Lessons Learned at the DARPA Urban Challenge

Carl Crane, David Armstrong, Antonio Arroyo, Antoin Baker, Doug Dankel,
Greg Garcia, Nicholas Johnson, Jaesang Lee, Shannon Ridgeway,
Eric Schwartz, Eric Thorn, Steve Velat, and Ji Hyun Yoon
University of Florida
Gainesville, FL 32611
352-392-9461
ccrane@ufl.edu

ABSTRACT
This paper describes the intelligence components associated with the system design developed for Team Gator Nation's submission to the 2007 DARPA Urban Challenge. In this event, vehicles had to navigate on city streets while obeying basic traffic laws. One of the major challenges was interacting with other vehicles such as at intersections. To address these challenges, a hybrid Toyota Highlander was automated and instrumented with pose estimation (GPS and inertial) and object detection (vision and ladar) sensors. A control architecture was developed which integrates planning, perception, decision making, and control elements. The intelligence element implements the Adaptive Planning Framework which was developed by researchers at the University of Florida. This framework provides a means for situation assessment, behavior mode evaluation, and behavior selection and execution. The paper describes this architecture and concludes with lessons learned from participation in the Urban Challenge event.

Keywords
autonomous ground vehicle navigation

1. INTRODUCTION
In DARPA's vision, "The Urban Challenge features autonomous ground vehicles maneuvering in a mock city environment, executing simulated military supply missions while merging into moving traffic, navigating traffic circles, negotiating busy intersections, and avoiding obstacles." Moving the challenge into an urban setting adds structure and complexity to the Grand Challenge problem. Previous success relied on a single mode of operation, without interaction with the environment beyond simple traversal. Success in the Urban Challenge will require numerous modes of operation and complex interaction with the environment. It is expected that the urban environment will also hamper the use of GPS for localization, further complicating the challenge.

The specific problem to be solved is detailed in the Urban Challenge Technical Evaluation Criteria document [1]. Here the problem is organized into four categories, i.e. Basic Navigation, Basic Traffic, Advanced Navigation, and Advanced Traffic, each of which is more complex than the previous. Upon reviewing this document, the authors identified the following set of technical challenges:

1. pavement (road) detection and lane detection
2. detection of static obstacles
3. detection and classification of dynamic objects
4. environment data representation and sensor integration
5. localization
6. reconciliation of differences in estimated global pose, a priori data, and sensed information
7. high level mission planning
8. determination of appropriate behavior mode
9. smooth transition of vehicle control between behavior modes
10. interprocess communication and coordination of multiple threads on multiple computers
11. fault tolerance

This paper documents some of the design choices that have been made to address these challenges with emphasis placed on items 8 and 9.

Much work has been done in the past twenty years to address many of the specific technical challenges listed in the previous section. Several references [2]-[7] provide excellent summaries of the advancements made by other teams competing in the 2005 DARPA Grand Challenge. The authors' work related to the 2005 event is published in two references [8]-[9]. Numerous additional references can be cited for each of the important technical challenges.

2. SYSTEM ARCHITECTURE
A hybrid Toyota Highlander was selected as the base platform for the system. Steering, throttle, braking, and transmission controls

Figure 1. NaviGATOR

were automated and vision, ladar, inertial, and GPS sensors were mounted to provide necessary information about the environment. The vehicle system is shown in Figure 1.

The 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. The actual core software to support the JAUS messaging system was developed and extensively tested for the previous Grand Challenge and supports the current effort with little or no modification required.

At the highest level, the architecture consists of four basic elements, which are depicted in Figure 2. The Planning Element contains the components that act as a repository for a priori data such as the Route Network Definition File (RNDF) which provides the overall database information about the roads, lanes, and intersections, and the Mission Data File (MDF) which provides the set of RNDF waypoints to traverse for a particular mission. This element also performs the high level route planning and re-planning based on that data plus real-time information provided by the rest of the system. The Control Element contains the Primitive Driver that performs closed-loop control on vehicle actuators to keep the vehicle on a specified path. The Perception Element contains the components that perform the sensing tasks required to determine the vehicle's position, to find a road, to find the lanes on a paved road, to locate both static and dynamic obstacles, and to evaluate the smoothness of terrain. Finally, the Intelligence Element contains the components that work together to determine the best course of action to navigate the vehicle in a complex environment based on the current mission and situation. An overview of a typical sequence of operations of the architecture is presented as follows (reference Figure 2):

(1) The High Level Planner component performs off-line path planning to generate a desired motion path based on the Route Network Definition File (RNDF) and the Mission Data File (MDF).

(2) A tessellated Local World Model (LWM) (300m × 300m grid with 0.5m resolution) is generated based on a priori road network data and the planned motion path. The center point of the LWM is located at the current location of the vehicle as determined from sensor positioning data.

(3) Data from ladar and vision sensors, which identify static obstacles, dynamic objects, smooth terrain, and road lane regions, is integrated as a layer into the LWM.

(4) Based on the a priori data and sensed data stored in the LWM, software components referred to as Situation Assessment Specialists focus on making specific findings (one simple example is the specialist that reports if the lane to the left, or right, is clear of other vehicles or obstacles).

(5) Seven software components referred to as Behavior Specialists then make an assessment of whether their corresponding behavior mode is appropriate at this moment. The six behavior modes are Roadway Navigation, Open Area Navigation, Pass Left and Pass Right, Reverse Direction, Intersection Traversal, Off Road, and Parking.

(6) A software component referred to as the Decision Broker selects the behavior mode for the system based on the recommendations of the Behavior Specialists.

(7) Based on the behavior mode, a software component called the Smart Arbiter then generates a 60m × 60m traversability grid that is formed to elicit a specific response from the vehicle (change lanes is an example).

(8) Finally, the Receding Horizon Controller component plans a suitable path through the grid that was output by the Smart Arbiter. Steering, throttle, and braking commands are generated to execute the planned path.

A description of the components associated with the Intelligence Element follows.

3. INTELLIGENCE COMPONENTS

Team Gator Nation has developed and deployed the Adaptive Planning Framework [19] to address the issues associated with behavior mode selection in complex or unstructured environments presented during the DARPA Urban Challenge. It enables the vehicle to intelligently select the most appropriate behavioral characteristics given the perceived operating environment. The framework is scalable to systems of varying complexity and size and is compatible with existing architectures such as JAUS RA-3.2, NIST 4D/RCS, and others. The Adaptive Planning Framework is composed of three principle elements tasked with assessing the situation, determining the suitability and viability of all possible solutions, and executing the most suitable of all recommended solutions. These three component types are

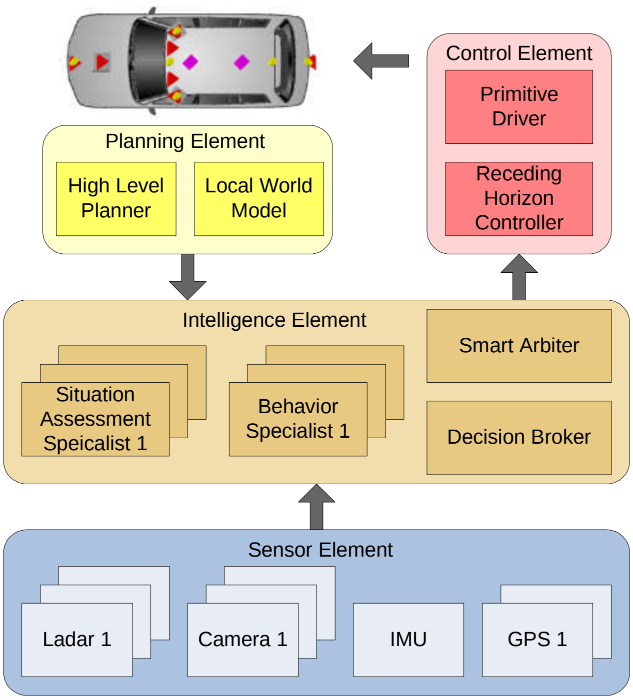

Figure 2. System Architecture

multiple Situation Assessment components, multiple Behavior Specialists, and one Decision Broker component.

3.1 Situation Assessment Specialist Components

Dynamic environment information, originating from any array of sensors is monitored and managed by the Situation Assessment Specialists. Each specialist design is tailored to the sensor or collection of sensors whose data it will be analyzing. While the inputs to the specialist can come from any data source, the output or "finding" must adhere to specific guidelines outlined by the framework. Findings can be in the form of conditions, state, or events. Once the findings have been generated the information is disseminated to all other components that might need it. An example of a situation assessment specialist would be a software component whose sole function was to determine if it is safe to move to the adjacent lane. This component would monitor sensor data as reported by the Moving Objects sensor and reach a Boolean conclusion that would be stored as metadata for use by other processes.

3.2 Behavior Specialist Components

The findings rendered by the Situation Assessment Specialists are consumed by the Behavior Specialists. There is a one-to-one mapping of each behavior with a Behavior Specialist. The role of this specialist is to monitor the findings and evaluate the suitability of its behavior under the current perceived operating conditions. An example of a behavior specialist is the Pass Left/Right behavior specialist. This specialist simultaneously monitors the desired travel lane for obstructions as well as adjacent travel lanes. Based on all the inputs the specialist recommends whether or not a lane change is an appropriate and safe option. As with the specialist findings, the default recommendation is unsuitable and must be proven appropriate at every iteration of the program to ensure truth of the results and operating safety. The Behavior Specialists do not possess the ability to activate or deactivate their associated behavior; such authority is only given to the Decision Broker.

3.3 Decision Broker Components

At the highest level of the framework lies the Decision Broker. Its role is to monitor all Behavior specialist recommendations. It assumes ultimate authority over how the Urban NaviGator will operate while in autonomous mode. Like the other entities within the framework the Decision Broker can base its conclusions on not only the recommendations and findings of other specialists, but it may also look at data from any other pertinent source. Team Gator Nation's implementation of the Adaptive Planning Framework centralizes all the Decision Broker functionality within the JAUS Subsystem Commander and has the added responsibility of selecting which component receives control of the vehicle's JAUS Primitive Driver. The framework architecture employs an asynchronous, iterative, forward chaining reasoning approach to decision making.

3.4 Behaviors Used during the Urban Challenge

The Urban NaviGator is programmed with seven operating behavior modes where each behavior is comprised of a series of sub-behavior modes. Some sub-behaviors may be optional, depending on the mission plan or ambient conditions. Vehicle performance is denoted by a sub-behavior status indicator. A failure protocol is incorporated into each sub-behavior should sufficient environmental changes warrant the current vehicle operation inappropriate or unsafe. In most cases the vehicle is able to recover to a default safe operational state. However, in some cases, such as a catastrophic system failure or an excessively hostile environment, the safest course of action is for the vehicle to pause and wait for more favorable conditions. The corresponding behavior specialist constantly evaluates the appropriateness of its behavior mode and the decision broker determines which mode will have operation of the vehicle. The seven behavior modes are described subsequently.

3.4.1 Roadway Navigation

The Roadway Navigation behavior is the primary driving behavior deriving commands to be sent to the vehicle actuators while the objective is lane following. This behavior will allow the vehicle to navigate the roadway within the lines of its desired lane. The default sub-behavior is to maintain a safe following distance behind any vehicles ahead. Other sub-behaviors include lane changes on a multi-lane road in order to pass through a mission goal point.

3.4.2 Open Area Navigation

Open area navigation is a behavior that should only be needed in special circumstances during the Urban Challenge event. This behavior allows the vehicle to move towards a goal location without striking any object, while avoiding any rough terrain. This is in effect the only behavior mode that was required in the 2005 DARPA Urban Challenge. It will be useful in the Urban Challenge when the vehicle is in an open area such as a parking lot or an obstacle field. The associated sub-behaviors are Enter Open Area, Exit Open Area, Enter Parking Space, and Exit Parking Space. Thus if the mission plan for the open area does not contain parking spaces the system would transition from Enter Open Area to Exit Open Area sub-behavior.

3.4.3 Pass Left and Pass Right

The Pass Left and Pass Right maneuvers will be used in passing situations where a static obstruction impedes progress in the desired lane but there exists an adjacent available travel lane. Successful Pass Left Behavior execution entails a Lane Change Left sub-behavior, Passing Vehicle sub-behavior, and Lane Change Right sub-behavior. This behavior implies a momentary lane change for obstacle avoidance.

3.4.4 Reverse Direction

This behavior is called whenever it is determined that the current lane is blocked and there is no alternate clear lane available for passing. It will also be applicable in cases where the vehicle has entered a 'dead end' road that it must 'escape' to reach a mission goal point. The default sub-behavior is to execute an N-point turn sub-behavior protocol.

3.4.5 Intersection Traversal

The intersection traversal behavior will be applicable when the vehicle enters the vicinity of an intersection. This is one of the most complicated behavior modes in that the system must rely on a series of situation assessment specialists to safely navigate the intersection. This behavior mode must handle queuing, stopping at the stop line, determining right of way, and ultimately traveling through the intersection while avoiding other vehicles. These steps are compartmentalized into five sub-behaviors: Queue to Intersection, Stop at Intersection, Queue Turn, Clear Intersection, and finally Traverse Intersection. It should be noted that if there is no stop at the intersection the sub-behavior will transition from Queue to Intersection to Queue Turn.

3.4.6 Off Road

This behavior is called when a sparse waypoint problem is identified or when the MDF indicates an unmarked or dirt road. The default sub-behavior is Defensive/Reflexive. In this sub-behavior the vehicle operates in a heightened state of cautiousness. The Subsystem Commander enforces more stringent speed limits based on inertial measurement sensor feedback, other perception algorithms are retuned for path finding as opposed to lane finding and line following, and sensor grid maps are arbitrated to give more freedom to the A-star vehicle path planner for reflexive obstacle avoidance.

3.4.7 Parking

This behavior must deal with the problems that arise in the parking lot scenario where precise motion is necessary. When the vehicle approaches the vicinity of an assigned parking space, precise path planning will be initiated to align the vehicle as required. Situation assessment specialists monitor the near surroundings of the vehicle to center the vehicle in its parking space while avoiding any static or dynamic objects.

3.5 Smart Arbiter Component

The purpose of the Smart Arbiter component is to generate a 60m × 60m traversability grid, centered at the vehicle's current position, which is used to implement a desired behavior. Motion execution, which is discussed in the next section, is accomplished via an A* search through this grid to determine the least cost path. In most cases, the least cost path will be obvious as the grid has been constructed to accomplish a desired action. An important feature of this entire approach is that specific behavior modes can be changed with smooth continual control of the vehicle.

The Smart Arbiter obtains inputs from the Terrain Smart Sensor, the Lane Finding Smart Sensor, the Path Finding Smart Sensor, and the Local World Model and builds its grid based on the current behavior mode of the system. For example, if the system is in the Roadway Navigation behavior, then the grid cells corresponding to the positions of the line on the edge of the lane as identified by the Lane Finding Smart Sensor will be marked as non-traversable regions in the Smart Arbiter grid. The cells corresponding to the road lane will be marked as highly traversable. This will prevent the planner from planning outside the current lane. The output grid of the Smart Arbiter is used by the Receding Horizon Controller component which plans the appropriate path for the vehicle.

4. RESULTS AND LESSONS LEARNED

The performance of the implemented architecture at the DARPA Urban Challenge was in most part satisfactory, but less than desired with respect to certain scenarios. The system performed most subtasks well, but failed to fully realize the potential of the design. The qualification event was comprised of missions planned on three courses. The vehicle ran on all three courses with some success.

The Adaptive Planning Framework correctly managed the system's behavior with respect to the sensed scenario. Low level control of the vehicle was maintained during imposed behaviors by the architecture, leading to smooth continuous driving behavior.

Course A exposed a deficiency in the persistence of moving objects in the implementation. This course simulated a two way traffic circle that the autonomous vehicle had to merge into and out of. Sometimes traffic vehicles became occluded by others, leading in these cases to the autonomous vehicle incorrectly determining it had right of way and could proceed. An attempt to tune this deficiency's effect down was not successful, mostly due to the lack of testing.

Course B exposed an error in the search methodology for the open area behavior. This test involved navigating through a large open area to a road network to complete a mission. The search space the algorithm considered for the open area was uniform in costs related to traversability, leading to an ill conditioned optimization problem. This situation was not handled, and the component controlling the vehicle became non responsive.

Courses A and C exposed a "ground strike" problem with moving object detection. Course C was designed to test intersection precedence and re-planning. In both A and C, ground strikes from the LADAR sensors were detected as fixed objects that had to be considered in the intersection and roadway navigation behaviors. These false detections lead to less than desirable behaviors for the scenarios encountered.

The Lane Correction Arbiter Smart Sensor (LCASS) concept provided accurate information concerning lane center relative to the vehicle location. The utilization of this information in a grid resolution of 0.5 m proved to be problematic. Typical lane widths encountered at the DUC site were 3+ m to 4 m. Typical lane center corrections were often smaller than the grid resolution, leading to situations where the vehicle left the lane due to the lack of precision in lane representation. These problems were observed in roadway navigation on lanes less than 5 m.

Most failures observed involved component implementation errors. The overall architecture worked as designed, given the performance of the components. The implementation does not have significant simulation capacities. Testing was performed with the system deployed in a suitable environment. The development of hardware in the loop simulation for the system could have allowed many of the shortcomings of the implementations to be identified and fixed in a shorter time.

5. CONCLUSION

The performance requirements identified in the Urban Challenge Technical Evaluation Criteria were challenging. The system had to be able to detect and model its environment and then plan and execute appropriate actions in real time.

The approach described in this paper was generated after careful consideration of the design requirements. The central concept is the integration of a priori and sensed information in a raster format in the Local World Model. Based on this information, an appropriate behavior is selected via arbitration. The behavior is executed by generation of a navigation grid coupled with metadata.

The primary new contribution of this approach is that related to solving the technical challenges of (a) the determination of the appropriate behavior mode, and (b) the smooth transition of vehicle control between behavior modes.

6. ACKNOWLEDGMENTS

The members of Team Gator Nation would like to gratefully acknowledge the sponsorship and support of GE Aviation, the Air Force Research Laboratory at Tyndall Air Force Base, and the University of Florida.

7. REFERENCES

[1] DARPA Urban Challenge Technical Evaluation Criteria, 16 Mar 2006.

[2] Urmson, C., et al., "A Robust Approach to High-Speed Navigation for Unrehearsed Dessert Terrain," *Journal of Field Robotics*, Vol. 23, Issue 8, Aug 2006.

[3] Trepagnier, P., et al., "KAT-5: Robust Systems for Autonomous Vehicle Navigation in Challenging and Unknown Terrain," *Journal of Field Robotics*, Vol. 23, Issue 8, Aug 2006.

[4] Mason, R., "The Golum Group/ University of California at Los Angeles Autonomous Ground Vehicle in the DARPA Grand Challenge," *Journal of Field Robotics*, Vol. 23, Issue 8, Aug 2006.

[5] Lakhotia, A., et al., "CajunBot: Architecture and Algorithms," *Journal of Field Robotics*, Vol. 23, Issue 8, Aug 2006.

[6] Travis, W., et al., "SciAutonics-Auburn Engineering's Low-Cost High-Speed ATV for the 2005 DARPA Grand Challenge," *Journal of Field Robotics*, Vol. 23, Issue 8, Aug 2006.

[7] Miller, I., et al., "Cornell University's 2005 DARPA Grand Challenge Entry," *Journal of Field Robotics*, Vol. 23, Issue 8, Aug 2006.

[8] Crane, C., et al., "Team CIMAR's NaviGATOR: An Unmanned Ground Vehicle for the 2005 DARPA Grand Challenge," *Journal of Field Robotics*, Vol. 23, Issue 8, Aug 2006.

[9] Touchton, R., Galluzzo, T., Kent, D., and Crane, C., "Perception and Planning Architecture for Autonomous Ground Vehicles," *IEEE Computer*, Dec. 2006, pp. 40 – 47.

[10] Coppin, Ben, *Artificial Intelligence Illuminated*, Jones and Bartlett Publishers, Sudbury, MA, 2004, ISBN 0-7637-3230-3, pp. 108-109.

[11] Touchton, R., et. al, "Planning and modeling extensions to the Joint Architecture for Unmanned Systems (JAUS) for application to unmanned ground vehicles," *SPIE Defense and Security Symposium*, Orlando, Mar 2005.

[12] Solanki, S., "Development of Sensor Component for Terrain Evaluation and Obstacle Detection for an Unmanned Off-Road Autonomous Vehicle," Ph.D. Disertation, Univ. of Florida, 2007.

[13] Nanning, Q., and Cheng, H.. "Springbot: a Prototype Autonomous Vehicle and Its Algorithms for Lane Detection." *IEEE Trans. on Intelligent Transportation Systems 5*, (2004).

[14] Odagiri, K., Kobayashi, K., Watanabe, K., Umino, H., and Nakaho, N., "Development of Active Contour Extraction for Autonomous Vehicle Lane Detection at Outdoor Environment." Society of Automotive Engineers, 2001.

[15] Matthews, John H., *Numerical Methods for Computer Science*, Engineering and Math, Englewood Cliffs, New Jersey: Prentice-Hall, 1987, 224-252.

[16] Lee, J., Crane, C., Kim, S., and Kim, J., "Road Following in an Unstructured Desert Environment using Monocular Color Vision as Applied to the DARPA Grand Challenge," *Proceedings of the International Conference on Control, Automation, and Systems (ICCAS 2005)*, Gyeong Gi, Korea, June 2005

[17] Gonzalez, R.C. and Woods, R.E., *Digital Image Processing*, Prentice Hall, 2002.

[18] Viola and Jones, "Rapid Object Detection Using a Boosted Cascade of Simple Features," *Proceedings IEEE Conf. on Computer Vision and Pattern Recognition*. 2001.

[19] Touchton, R. "An Adaptive Planning Framework for Situation Assessment and Decision-Making on an Autonomous Ground Vehicle," Ph.D. Dissertation, University of Florida, 2006.

[20] Mayne, D. Q., Rawlings, J. B., Rao, C. V., Scokaert, P. O. M., "Constrained model predictive control: Stability and optimality," *Automatica*, Volume 36, Issue 6, June 2000, pp. 789-814.

[21] Scokaert, P.O.M, Mayne, D.Q., Rawlings, J.B., "Suboptimal Model Predictive Control (Feasibility Implies Stability)," *IEEE Transactions on Automatic Control*, Volume 44, Issue 3, March 1999, pp. 648-654.

[22] Hart, P.E., Nilsson, N.J., Raphael, B., "A Formal Basis for the Heuristic Determination of Minimum Cost Paths," *IEEE Transactions on Systems Science and Cybernetics*, SSC4 (2), (1968), pp. 100-107.