Team AnnieWAY's Autonomous System

Christoph Stiller[1], Sören Kammel[1], Benjamin Pitzer[1], Julius Ziegler[1], Moritz Werling[2], Tobias Gindele[3], and Daniel Jagszent[3]

[1] Institut für Mess- und Regelungstechnik
[2] Institut für Angewandte Informatik/Automatisierungstechnik,
[3] Institut für Technische Informatik,
Universität Karlsruhe (TH), 76131 Karlsruhe, Germany
stiller@mrt.uka.de

Abstract. This paper reports on AnnieWAY, an autonomous vehicle that is capable of driving through urban scenarios and that has successfully entered the finals of the *DARPA Urban Challenge 2007* competition. After describing the main challenges imposed and the major hardware components, we outline the underlying software structure and focus on selected algorithms. A recent laser scanner plays the prominent role in the perception of the environment. It measures range and reflectivity for each pixel. While the former is used to provide 3D scene geometry, the latter allows robust lane marker detection. Mission and maneuver selection is conducted via a concurrent hierarchical state machine that specifically ascertains behavior in accordance with California traffic rules. We conclude with a report of the results achieved during the competition.

1 Introduction

The capability to seamlessly perceive the vehicle environment, to stabilize the vehicle, and to plan and conduct suitable driving maneuvers is a remarkable competence of human drivers. For the sake of vehicular comfort, efficiency, and safety, research groups all over the world have worked on building autonomous technical systems that resemble such capability (cf. e.g. [1,2,3,4,5]).

The *DARPA Urban Challenge 2007* has been a competition introduced to expedite mainly US research on autonomous vehicles. Its finals took place on Nov. 3rd, 2007 in Victorville, CA, USA. As in its predecessors, the Grand Challenges 2004 and 2005 [6,7], the vehicles had to conduct missions fully autonomously and unmanned without any intervention of or interaction with the teams. In contrast to the earlier competitions, the Urban Challenge required operation in 'urban' traffic, i.e. in the presence of other vehicles operated either autonomously themselves or by the organizer. The major challenge imposed was collision-free driving in traffic in compliance with traffic rules (e.g. right of way at intersections) while completing the given missions that included overtaking maneuvers, U-turns, parking, and merging into regular flow of traffic. Finally, recovery strategies had to be demonstrated in deadlock situations or in traffic congestions that

G. Sommer and R. Klette (Eds.): RobVis 2008, LNCS 4931, pp. 248–259, 2008.
© Springer-Verlag Berlin Heidelberg 2008

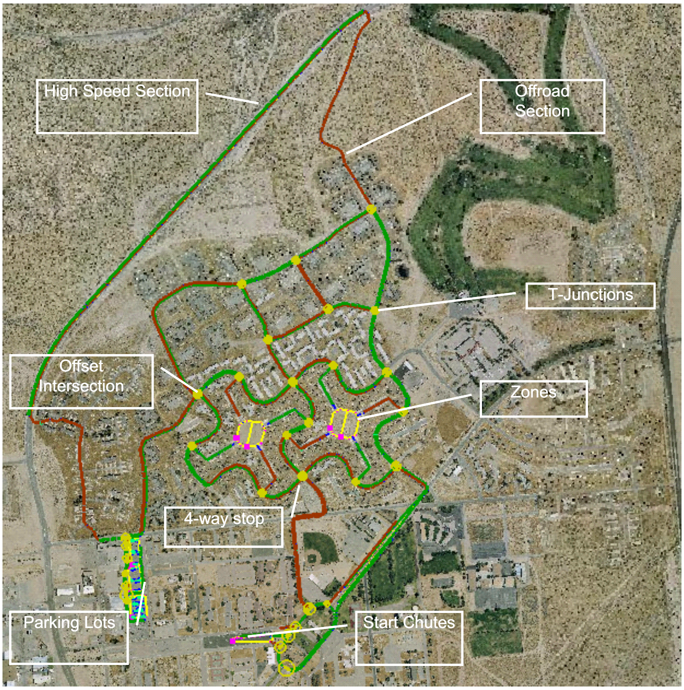

Fig. 1. Terrain of the Urban Challenge finals

cannot solely be handled with traffic rules. Fig. 1 shows the site of the finals in Victorville.

This competition has significantly enhanced research, progress, and public awareness in the field of autonomous driving.

Autonomous driving requires substantiated knowledge in different domains like engineering, computer and cognitive sciences: A robust vehicle architecture and design plays a central role. On-board sensor technology, sensor data analysis including localization and sensor fusion techniques, are required for a consistent perception of the environment and the ego pose of the vehicle therein. This information is eventually used for decision on and planning of a suitable behavior that is appropriate in a given traffic situation. The outcome of the planning is fed in real-time to a robust controller that keeps the vehicle in a stable condition under all circumstances. Fault and error detection and recovery procedures need to supervise both hardware and software of the autonomous vehicle. Omnipresent measurement uncertainties as well as contradictory sensor data have to be handled consistently. This broad variety of tasks shows that the Urban Challenge connects interdisciplinary areas of research with relevance to science, industry and community that may hardly be overestimated.

Team AnnieWAY mainly gathers partners from the Collaborative Research Centre on *Cognitive Automobiles* established by the German Research Foundation (DFG) on January 2006. This Collaborative Research Centre - like Team AnnieWAY - comprises partners from Universität Karlsruhe (TH), Forschungszentrum Karlsruhe, Technische Universität München, Fraunhofer-Gesellschaft IITB Karlsruhe, and Universität der Bundeswehr München.

The *Cognitive Automobiles* team conducts fundamental and interdisciplinary research of machine cognition in the context of mobile systems as the basis of machine action and the development of a scientific theory thereof. Its verifiability will exemplarily be proven in the way that the behavior of automobiles in traffic will be perceived, interpreted and even automatically generated. A cognitive automobile will be capable of individual as well as cooperative perception and interaction. A theory of machine cognition will make it feasible to propagate measurement uncertainty and symbolic vagueness throughout the complete cognition chain to give measures of certainty in order to quantify the trust in a specific generated behavior. Cognition includes perception, deduction and recognition and thereby supplies automobiles with completely novel abilities. Cognitive automobiles are able to sense themselves and their surroundings, as well as to accumulate and organize knowledge [8].

The scope of Team AnnieWAY was to extract early research results from the *Cognitive Automobiles'* project that would allow real-time operation of the vehicle under the restricted traffic environment in the Urban Challenge. Its team members are professionals in the fields of image processing, 3D perception, knowledge representation, reasoning, real time system design, driver assistance systems and autonomous driving. Some of the team members were in the 'Desert Buckeyes' team of Ohio State University and Universitt Karlsruhe (TH) and developed the 3D vision system for the Intelligent Offroad Navigator (ION) that traveled successfully 29 miles through the desert during the Grand Challenge 2005 [9,10].

2 Hardware Architecture

The basis of the AnnieWAY automobile is a 2006 VW Passat Variant B4 (Figure 2). The Passat has been selected for its ability to be easily updated for drive-by-wire use by the manufacturer.

AnnieWAY relies on an off-the-shelf AMD dual-core Opteron multiprocessor PC main computer whose computing power is comparable to a small cluster, yet offers low latencies and high bandwidth for interprocess communication. All sensors connect directly to the main computer which offers enough processing capacity to run almost all software components. The main computer is augmented by a dSpace AutoBox that operates as electronic control unit (ECU) for low-level control algorithms. It directly drives the vehicle's actuators. Both computer systems communicate over a 1 Gbit/s Ethernet network. The drive by wire system as well as the car odometry are interfaced via the Controller Area Network (CAN) bus. The DGPS/INS system allows for precise localization and connects to the main computer and to the low-level ECU (AutoBox). The chosen hardware architecture is

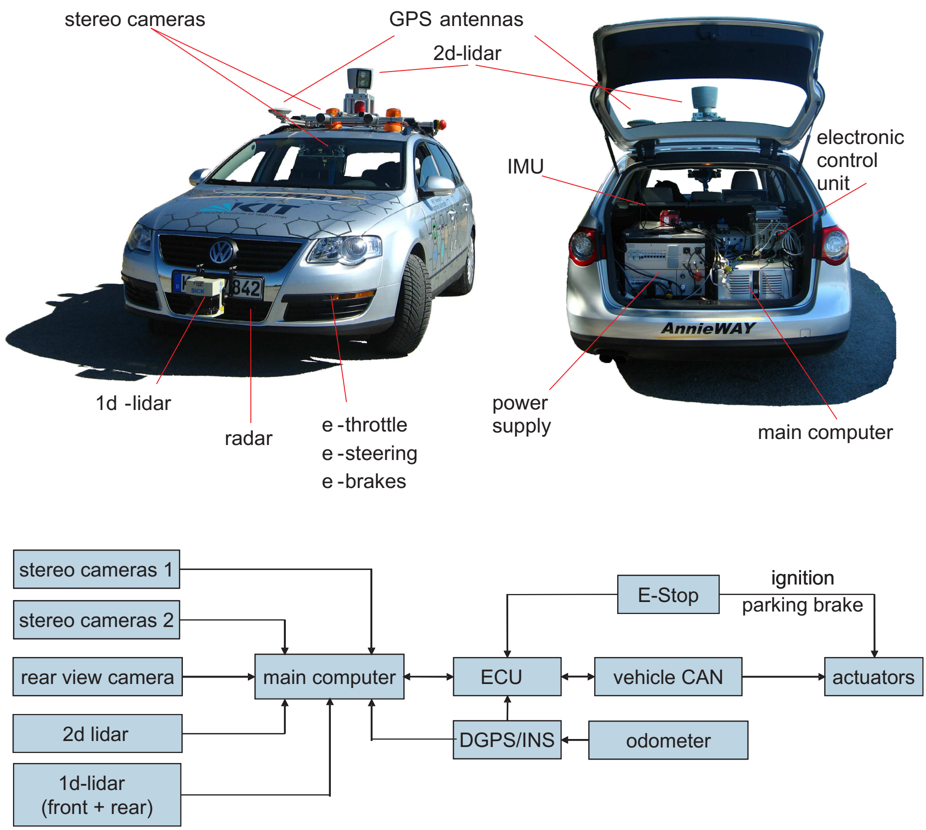

Fig. 2. Architecture and hardware components of the vehicle

supported by a real-time-capable software architecture which will be described in the next section [11].

2.1 Long Distance 2d Lidar

Since lidar units produce their own light, low light conditions have no effect on this kind of sensor. The Velodyne HDL-64E rotating laser scanner (www.velodyne.com/lidar/index.html) comprises 64 avalanche photodiodes that are oriented with constant azimuth and increasing elevation covering a 26.5-degree vertical field of view. The lasers and diodes are mounted on a spinning platform that rotates at a rate of 600 rpm. Thus, the HDL-64E provides a 360 degree field of view around the vehicle producing more than 1 million points per second at an angular resolution of 0.09 degrees horizontally and a distance resolution of 5cm with distances up to 70m. The result is a dense, highly accurate scan representation of almost the entire scene surrounding the vehicle. For each point, the sensor measures range and reflectivity. The reflectivity map is well suited for monoscopic image analysis tasks like lane marker detection. The inherent association of each reflectivity pixel with a range measurement alleviates information fusion of these data significantly.

2.2 Stereo Vision

An active camera platform with two cameras is mounted in the vehicle that surveys the forward 180 field of view. The system is able to conduct on-line self-calibration while the individual cameras are steered individually in the direction of interest [12], i.e. the system continuously calculates the precise orientation of its two cameras. This capability is essential for 3d stereoscopic vision.

2.3 DGPS/INS

For precise localization we use the OXTS RT3003 Inertial and GPS Navigation System which is an advanced six-axis inertial navigation system that incorporates a Novatel L1/L2 RTKGPS receiver for position and a second GPS Receiver for accurate Heading measurements. Odometry is taken directly from AnnieWAY's wheel encoders. The RT3003delivers better than 0.02m positioning accuracy under dynamic conditions using differential corrections and 0.1° heading accuracy using a 2m separation between the GPS antennas. The RT3003 Inertial and GPS Navigation System includes three angular rate sensors (gyros), three servo-grade accelerometers, the GPS receiver and the required processing. It works as a standalone, autonomous unit and requires no user input for operation.

2.4 Parking Lidars

Two additional Sick LMS 291 1D lidar scanners are mounted horizontally on the front and rear bumper to support the HDL-64E sensor and stereo cameras during parking maneuvers.

2.5 E-Stop System

As the vehicle had to operate unmanned, a remote stop system has been integrated for safety reasons as required by the organizer. This E-Stop system allows to remotely command run-, pause-, or emergency-stop mode via a wireless transmission. The system is connected directly to the ignition and the parking brake to ascertain appropriate emergency stop regardless of the state of the computer system. Run and pause mode are signalled to the low-level control computer.

2.6 Actuators

To actually control the car, actuators for steering, brake, throttle and gear shifting have been mounted. The steering actuator consists of a small electric motor attached to the steering column while the brakes are controlled by an additional brake booster.

3 Software Architecture

The core components of the vehicle are the perception of the environment, an interpretation of the situation in order to select the appropriate behavior, a path planning component and an interface to the vehicle control. Figure 3 depicts a block diagram of the information flow in the autonomous system. Spatial information from the sensors are combined to a static 2D map of the environment. Moving objects are treated differently. Such dynamic objects also include traffic participants that are able to move but have zero velocity at the moment. To detect moving objects, the spatial measurements of the lidar sensor are clustered and tracked with a multi-hypothesis approach. To detect possibly moving objects (which are standing still right now), a simple form of reasoning is used: If an object has the size of a car and is located on a detected lane, it is considered to be probably moving. Lane markers are detected in the reflectance data of the main lidar. Together with the road network map (RNDF), the absolute position obtained from the DGPS/INS system and the mission plan (MDF), this information serves as input for the situation assessment and the subsequent behavior generation. Most of the time, the behavior will result in a driveable trajectory. If a road is blocked or the car has to be parked, modules for special maneuvers, like the parking lot navigation module, are activated.

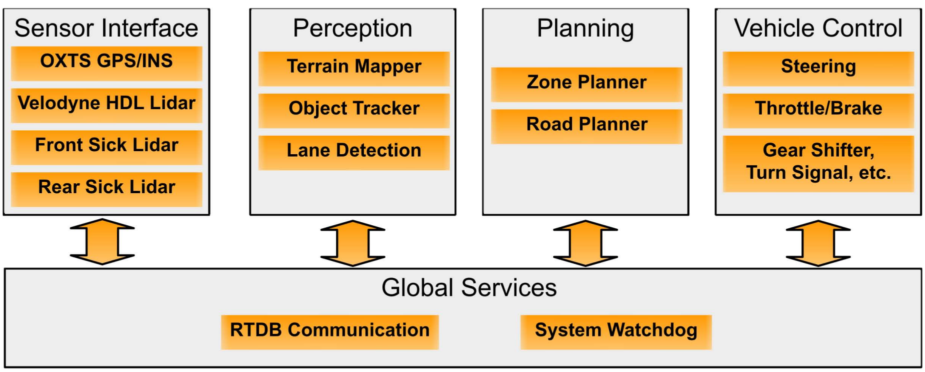

Fig. 3. Overview of the software architecture and the information flow

3.1 Environmental Mapping

Accurate and robust detection of obstacles at a sufficient range is an essential prerequisite to avoid obstacles on the road and in unstructured environments like parking lots. AnnieWAY uses a 2D grid structure $g_E(\boldsymbol{u})$ where each cell stores information on the local elevation. The grid map is constructed and updated asynchronously by the lidar sensors. The elevation $g_E(\boldsymbol{u})$ of the grid is the height range at a given cell position:

$$g_E(\boldsymbol{u}) = \max_{\boldsymbol{u}\in\mathcal{U}, l\in L} h(\boldsymbol{u}, l) - \min_{\boldsymbol{u}\in\mathcal{U}, l\in L} h(\boldsymbol{u}, l) \ ,$$

where L is the number of points recorded in a single rotation, $\mathcal{U}$ is the region representing the size of one grid cell (30cm$\times$30cm) and h is the vertical component of each scan point. The map is shifted along with the motion of the vehicle. This restricts the size of the map to an area around the vehicle while the cells are bound to an absolute position.

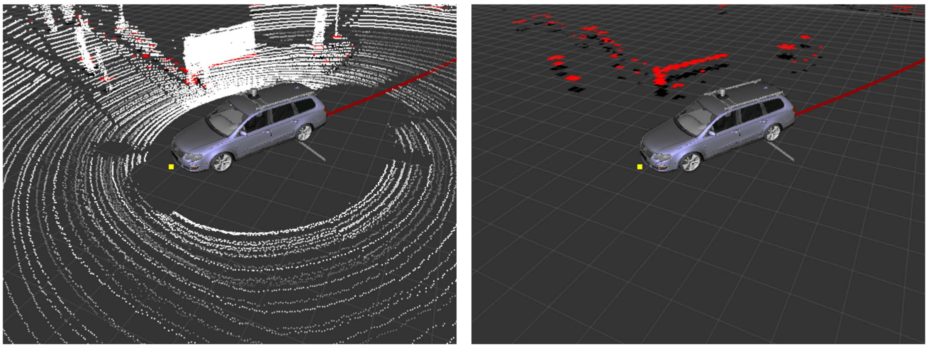

Fig. 4. Example for the mapping of 3D lidar data (left) onto a 2D grid (right)

Figure 4 shows an example for our mapping algorithm. On the left side, the raw sensor information is shown, the right side shows the map resulting from this data.

3.2 Lane Marker Detection

Lane markers are detected from the intensity readings of the lidar as depicted in the top of Fig. 5. In contrast to camera images, the laser reflectivity map is insensitive to background light and shadows while producing only a sparse intensity image. In order to increase the density of the lane marker information, subsequent scans are registered and accumulated employing INS information similar to the obstacle map described in the last subsection. An example of a resulting map is shown in the bottom of Fig. 5.

Lane markers are detected applying in the Radon transform domain of the accumulated reflectivity data. While Fig. 5 illustrates that the algorithm works well in real environments, it also clearly reveals a significant offset between the road network provided by the organizer and the visible lane markers. Thus even the correct detection of visible lane markers led to potential wrong or misleading lane positions in the competition. Hence the algorithm was suspended during the Urban Challenge.

3.3 Tracking of Dynamic Objects

Driving in urban environments requires to capture and estimate the dynamics of other traffic participants in real time. In our vehicle we use a detection and

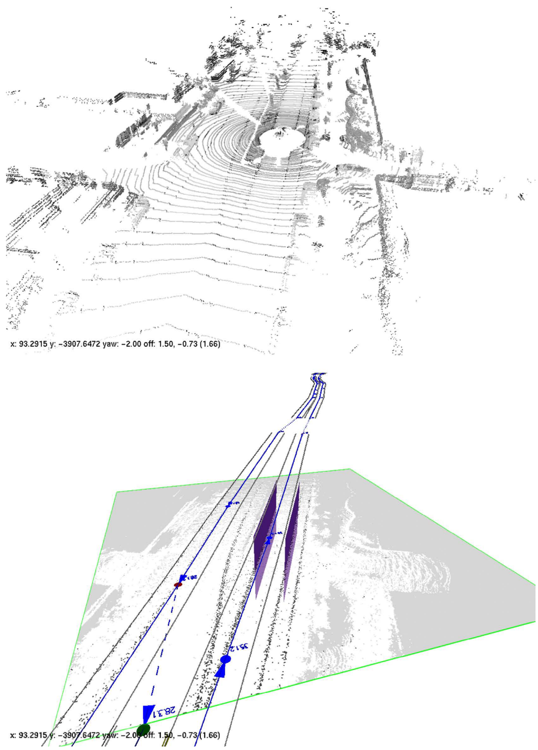

Fig. 5. Intensity readings of the lidar (top). Lane marker map, the estimated current lane segment and a part of the original road net (bottom).

tracking system that consists of two parts: Clustering of point cloud/image data into objects and tracking of these objects in subsequent sensor frames. Every cluster that cannot be associated to a known object yields initialization of a new object on the map. We use a linear Kalman filter to both predict the next location of a moving object, and to observe useful information in the tracking process which cannot be measured directly with our sensors like velocities or accelerations.

3.4 Mission and Maneuver Planning

As a first preprocessing step for the planning, all elements of the RNDF (lanes, checkpoints, exits, etc.) are converted to a graph-based representation. RNDF waypoints form the vertices of the graph; lanes and exits are represented by graph edges. Information such as distances, lane boundaries, and speed limits annotate the respective graph edges. These annotations can be updated dynamically, e.g. to describe a road blockage. The graph is used as input to the mission planner. It finds the optimal route from one checkpoint to another using an A* graph search algorithm. The search process is repeated for every pair of subsequent checkpoints in the MDF. The resulting route is written into the central real-time database together with all relevant information. Besides the obvious application of driving the shortest route, this enables other modules to make use of the map data, e.g. the lane recognition is supported by the a priori knowledge on the existence and type of lane markings. In addition to the information provided by the RNDF, the map also contains virtual turnoff lanes at intersections. A virtual lane connects each exit waypoint with a corresponding entry waypoint. It is generated assuming standard intersection geometry. If deviations between road network map and reality are detected, the map can be updated dynamically, enabling a safer and faster rerun of the respective lane.

High-level behavior decisions of the car are derived from the mission plan, together with the map data and the car's current position by a concurrent hierarchical state machine. The different states are also responsible for a behavior in accordance with the California traffic rules. Figure 6 shows the main level of the used state machine and the sublevel for intersection handling. The hierarchical structure simplifies the design and debugging of the state machine.

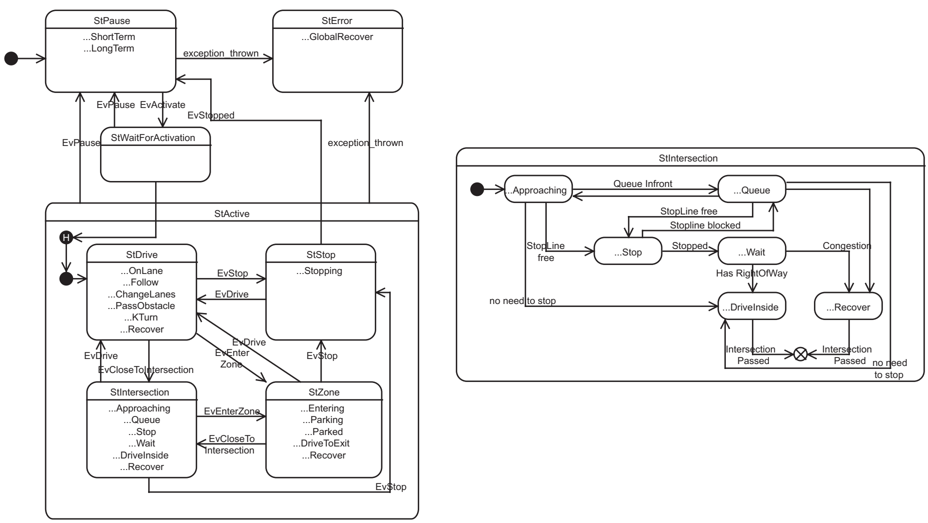

Fig. 6. Example from the concurrent hierarchical state machine used to model traffic situations and the behavior

3.5 Vehicle Control

The last step of the processing chain is the vehicle control which can be separated into lateral and longitudinal control.

For the lateral movement of the vehicle a software state space controller was chosen. This controller minimizes the distance and heading error of the car as it moves along the planned curve. A feed-forward term, that uses the curvature of the planned curve, strongly increases accuracy.

Since the vehicle's longitudinal dynamic is nonlinear, a compensation algorithm was chosen, which converts a requested acceleration into brake pressure and acceleration pedal values. The nonlinearity was compensated with the inverse of the static characteristic (engine and brake system). Velocity-dependant disturbances like wind drag and roll resistance of the wheels were measured in advance and accounted for in the control structure. For unpredictable disturbances, such as additional wind, an integrator was added. A higher control strategy asserts the bumpless transfer between a velocity, a stopping, and a following controller whereby the output of each controller is the desired acceleration (s. Fig. 7).

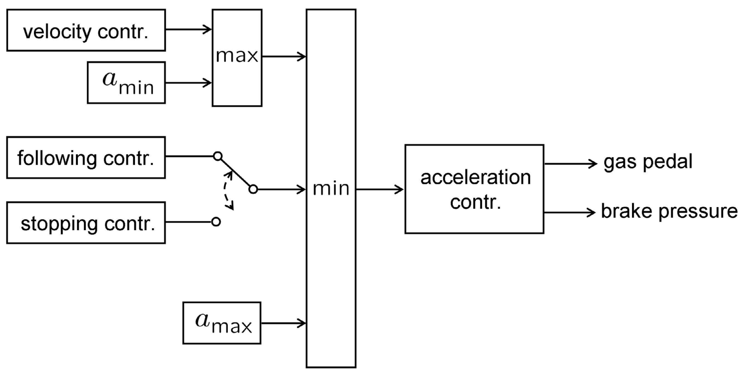

Fig. 7. High level strategy for longitudinal control

4 Results

Originally 89 teams have entered the competition, 11 of which were sponsored by the organizer. After several stages, 36 of those teams were selected for the semi-final. There, AnnieWAY has accomplished safe conduction of a variety of maneuvers including

- regular driving on lanes
- turning at intersections with oncoming traffic
- lane change maneuvers
- vehicle following and passing
- following order of precedence at 4-way stops
- merging into moving traffic

Although the final event was originally planned to challenge 20 teams, only 11 finalists were selected by the organizers due to safety issues. AnnieWAY has entered the final and was able to conduct a variety of driving maneuvers. It drove collision-free, but stopped due to a software exception in one of the modules.

Figure 8 depicts three examples of the vehicles actual course taken from a log-file and superimposed on an aerial image. The rightmost figure shows the stopping position in the finals.

Fig. 8. Three steps of AnnieWAYs course driven autonomously in the finals

The competition was won by CMU, followed by the teams of Stanford and Virginia Tech.

5 Conclusions

The autonomous vehicle AnnieWAY is capable of driving through urban scenarios and has successfully entered the finals of the *DARPA Urban Challenge 2007* competition. In contrast to earlier competitions, the Urban Challenge required to conduct missions in 'urban' traffic, i.e. in the presence of other autonomous and human-operated vehicles. The major challenge imposed was collision-free and rule-compliant driving in traffic. AnnieWAY is based on a simple and robust hardware architecture. In particular, we rely on a single computer system for all tasks but low level control. Environment perception is mainly conducted by a roof-mounted laser scanner that measures range and reflectivity for each pixel. While the former is used to provide 3D scene geometry, the latter allows robust lane marker detection. Mission and maneuver selection is conducted via a concurrent hierarchical state machine that specifically ascertains behavior in accordance with California traffic rules. More than 100 hours of urban driving without human intervention in complex urban settings with multiple cars, correct precedence order decision at intersections and - last not least - the entry in the finals underline the performance of the overall system.

Acknowledgements

The authors gratefully acknowledge the fruitful collaboration of their partners from Universität Karlsruhe, Technische Universität München and Universität der Bundeswehr München. Special thanks are directed to Annie Lien for her instant willingness and dedication as our official team leader. This work had not been possible without the ongoing research of the Transregional Collaborative Research

Centre 28 Cognitive Automobiles. Both projects cross-fertilized each other and revealed significant synergy. The authors gratefully acknowledge support of the TCRC by the Deutsche Forschungsgemeinschaft (German Research Foundation).

References

1. Bertozzi, M., Broggi, A., Fascioli, A.: Vision-based intelligent vehicles: State of the art and perspectives. J. of Robotics and Autonomous Systems 32, 1–16 (2000)
2. Franke, U., et al.: From door to door – Principles and applications of computer vision for driver assistant systems. In: Vlacic, L., Harashima, F., Parent, M. (eds.) Intelligent Vehicle Technologies: Theory and Applications, ch. 6, pp. 131–188. Butterworth Heinemann, Oxford (2001)
3. Nagel, H.-H., Enkelmann, W., Struck, G.: FhG-Co-Driver: From map-guided automatic driving by machine vision to a cooperative driver support. Mathematical and Computer Modelling 22, 185–212 (1995)
4. Thorpe, C.E.: Vision and navigation - The Carnegie Mellon Navlab. Kluwer Academic Publishers, Dordrecht (1990)
5. Dickmanns, E.D., et al.: The seeing passenger car "VaMoRs-P". In: Proc. Symp. On Intelligent Vehicles, Paris, pp. 68–73 (October 1994)
6. Darpa. Grand Challenge 2005 (October 2005), `http://www.darpa.mil/grandchallenge/overview.html`
7. Thrun, S., et al.: Stanley: The robot that won the DARPA Grand Challenge. Journal of Field Robotics 23(9), 661–692 (2006)
8. Stiller, C., Färber, G., Kammel, S.: Cooperative cognitive automobiles. In: Proc. IEEE Intelligent Vehicles Symposium, Istanbul, Turkey, pp. 215–220 (2007)
9. Özgüner, Ü., Stiller, C., Redmill, K.: Systems for safety and autonomous behavior in cars: The DARPA Grand Challenge experience. IEEE Proceedings 95(2), 1–16 (2007)
10. Hummel, B., et al.: Vision-based path-planning in unstructured environments. In: IEEE Intelligent Vehicles Symposium, Tokyo, Japan (June 2006)
11. Goebl, M., Färber, G.: A real-time-capable hard- and software architecture for joint image and knowledge processing in cognitive automobiles. In: Proc. IEEE Intelligent Vehicles Symposium, Istanbul, Turkey, pp. 734–739 (June 2007)
12. Dang, T., Hoffmann, C., Stiller, C.: Self-calibration for active automotive stereo vision. In: Proc. IEEE Intelligent Vehicles Symposium, Tokyo, Japan, pp. 364–369 (June 2006)