to work independently and the components to be asynchronous. Many of aspects of sensor development, LWM, and path planning for Urban Challenge were developed and improved during the previous two Grand Challenge competitions. Crane, et al., (2006) and others have reported on vehicle design, system localization, perception sensors, and dynamic planning algorithms.

The largest deviation from the previous competitions was the addition of manned and unmanned moving traffic. Moving obstacles were found by examining LADAR point clouds and finding clusters that were related. The velocity state of every obstacle in sensing range was localized, tracked, and reported to the LWM.

INTELLIGENCE: The sensors generated a perceived operating environment. But it was then up to the adaptive planning framework to select the most appropriate behavior characteristics. The framework first used specialists to assess the situation. Findings were generated in the form of conditions, states, or events. For example, a software component could decide whether a move to an adjacent lane was safe based upon sensor data.

Behavior specialists monitor findings and evaluate the suitability of a behavior. For example, the Pass Left/Right behavior specialist monitors the travel lane and adjacent lanes for obstacles and recommends lane changes.

The decision broker monitors all the behavior specialists and assumes ultimate authority over how the vehicle will operate. The seven behavior modes which were used were roadway navigation, open area navigation, pass left/right, reverse direction, intersection traversal, off-road, and parking.

The smart arbiter obtained imputs from the terrain smart sensor, the lane finding smart sensor, the path finding smart sensor, and the LWM and built a 60m x 60m transversability grid centered at the vehicle's current position based upon the current behavior mode of the system.

(a)

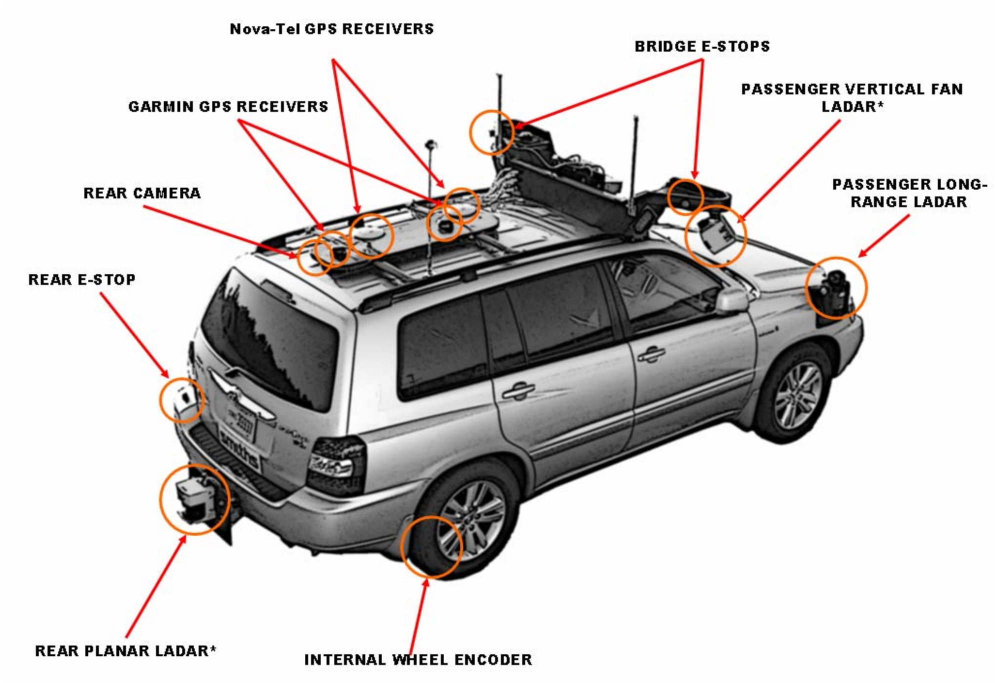

(b)

Figure 2: Vehicle sensor package

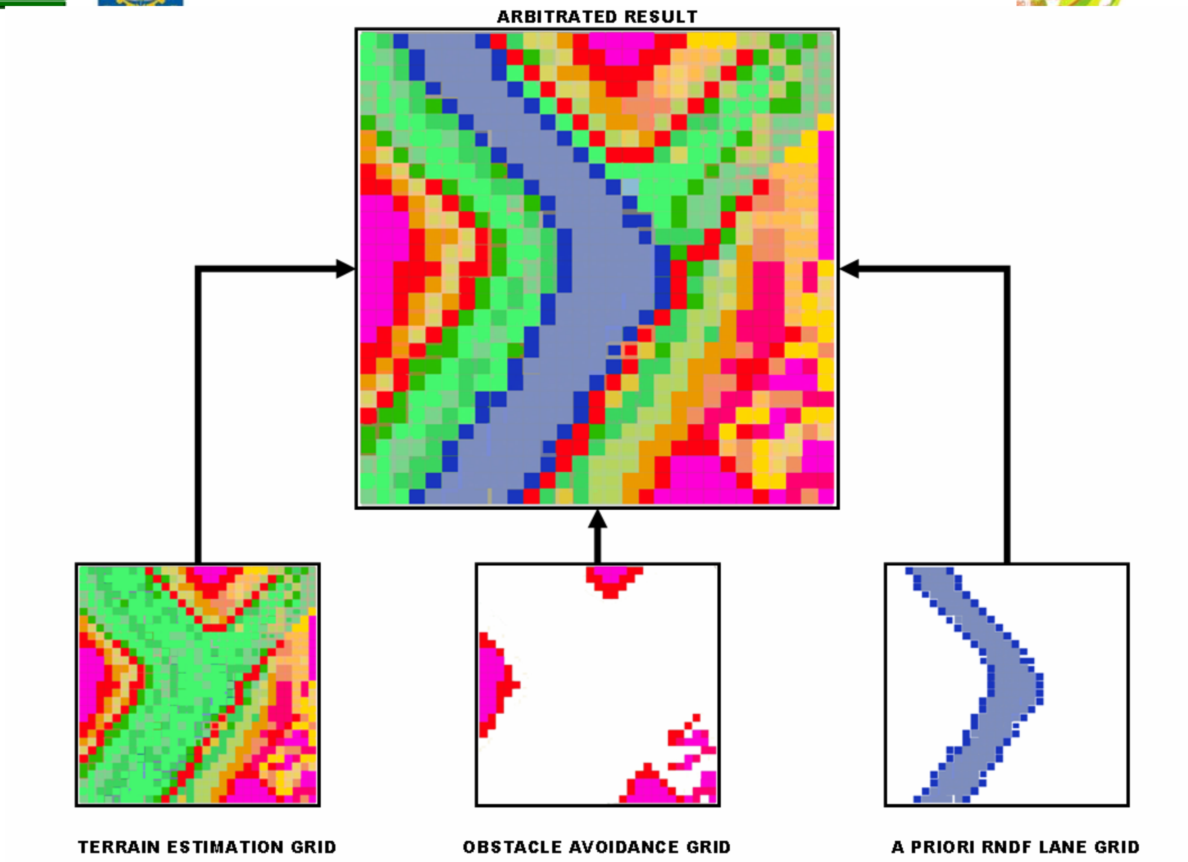

Figure 3: Sensor fusion for transversability grid

CONTROLLER: Low level path planning uses a modified form of model predictive control to find a optimal trajectory which minimizes a cost function of the form $J = \Phi(x(t_f), t_f) + \int_{t_0}^{t_f} g(x(t), u(t), t)dt$.

Goal points are calculated from a list of path segments constructed by the high level planner and the LWM. A vehicle kinematic model is used and closed-loop steering control with position and actuation feedback is repeated at 40 Hz. The PID speed controller is limited by such items as intersection proximity, static or moving obstacles, path curvature, and transversability. The output has the form

$$linear\_effort = K_{FF}s + Bias_{FF} + K_p e(t) + K_i \int_0^t e(\tau)d\tau + K_d \frac{de}{dt}$$

which is looped at 40 Hz with velocity state sensor feedback.

The primitive driver closes the loop between the receding horizon controller and the vehicle. Steering and shifter commands are converted to absolute positions for the PID-tuned motors connected to the steering column and the shifting mechanism. Throttle and brake wrench commands are converted to voltages fed to the throttle and brake drive-by-wire electronic control units.

RESULTS AND LESSONS LEARNED: The adaptive planning framework correctly managed the system's behavior with respect to the sensed scenarios. Low level control was maintained during imposed behaviors leading to continuous driving behavior.

The qualification event for the Urban Challenge consisted of missions planned on three different courses. The vehicle ran on all three courses with some success, but not with total success. On course A, the vehicle had difficulty detecting vehicles partially occluded by other vehicles, leading to incorrect assumption of right-of-way at a two way traffic circle. On course B, a transition from an open area to a road network resulted in an ill-conditioned optimization problem. On both course C (which tested intersection precedence and re-planning) and course A, ground strikes from LADAR sensors were detected as fixed objects. The grid resolution of 0.5 m proved problematic on lanes less than 5 m wide due to lack of precision in lane representation.

Most of these problems involved component implementation errors which should be solvable with additional development. The overall architecture worked as designed. The system was able to detect and model its environment and then plan and execute appropriate actions in real time. The central concept is the integration of a priori and sensed information in a raster format in the local world model. Based upon this information, an appropriate behavior is selected via arbitration. The vehicle behavior is executed by generation of a navigation grid coupled with metadata.

IMPLICATIONS FOR AUTONOMOUS OFF-ROAD VEHICLES: At first glance, it might seem that the Grand Challenge events with vehicles running relatively free across the desert might be more applicable to off-road applications such as agriculture, mining, and construction. But those off-road practical applications will demand vehicles which can follow precise paths, perform detailed missions, and accommodate many static and moving obstacles. The technologies demonstrated in the Urban Challenge competition are very relevant.

The architecture used on the NaviGator vehicle should be applicable to off-road vehicles. Its extensibility will accommodate the operational behaviors required of those vehicles. At a minimum, the JAUS architecture should be used to take advantage of the advancements in other vehicle applications.

The LADAR and camera sensor systems should be applicable to off-road vehicles, along with the perception algorithms and systems. Concepts such as the local world model, transversability grid, and receeding horizon control should also work well in the off-road environment. Overall, the NaviGator Urban Challenge vehicle showed both the successes and challenges of autonomous vehicle control. Similar results may be achievable with off-road vehicles.

REFERENCES:

CRANE, C.D.; ARMSTRONG III, D.G.; TOUCHTON, R.; GALLUZZO, T.; SOLANKI, S.; LEE, J.; KENT, D.; AHMED, M.; MONTANE, R.; RIDGEWAY, S.; VELAT, S.; GARCIA, G.; GRIFFIS, M.; GRAY, S.; WASHBURN, J.; ROUTSON, R. Team CIMAR's NaviGator: An unmanned ground vehicle for the 2005 DARPA grand challenge. Journal of Field Robotics. 23(8):599-623, 2006

TOUCHTON, R.; GALLUZO, T.; KENT, D.; CRANE, C.D. Perception and planning architecture for autonomous ground vehicles. IEEE Computer. 40-47, December 2006.