

Integrated Air/Ground Vehicle System for Semi-Autonomous Off-Road Navigation

Tony Stentz (Carnegie Mellon University); axs@cmu.edu
Alonzo Kelly (Carnegie Mellon University); alonzo@ri.cmu.edu
Herman Herman (Carnegie Mellon University); herman@ri.cmu.edu
Peter Rander (Carnegie Mellon University); rander@ri.cmu.edu
Omead Amidi (Carnegie Mellon University); amidi@ri.cmu.edu
Robert Mandelbaum (Sarnoff Corporation); rmandelbaum@sarnoff.com

Program Category: Command and Control
Presenters: Tony Stentz, Ph.D., Peter Rander, Ph.D.
National Robotics Engineering Consortium
Robotics Institute
Carnegie Mellon University
Ten 40th Street
Pittsburgh, PA 15201
412-268-8155
412-681-6961 (FAX)
www.rec.ri.cmu.edu

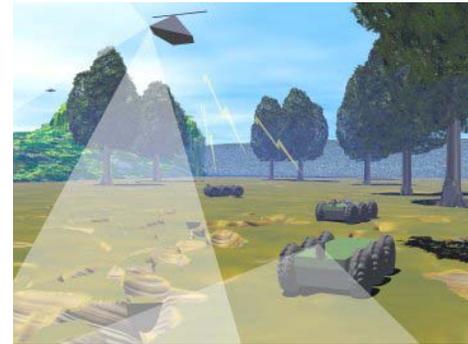
1. Abstract

Current unmanned vehicle systems enable exploration of and travel through remote areas, but demand significant communications resources and constant human operation. DARPA and the US Army have recognized these limitations and are now pursuing semi-autonomous vehicle systems in the Future Combat Systems (FCS) program. FCS places high demands on robotic systems, which must assess mobility hazards under all weather conditions, day and night, in the presence of smoke and other airborne obscurants, and with the possibility of communications dropouts. Perhaps the most challenging mobility hazard is the "negative obstacle", such as a hole or ditch. These hazards are difficult to see from the ground, limiting maximum vehicle speed. From the air, however, these obstacles are often far easier to detect.

In this context, we present a novel semi-autonomous unmanned ground vehicle (UGV) that includes a dedicated unmanned air vehicle - a "Flying Eye" (FE) - that flies ahead of the UGV to detect holes

and other hazards before the onboard UGV sensors would otherwise be able to detect them. This concept is shown in Figure 1. The FE can be considered a longer-range “scout” to explore terrain before the UGV must traverse it, allowing the UGV to avoid certain areas entirely because of hazards or cul-de-sacs detected by the FE. The UGV itself carries its own extensive sensor suite to address the broad range of

Figure 1: Blitz concept, with Flying Eye scouting for UGV.



operating conditions at slower speeds and to confirm the hazards seen from the air. In the fully developed system, the UGV will deploy the FE from a landing bay on the back of the UGV. For covert operations, the FE will return to its landing bay on the back of the UGV. This paper presents the current prototype system as well as initial field experiments on the performance of the system.

2. Perception for Off-Road Mobility (PerceptOR) Program

Team Blitz, consisting of the National Robotics Engineering Consortium at Carnegie Mellon University, Sarnoff Corporation, Robotics Engineering Excellence, Rockwell Science Center, Redzone Robotics, and The Boeing Company, has designed and developed an innovative unmanned vehicle semi-autonomous perception and navigation system under the DARPA PerceptOR Program. PerceptOR is a technology program addressing the Robotics Supporting Technology area of the US Army’s Future Combat System (FCS). The PerceptOR program seeks to remedy the poor performance of extended tele-operation by introducing significant autonomous perception, reasoning, and planning onboard the UGV, while directed by a remote human operator who can assist the UGV when it is unable to determine the best course of action. By introducing this autonomy, it is expected that the frequency and intensity of human involvement at the remote command post can be greatly

reduced, enabling a single operator to control multiple vehicles and reducing the bandwidth required between the UGV and the operator. To quantify these expectations, the PerceptOR program emphasizes real-world, unrehearsed field experiments under different terrain, weather, and communications conditions, both at night and in daylight.

In Phase I (Mar 01 - Nov 01), DARPA selected four teams to prepare and test critical elements of their respective designs. Three teams were selected for Phase II (Dec 01 – Nov 02), which involves deploying each prototype each quarter for field experiments. The experiment sites were chosen to present a broad range of terrain and weather conditions: Ft. AP Hill, Virginia (Feb 02); Yuma Proving Grounds, Arizona (May 02); Pickel Meadows US Marine Corps High-Altitude Training Center, California (Aug 02); and Ft. Polk, Louisiana (Nov 02). Phase III (starting Dec 02), will include more field experiments, with added focus on loss of assets such as GPS and communications.

3. Team Blitz Design Philosophy

As already noted, FCS places high demands on robotics perception systems. In fact, current perception technology leaves a gap in obstacle detection: intermediate obstacle sizes that are too large to negotiate with the inherent traversability and crash survivability of current vehicles are also too small to be seen at the stopping distances of those vehicles. The problem is compounded by the desire for increased vehicle speed, which dramatically increases stopping distances.

To address this “crash gap”, the conventional wisdom is to pursue ever-greater precision devices, to resolve smaller obstacles at greater distances. But these efforts produce diminishing returns. For example, a vehicle traveling 60 km/hr, with a half-second combined sensor and computer reaction time and a coefficient of friction of 0.6, requires nearly 30 m to completely stop. At this speed, a sensor must be able to detect debilitating hazards at least 30 m ahead of the vehicle. And given the

speed, even obstacles as small as 20 cm could be debilitating. While sensors such as ladars and radars can provide such ranging capability, contemporary versions cannot simultaneously supply both the angular resolution and the data rate required. Furthermore, negative obstacles – e.g., potholes and ditches – cannot be detected at all, regardless of sensor precision, because the obstacle is occluded until it is too late to react. Finally, the development of extremely high-resolution sensor technology usually leads to expensive, complex, fragile sensors, easily destroyed or knocked out of calibration.

Believing this approach to be misguided, Team Blitz has pursued a design based on three principles:

- Optimal Perspective: For best performance, view terrain from the best perspective possible: from the air. For negative obstacles, this is critical, but airborne sensors also are free to scout ahead for any hazards, both increasing UGV safety and helping reduced total mission duration.
- Data Fusion: Combine data from multiple types of multiple sensors from multiple perspectives. Different types of sensors often have complementary capabilities, leading to higher probability of detecting hazards. Multiple sensors observe a wider field of view, and offer redundancy. Multiple perspectives provide considerable enhancement in resolution, and reduce “blind spots”.
- Automatic Learning: Design the system to learn from past experience, to enable increased performance. Attempting to pre-program all possible conditions – terrain, weather, time of day, etc – into a perception system is an overwhelming proposition unlikely to succeed.

4. Blitz Design Description

The goal of PerceptOR is to develop autonomous perception systems that safely navigate a ground vehicle through a series of waypoints. To meet this goal, the Blitz design incorporates three hardware systems: the UGV itself, the FE, and an Operator Control Station (OCS). The use of the Flying Eye

directly addresses design principle 1, to acquire optimal perspectives of the terrain. The FE and the use of many sensors on the UGV together address design principle 2, to fuse multiple observations of terrain. Design principle 3, to allow the system to learn, is incorporated into the software used for perception and navigation. For brevity, this document will not discuss the base surrogate vehicle, a Honda Rubicon ATV, or the retrofit performed to make the vehicle computer controllable. Similarly, we will not discuss the process of making the air vehicle computer controllable. We will instead focus exclusively on the higher level design implemented to make the system autonomous.

The autonomous system design can be broadly divided into two categories: perception and navigation. Perception uses sensors and sensor data processing to observe the surrounding environment and to construct maps of terrain and obstacles. This process produces a map of cost, or risk, associated with the UGV passing through a given region. Navigation assesses the cost map to determine paths both safe and still leading to the waypoints, and directs the fly-by-wire vehicle to follow the best path.

It is important to note that the autonomous system treats the FE as a “programmable” sensor that can be directed to analyze terrain from viewpoints independent of the UGV location. The perception system treats the FE sensors as it does all others, albeit with a greater lookahead distance. The navigation system uses the current maps to plan a path for the FE in addition to the UGV.

4.1. Sensors and Sensor Processing

In keeping with the philosophy of fusing complementary perception sensor modalities, ground and air vehicles host several different sensors. Throughout the architecture, sensor processing is expected to provide both the best estimate of the terrain’s 3D structure and an analysis of this data for obstacles.

4.1.1. Flying Eye

The flying eye provides geo-referenced 3D geometry, which is analyzed for obstacles and terrain hazards. To get this information, the FE includes both terrain sensors and FE pose estimation sensors. For terrain sensing, we have systems that use either binocular stereo from video, or ladar. The stereo system, mounted on a Bergen

Figure 2: Manually-flown Stereo Flying Eye.



Industrial Twin helicopter (see Figure 2), consists of two Sony XC-999 cameras, a Horita BSG-50 sync generator, and one Acadia board – a purpose-built image processing board from Sarnoff Corporation. Pose is measured with a Novatel OEM4 RT-2 GPS receiver and a Watson Industries E304 AHRS. An onboard computer processes and geo-references the stereo data for hazards and handles communications with the UGV through IEEE 802.11b wireless Ethernet.

The second flight platform (see Figure 3) contains a custom-designed ladar system mounted on a Honda R50 helicopter that has been modified for computer control. Both the sensor and the autonomous flight system have been developed in another program at Carnegie Mellon University [2]. The ladar contains a Riegl LD-90 laser rangefinder with an externally mounted scanning mirror controlled by custom electronics. This helicopter also includes a pose estimation system, including a Novatel RT-2 GPS, a Litton LN-100 IMU, and a KVH flux-gate compass, fused with a custom processing module to produce the pose estimate. Onboard TI 'C44 DSP processors process the ladar data into geo-referenced point clouds. Communications is handled with a wireless Ethernet connection.

Figure 3: Automated Ladar Flying Eye.



4.1.2. Unmanned Ground Vehicle

The UGV also includes ladar and stereo for terrain analysis and its own pose estimation system. While ladar is an active modality, it provides imagery of unequalled fidelity under many operating conditions, including night operations. We use four single axis industrial ladars (Sick Industries LMS-200), each mounted on a separately actuated orthogonal second axis. Two sensors look forward, and two backward, giving the UGV better ability to drive in reverse. This design represents a highly affordable 2-axis ladar system (\$6K hardware cost) with little or no reduction in data quality over higher cost units with internal 2-axis scanning and over experimental solid-state units. Once proven, solid-state units are preferred because of the expected increase vibration and shock resilience.

Stereo sensing, provided by team partner Sarnoff Corporation, uses the purpose built Acadia image processor board [6] to perform intensive stereo matching calculations. Stereo is a passive sensing option, for increased stealth, and complements the terrain analysis of ladar. We currently use three forward-looking binocular stereo pairs, to give both high resolution and large FOV. Infrared cameras, for passive night vision and improved terrain discrimination, are planned.

For remote operator assistance, we employ two video cameras that stream video over Ethernet and provide pan, tilt, and zoom capabilities to the remote operator. If the autonomous system becomes confused, the remote operator can use these features to gain situational awareness without having to drive the vehicle around and examine the output of the sensors used by the autonomous system.

To maintain awareness of vehicle pose – position, attitude, and heading – we employ a system containing GPS, attitude sensing, and a single-axis gyro mounted horizontally to monitor UGV heading. The system also makes use of odometry from the vehicle tires. A Kalman filter is used to fuse the sensor indications into a global estimate of vehicle position. When high-accuracy GPS is

available, it keeps the absolute accuracy high, but if GPS drops out, the global estimate is still able to run on odometry. The system also maintains an independent local pose estimate generated purely from 3D odometry. The local estimate is guaranteed to be smooth, an important property in obstacle avoidance. The global estimate may contain position jumps associated with GPS returning after dropout, but is needed for waypoint following and for all references to the global map. In this case, the associated position jump is advantageous and should be passed to the output.

4.2. Mapping for Navigation

Two levels of environmental maps are maintained in the system. At the highest level, the global map incorporates all information available about the test range in the form a two-dimensional cost field. In this map, positions of progressively higher cost are progressively less favorable and should be more actively avoided. Depending on the size of this map, pieces of it may be loaded periodically into a local sub-map used for planning to the next waypoint. Map information originates from three sources.

- High fidelity terrain maps – from pre-mission over-flights or from DTED data – of the test site are intermittently available. When they are, the geometric and radiometric information that they contain is converted to cost information by applying our Flying Eye sensor processing.
- Flying Eye: Our own flying sensor generates data similar to the above with the added advantages that the data is available on demand and is potentially of both higher fidelity and reduced age.
- UGV sensors: Local mapping algorithms based on processing the UGV perception sensors also provide information that is merged with the above two sources. In this way, the effects of any omissions, changes, or errors in the aerial data can be potentially mitigated.

At the local level, UGV perception algorithms maintain a high-resolution multi-faceted terrain map solely for the purpose of obstacle avoidance. This map receives information from the same three sources as the global map. Unlike the global map, however, this map uses the local pose estimate discussed earlier. By doing so, the perception system maintains a smooth relationship between observed terrain and UGV position. If the global pose were used, the perception system would see all the obstacles “move” if the global pose jumps, even if the UGV is physically motionless. This could lead the UGV to plan around an “echo” of an obstacle and actually run into the real one.

As this phenomena shows, generally one of the key issues in mapping is localization. The data from all sensors must be fused into a coherent representation of the environment that is larger in extent than any single sensor frame can observe at the desired resolution. Our present approach relies on the dual pose estimate for the UGV, and the single, global estimate from the FE. We are actively exploring visual odometry on the FE, to have the same abilities as on the UGV. Data registration remains an interesting alternative but is not one we have invested in heavily.

4.3. Global Navigation

Our global navigation system is based on the most recent version of the D* (for “dynamic A*”) global planning algorithm originally developed on the Demo II program [3][4][5]. The path planned is the lowest cost path available that moves the vehicle from its current position through a sequence of waypoint goals. The initial path is based on any available prior information. The path is recomputed continuously in order to preserve optimality by incorporating both new sensor information about the locations of obstacles and the current location of the vehicle, which may have deviated from the previous path to avoid obstacles.

4.4. Local Navigation and Obstacle Avoidance

Our local navigation is based on the most recent version of the RANGER local navigation system originally developed on the Demo II program [1]. This system operates on shorter time scales and uses much shorter lookaheads than the global planner. Based on an optimal control formulation of obstacle avoidance, it samples the set of presently feasible vehicle commands and evaluates them for encounters with hazards such as discrete obstacles, high centering, tipover, etc. Hazard information is then converted to equivalent costs, which are combined with the global planner costs in order to decide on which command to issue to the vehicle.

4.5. OCS

While the highest possible levels of autonomy are being sought, human interfaces for remote monitoring and troubleshooting are fundamental to the system design. Our operator control station is designed to be portable while emphasizing situational awareness and reduced communication signature. In addition to various textual and graphical status displays, a synthetic overhead perspective rendering of the vehicle in-situ is annotated with various forms of useful information, such as the positions of nearby waypoints and obstacles, an overhead color image (if available), and the paths that the vehicle is considering. An intuitive steering wheel and accelerator pedal interface is available but its use is discouraged because of the large bandwidth required for safe vehicle control.

5. Blitz Experimental Results

To demonstrate the performance of the UGV system, we now show experimental results of the real system in three different mission types. In the first, the UGV is sent on a mission without any prior knowledge of the terrain it must traverse – we call this the UGV Sensors Only mission. In the second,

the UGV may be sent out with a prior map of the area – the UGV Sensors and Prior mission. Note that the prior map is actually produced with our Flying Eye, with all terrain mapped before the missions. The third mission uses sensors and prior data, but the prior data is “stale”, i.e., the scene has changed between the prior data collection and the actual mission – the UGV Sensors and Stale Prior mission.

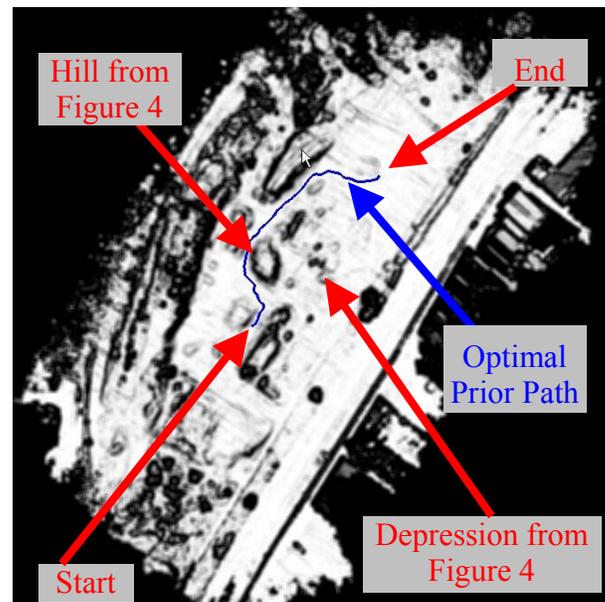
To date, the Blitz autonomous UGV has been used for 6 months, while the FE system has been used for 4 months. The integrated UGV-FE system, with both vehicles simultaneously moving, has operated for limited technical experiments. To show the full benefit here, we add a simulated fourth UGV mission, with performance of the integrated system traversing a real map – the UGV-FE mission.

5.1. Mission Overview

Figure 4 shows the Blitz UGV in several positions on the test course. Figure 5 shows an overhead map of the course. The overhead map is the prior data used in runs with prior data available, collected with



Figure 5: Overhead prior map of course, computed with our Flying Eye. Whiter is safer.



our FE. The map shows safe areas as white, impassible areas as black, and less severe hazards as gray. The map also shows the route that the UGV will want to take if prior data alone is available. For reference, the map includes annotations showing the areas corresponding to the photos of Figure 4.

5.2. Mission: UGV Sensors Only

For this run the UGV relied entirely on its onboard sensors, with no FE support and no prior data. The actual path taken is shown in Figure 6. The path is approximately 315 meters long, and took the UGV approximately 13.5 minutes to complete. The figure shows the UGV looping back and re-approaching some terrain, then choosing a sub-optimal route. The remote operator intervened twice in order to minimize further exploration on the part of the UGV. Clearly, without prior knowledge, the UGV is able to keep itself relatively safe, but lacks sufficient context to plan a good path to the goal.

5.3. Mission: UGV Sensors and Prior from FE

For this run, the UGV was also supplied with accurate prior data from the FE – i.e., the cost map as shown in Figure 5. The actual path taken in this case is shown in Figure 7. The path is approximately 244 meters long and took the UGV 8.5 minutes to complete – about 20% less distance traveled and 40% faster than with sensors alone. The route followed is far

Figure 6: Path of UGV with onboard sensors only.

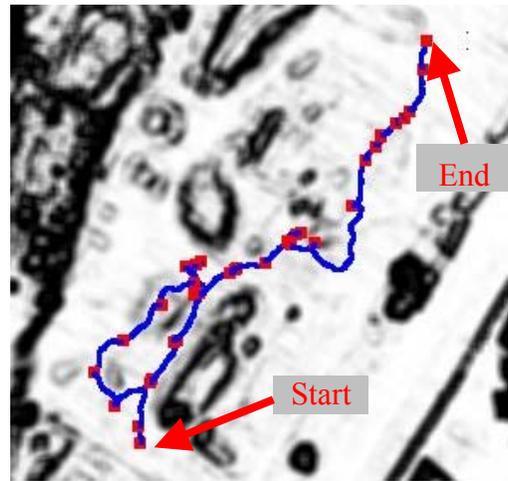
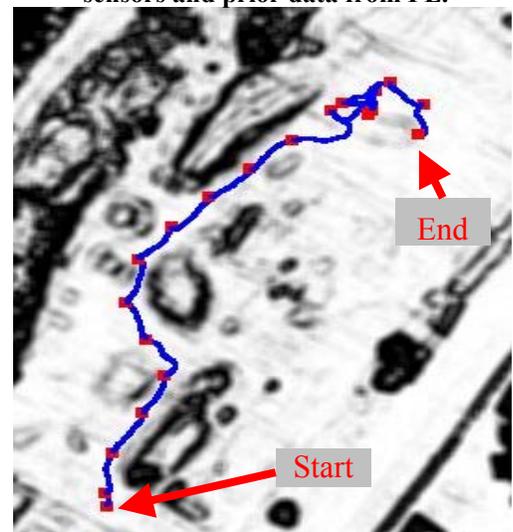


Figure 7: Path of UGV with onboard sensors and prior data from FE.



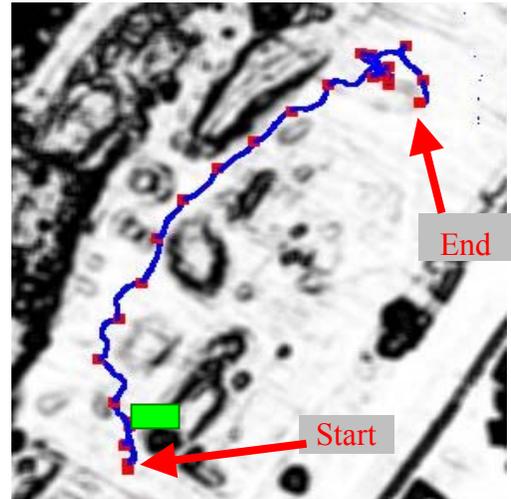
better.

5.4. Mission: UGV Sensors and Stale Prior

For this run, the UGV again uses onboard sensors and prior data, but the actual terrain is intentionally altered to block the desired path. This can be seen in Figure 8, by the large green obstacle inserted in the gap the UGV passed through in Figure 7. This path is

approximately 277 meters long, and took approximately 12.5 minutes to traverse. The path is still more than 10% shorter (distance and time) than the sensors-only path without changes. Without the onboard sensors, the UGV would not have seen the new obstacle, so onboard sensing is essential, but the prior data is also needed to perform large-scale, intelligent planning.

Figure 8: Path of UGV with onboard sensors and stale prior data, with new obstacle (shown in green).



5.5. Mission: UGV-FE

As of this writing, the FE systems have been in operation as prior-data collection systems, with the UGV receiving, processing, and logging the FE terrain telemetry. We have also had initial success with integrated UGV-FE activities in which the FE flies under manual control as the UGV follows along behind, but these missions have all been short tests, perhaps 30 meters in length. We therefore use a simulation to show the benefit of integrated flights. For this mission, the UGV begins in the middle-left of the cost map, rather than the previous starting location. The UGV has no prior map available. In the simulation, the UGV and FE each see ahead of themselves several meters, and the UGV directs the FE to explore regions on the path up to 30 meters ahead of the UGV. In this case, the total cost of the route taken is about 20% lower than without the FE – meaning that the UGV is

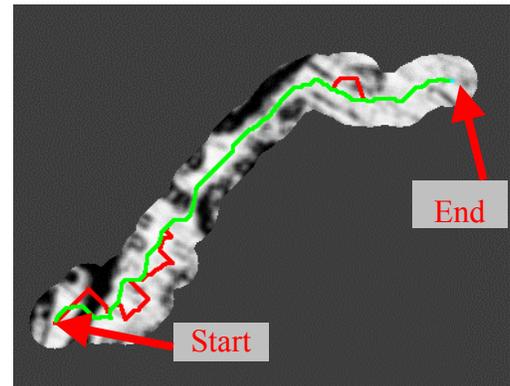
more likely to survive the mission. Figure 9 shows the map of the terrain covered by the UGV and FE. Note that this map includes only the information directly observed from the two systems, not the full map. We expect to have significant integrated tests soon.

6. Summary

We have presented the design, implementation, and initial experimental results of a novel semi-autonomous unmanned ground vehicle (UGV) that includes a dedicated unmanned air vehicle – a "Flying Eye" (FE) – that flies ahead of the UGV to detect holes and other hazards before the onboard UGV sensors would otherwise be able to detect them. The UGV itself is equipped with onboard sensors to be self-sufficient, and is able to use prior maps for global planning. Experimental results suggest that while the UGV is generally able to maintain its own safety at slow speeds, the prior data is extremely useful in the presence of complex terrain, substantially improving performance metrics.

Experimental results also suggest that the Flying Eye, a small UAV dedicated to and eventually carried on the UGV, can provide high-quality information useful to the UGV. All "prior" data used in the experimental results section were actually "batch mode" processing of a sweep of the course using our automated Flying Eye. Simulations of integrated UGV-FE operations suggest that the air-ground tandem of an FE scouting ahead of a UGV, with both in motion, is an effective way to improve the safety and efficiency of ground vehicle missions. Team Blitz is actively working to integrate and better demonstrate this integrated capability, and to explore the varied ways in which the FE can work with the UGV to improve performance.

Figure 9: Simulated path of FE (red) and UGV (green) with onboard sensors.



7. Acknowledgements

Although the author list is long, there is an even longer list of people who have been essential for Team Blitz's achievements: Michael Bode, Mark DeLouis, Antonio Diaz-Calderon, Robert Derham, Todd Dudek, Keith Hanna, Jeffrey Hibner, James Ketterer, Eric Kratzer, Jeff McMahill, Ryan Miller, Jorgen Pedersen, Adam Petruszka, André Rieder, Harold Rosenstein, Garbis Salgian, John Southall, Randon Warner, and Jonathan Woytek. Team Blitz is funded by the DARPA Perception for Off-Road Mobility (PerceptOR) program under contract number MDA972-01-9-0016.

8. References

- [1] Kelly, A., Stentz, A. "Rough Terrain Autonomous Mobility - Part 2: An Active Vision, Predictive Control Approach," *Autonomous Robots*, 5, 163-198, 1998.
- [2] Miller, R., Amidi, O. "3-D Site Mapping with the CMU Autonomous Helicopter," *Proceedings of 5th International Conference On Intelligent Autonomous Systems (IAS-5)*, June 1998.
- [3] Stentz, A. "Optimal and Efficient Path Planning for Partially-Known Environments," *Proceedings of the IEEE International Conference on Robotics and Automation*, May 1994.
- [4] Stentz, A. "The Focused D* Algorithm for Real-Time Replanning," *Proceedings of the 14th International Joint Conference on Artificial Intelligence (IJCAI-95)*, August 1995.
- [5] Stentz, A., Hebert, M., "A Complete Navigation System for Goal Acquisition in Unknown Environments," *Autonomous Robots*, Vol. 2, No. 2, August 1995.
- [6] van der Wal, G., Hansen, M., Piacentino, M. "The Acadia Vision Processor," *International Workshop on Computer Architectures for Machine Perception*, September 2000.