

***Prospect Eleven's* Autonomous Accomplishments while relying on Stereo Vision for Object and Physical Boundary Detection**

by

Alain L. Kornhauser¹, Brendan Collins², Gordon Franken², Andrew Saxe², Anand Atreya³, Bryan Cattle³, and Scott Schiffres⁴

Abstract

Presented is a concise summary of the stereo vision object detection system and achievements of *Prospect Eleven*, Princeton University's entry in the 2005 DARPA Grand Challenge. Described are the simplifying assumptions and limitations of the system that enabled *Prospect Eleven* to complete 9.4 miles during the Grand Challenge. More importantly, after changing only "one line of code" it returned to the desert at the end of October to autonomously re-run of not only the 2005, but also the 2004 Grand Challenge course.

1. Background

Prospect Eleven is the vehicle name for Princeton University's entry in [DARPA's Grand Challenge of 2005](#), a competition of truly autonomous vehicles traveling a prescribed course. The challenge of the "Challenge" was to "build or modify" an automobile-sized vehicle that can negotiate a prescribe course containing randomly place obstacles without any human intervention. The Challenge was originally contested in March, 2004; however, none of the entries completed the course. In fact, the most successful vehicle completed only the first seven miles of the more than 140 mile course. As a result, [DARPA](#) (Defense Advanced Research Projects Administration), decided to organize a second Challenge which took place on October 8, 2005.

Princeton University did not participate in the first Challenge; however, upon learning, in April 2004, that a second Challenge would be contested, a group of undergraduates led by Ben Klaber'05 approached Professor Alain Kornhauser with a desire participate in the 2005 Challenge. In May 2004, Princeton University officially enter a vehicle named "*Prospect Eleven*" with the stipulations that it be an undergraduate student, as opposed to a professional staff, activity **and** that the principle objective be that it complement to the highest degree possible the students' academic experience at Princeton. Throughout the design, build and test process leading to the competition and beyond, the guiding objectives have been academic relevance, simplicity, elegance and minimal expenditure of funds. Academic relevance because it is Princeton; simplicity because it is too easy to make this initiative so hard that it is undoable; thus, the need for elegance. Finally, funds were to be used only for summer stipends for participating students, the purchase of needed computing, command and control apparatus and travel expenses associated with the competition and post-competition activities. No funds were available for students during the academic year, professional staff, nor advising faculty. Student participation during the academic year had to be justifiable for academic credit or as an extra-curricular activity. In total, approximately \$125,000 was spent over a period of eighteen months.

¹ Team leader, Professor Operations Research & Financial Engineering , Princeton University

² Class 2008, Princeton University

³ Class 2007, Princeton University

⁴ Class 2006, Princeton University

3. Obstacle Detection Using Stereo Vision.

Among Prospect Eleven's most distinctive features is its reliance on stereo vision. Indeed, it was the only vehicle in the 2005 Grand Challenge Event which relied exclusively on stereo vision for the detection of obstacles and physical elevation changes marking road edges. It, along with the GPS-based inertial navigation array, is the only two sensors of the external world added to the base vehicle. This section discusses the challenges of using stereo vision for an Autonomous Ground Vehicles, describes the hardware and algorithms Prospect Eleven uses to detect and track obstacles, and considers future work in the field.

Stereo vision, the process of converting two simultaneous images from synchronized spatially-separated cameras into a depth image, is a well-studied problem, see [5] Forsyth, D. & Ponce, J., (2003). The key challenge is to reliably determine the correspondence between the features of an object as they are captured on each of the two images. The separation distance between the corresponding features on each of the images is inversely proportional to the distance from the cameras. A depth image, or disparity map, is the ensemble of the distances of all corresponding features in a scene. Several vendors sell systems which include synchronized cameras and SDK that can produce a disparity map tuned to several parameters. Purchased was one such system, known as a Bumblebee™, from Point Grey Research (PGR). It contains two black and white CCD cameras at a 12 cm baseline separation. The PGR Software Development Kit (SDK) produces a disparity map at a rate of approximately 16Hz from synchronized images. Given the disparity maps, we focused on the problem of obstacle detection and tracking given the capability to generate a range of disparity maps as a function of several parameters.

Stereo's reputation for producing noisy results is well-deserved. Data is heavily quantized due to pixilation of the image— that is, a point can take on only a small number of possible depth values. Moreover, in areas of low texture, correspondence of features can become very unreliable. Though the PGR SDK had fairly robust validation routines and would generally not report false matches, many environments generated sparse disparity maps. Lighting conditions also present a problem. Images with shadows, for instance, are frequently either excessively dark in the shadowed region or are washed out in areas the lighted region. These are issues that can not be ignored. Our efforts to deal with them fall into three main categories:

1. Ensuring that images of the scene have sufficient contrast to generate dense disparity maps,
2. Using obstacle detection algorithms which are robust to noise and can take advantage of quantization, and
3. Tracking known obstacles in the time domain.

3.1 Generating disparity maps. Several strategies proved effective for improving the quality of scene images. Red and UV photographic filters, mounted in front of each lens, help increase contrast and remove specular features. In particular, red filters helped reduce the intensity of the sky and sun, preventing issues such as CCD "bleeding." Though the CCDs have their own autogain control, it seems to be tuned to generate a contrast level more appropriate for human viewing than for disparity processing. Fortunately, the PGR SDK allows the camera's gain to be controlled in code. By experimentation, it was found that the best depth maps were generated by relatively dark images. Implemented was the following simple control law to govern the camera's gain:

$$G' = G + k(C - T),$$

where G' is the new gain at a given iteration, G is the current gain, k is a gain term,

C is the current average intensity value sampled over some region of interest, and
T is the target intensity value.

This simple control law is far from the state of the art in control theory. The camera was sometimes slow to adjust to sudden changes in brightness such as a transition in and out of a dark tunnel. However, it was adequate for nearly all situations Prospect Eleven encountered. The combination of photographic filters and improved control of the camera's gain dramatically improves the range of situations in which the PGR SDK can generate full disparity maps.

Despite these improvements in image quality pre-disparity matching, problems with the depth images remain. Ideal lighting conditions do not guarantee the accuracy of the correspondence matching throughout the image plane. Fortunately, the PGR SDK's validation routines are quite robust, and tend to reported only reliable matches. Thus, the disparity maps can at times be quite sparse; however, one can reliably assume that data reported is accurate, at least to the extent possible within the constraint of heavy quantization.

3.2 Obstacle detection Given the disparity map, the problem at hand was the detection of obstacles in the field of view. Implicitly, it was assumed the viewed landscape consisted of a plane surface with substantial "obstacles" and travel lane edges seated on that surface extending above (and possibly below) that ground plane. The implication of such a simplified view of the world is that an obstacle free surface would exhibit a map whose disparity would monotonically increase as one moved up the map and be constant across the map. Obstacles perpendicular to that surface would exhibit constant disparity throughout. With this in mind, Prospect Eleven used the following algorithm to isolate obstacle:

- 1) FOR each column in the disparity map
 - a) Consider a DFA with states {IN OBSTACLE, NOT IN OBSTACLE}
 - b) Begin in state NOT IN OBSTACLE
 - c) FOR each pixel in the column, starting at the top of the disparity map
 - i) IF state is NOT IN OBSTACLE
 - (1) Consider the next m pixels. If they are all the same, transition to state IN OBSTACLE
 - (2) Consider the difference, disparity for the next pixel minus disparity for this pixel. If this is greater than k_D , transition to state IN OBSTACLE
 - ii) IF state is IN OBSTACLE
 - (1) Calculate the variance over the next m pixels.
 - (2) IF this quantity is greater than k_V , transition to state NOT IN OBSTACLE.
 - (a) Calculate the variance over the last span of pixels for which the DFA was in state IN OBSTACLE.
 - (b) Add this span, and its variance, to a list of obstacles for this column

(This algorithm is a bit simplified, as the actual implementation includes logic to deal with invalidated pixels. However, this logic greatly complicates the algorithm without adding to the discussion.)

It is worth discussing the decision to process on disparity maps directly. As [1] notes, processing a disparity map is substantially faster than working on an elevation map (as does [2]), and maintains the highest possible resolution of data.

Several properties of the simple column-detection algorithm are advantageous. First, it is robust to sparse disparity maps and high quantization. Because it does not rely on any global characteristics of the image, such as the computation of a global ground-plane, sparse disparity maps do not greatly inhibit performance. Instead, two conditions are sufficient: an abrupt change from background to a foreground object, or an object which is (approximately) parallel to the image plane. The algorithm takes advantage of heavy quantization, as objects which are approximately parallel to the image plane will have the same disparity value throughout.

There are trade-offs to the simplicity of this algorithm, however. Because distance is inversely related to disparity, we should expect that vertical bands of the same disparity value will grow smaller as disparity increases. The constant m assumes that this value remains constant. Approaches like [1] do not suffer from this problem, but are slower as a result. Figure 3.1 below provides an example of a scene, the corresponding disparity map, the generated columns and the resulting bounding rectangles representing the obstacles.

This simple algorithm runs in time linear to the number of pixels. At the termination of the algorithm, a list of row-spans in each column is classified as obstacles. The variance serves as a measure of the confidence in an obstacle's existence. A local search algorithm bounds these spans with rectangles, ensuring that there exists a degree of uniformity across columns. This step also helps with much of the noise in the disparity map, as a great deal of noise does not exist in several contiguous columns. Once a bounding rectangle is determined, points in the rectangle are ranged, and averaged. Assigned is an (x,y) location for each obstacle relative to the nose of the vehicle.

3.3 Filtering in the time domain. Unfortunately, expected error in range measurements naturally increase as the square of the range of those elements. Thus, accuracy of distant objects gets poor very quickly. For Prospect Eleven, practical ranging was limited to a depth of less than 75 feet. However, the nature of the problem is such that as distant objects are approached, they will be imaged many times. Thus, performing "correspondence" in the time domain can become very helpful. By merging assumed kinematics for the ranged object (say, stationary) and the known kinematics of the vehicles (camera system), one can perform an appropriate weighted "average" of the time sequenced range values to obtain not only a confidence value on the existence of an object but also a best estimate of its latest location. To this end, Prospect Eleven maintains a list of all obstacles of concern to the collision avoidance routine. Obstacles isolated by a new frame are compared to the list of existing obstacles. Each new obstacle is matched to an existing obstacle, or declared to be a new obstacle. If it is matched with an old obstacle, the position is updated as a linear combination of the old and new position. In addition to providing dramatically increased accuracy in positioning, this allows the removal of many false positives, as their randomness tends to not have them appear in the same location in multiple frames.

3.4 Limitations and Extensions of Stereo Vision. Prospect Eleven benefited substantially from its assumption that the environment ahead was very simple. That it was simply an infinite plane with stationary obstacles standing perpendicular to that planar surface. This enabled Prospect Eleven to easily discern posts, walls, parked vehicles, shrubs lining the side of the road and even deep crevasses. Objects used in the NQE such as tire stacks, hay bales and Normandy-style tank traps were ideal for the vision system because they contained sharp discernable edges that were ideal for the correspondence algorithm. Also shadows cast on the surface ahead served to provide enhanced

correspondence and in no way fooled the vision system. Items in the natural environment such as tree trunks, fences, “New Jersey barriers” and the sheer grade changes encountered in the most treacherous mountain passes all were ideal “objects” for the correspondence algorithm.

What was more challenging was the limited range of the vision system. This range could have been increased by increasing the separation of the cameras. This would however, have reduced the field of view. Some preliminary experiments were conducted aimed at using a pair of stereo cameras separated across the hood of the vehicle. Technical difficulties in managing the data from the four cameras caused the effort to be abandoned; however, this is believed to be a promising area for future research. What may also prove to be more valuable for increasing the range is to perform correspondence in the time domain rather than the time-synchronized image plane. Tracking corresponding objects over time in the image plane near the horizon is expected to provide an opportunity to increase the reliable range of the vision system.

Increasing the reliable range of the image system is the key enabler to increasing maximum operating speed. In order to avoid collisions, obstacles must either be dodged in a stable manner or the vehicle must be able to come to rest prior to reaching the obstacle. Each condition defines a minimum object detection distance that increases as the square of the speed, and dominated by the stopping criteria. Given that Prospect Eleven’s stereo vision system had a maximum effective range of less than 75 feet, its speed is limited to approximately 25 mph in the presence of unexpected obstacles. Achievement of a 75 foot range requires that the scene ahead is rather simple to simplify correspondence image analysis and the utilization of tracking in the time domain to build confidence and filter false alarms. In total, Prospect Eleven demonstrated that stereo vision is a feasible platform for mid-speed autonomous ground vehicle navigation. Ongoing research by the authors aims to extend stereo’s range through the time-domain correspondence of features, such that it will be a suitable sensing platform for high-speed navigation.

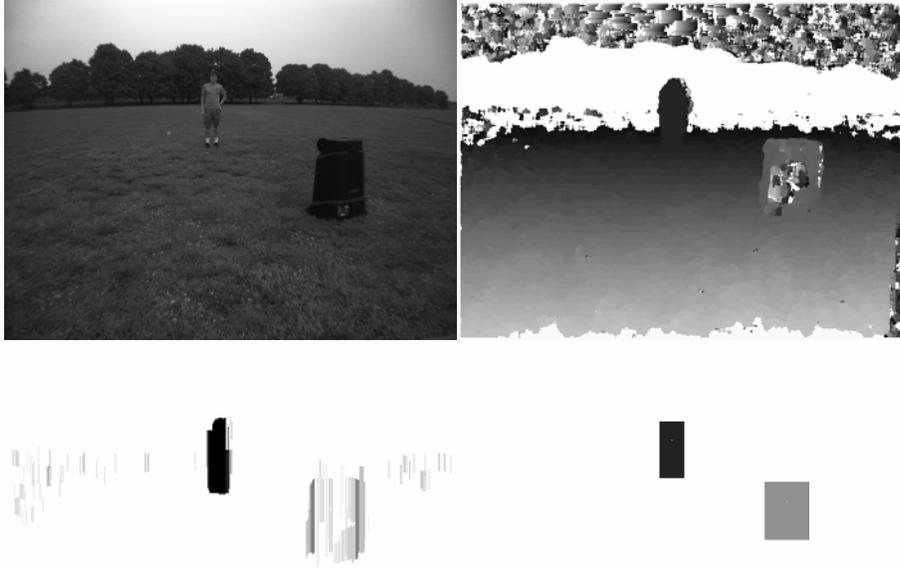


Figure 3.1 Example of stereo vision algorithm. Upper left: Original image; Upper right: Disparity map; Lower left: Processing in columns; Lower right: Bounding with rectangles

4. Accomplishments

4.1 Pre-National Qualifying Event (NQE) Prospect Eleven’s road to the 2005 Challenge was fraught with “do-overs” that cycled between failure and success. The vehicle failed the first site visit during May 2005. Table 4.1 below contains a summary of its performance during the three required and one optional runs. While its performance was poor, it did demonstrate potential and was accorded a second chance as an “Alternate”. An essentially flawless performance during the second site visit on August 16, 2005 as summarized in Table 4.2 earned Prospect Eleven an invitation to the National Qualifying Event (NQE). Included in Table 4.2 are links to depictions of the GPS data and a video documenting the performance.

| 1st Site Visit: May 3, 2005, West Windsor Fields, Princeton, NJ | | | | |
|---|---|--|--|-----------------------------|
| | Run #1 | Run #2 | Run#3 | Optional |
| Course | 220 Meter, S-shaped, 2 obstacles | 220 Meter, S-shaped, 2 obstacles | 220 Meter, S-shaped, 2 obstacles | Serpentine, under trees |
| Obstacles avoided | none | 1 | 1, nipped the other | |
| excursions outside lateral bounds | none | none | none | |
| Competed course | No, emergency stop after hitting 1 st obstacle | Yes, good time | Yes, good time | No, lost GPS under trees |

Table 4.1. Summary of runs during 1st site visit

| 2st Site Visit: August 16, 2005, West Windsor Fields, Princeton,NJ | | | | |
|--|---|---|--|---|
| | Run #1 | Run #2 | Run#3 | Optional |
| Course | 220 Meter, S-shaped, 2 obstacles | 220 Meter, S-shaped, 2 obstacles | 220 Meter, S-shaped, 2 obstacles | 729 Meter serpentine , 2 gates, 6 obstacles & narrow lanes |
| Obstacles avoided | 2 | 2 | 2 | All |
| excursions outside lateral bounds | none | none | none | none |
| Competed course | Yes, 38 seconds data | Yes, 40 seconds data | Yes, good time data | Perfect data ; Movie |

Table 4.2. Summary of runs during 2nd site visit

4.2 Pre-National Qualifying Event (NQE) The NQE involved the running of a [2.2 mile closed serpentine course](#) defined by GPS waypoints with associated lateral boundary off-sets and speed limits that were to constrain the “bots”. Fifty (50) “gates” marked some of the boundaries and five (5) obstacles (four (4) vehicles and a Normandy-style tank trap) were place in the “center” of the desired travel lanes. The course also included a 100 foot-long tunnel (no GPS availability), rumble strips, hay-bale-lined boundaries and a concrete barriers simulated some of the narrow and rough conditions that were to be expected during the Grand Challenge. At the NQE Prospect Eleven made five runs of the 2.2 mile qualifying course. It performed spectacularly well in each odd-numbered run, and disastrously in runs 2 and 4. In fact, its first run may have been the best run of any vehicle at the NQE. Not until after the 3rd run was it realized that Prospect Eleven’s GPS array was misaligned. This caused a biased shift in its perception of the GPS course boundaries. This bias, when combined with the placement of the physical gates meant that many of the gates were perceived to be obstacles within the course boundaries rather than physical markers of the course boundary. Thus, the autonomous path perceived to be available was substantially narrower and had more obstacles than the actual course. The fact that it successfully completed that run, while only nipping two gates and avoiding all other obstacles, was truly remarkable.

The GPS alignment problem and a sluggish response of the braking system caused crashes with the first gate and a parked vehicle and an early termination of the second run. Prior to the third run improvements were made to braking system’s dynamic response and, suspecting bias in GPS, an offset translation was made in the GPS code. This kluge allowed Prospect Eleven to successfully traverse all 50 gates and avoid all 5 obstacles; however, it did nick a couple of the tire stacks.

The 4th run ended in failure due to instability in the steering controller. It turned out that two inappropriate remote processes had been left running on Prospect Eleven’s vision computer. Not long into the run, these processes consumed essentially all of the computing resources, such that the vision computer was processing less than one frame per second, much too sluggish to remain stable.

For the 5th run, “everything” was fixed and Prospect Eleven had an essentially perfect run. It earned a 10th seed in the Grand Challenge Event. Table 4.3 below summarizes the performance and links movies of several of the runs.

| National Qualifying Event, NQE, California Speedway Sept 27-Oct 5, 2005 | | | | | |
|--|---|--|---|--|---|
| <i>Course</i> | <i>2.1 miles, serpentine, 5 obstacles, 50 gates, 1 tunnel</i> | | | | |
| | <i>Run #1</i> | <i>Run #2</i> | <i>Run#3</i> | <i>Run #4</i> | <i>Run #5</i> |
| <i>Obstacles avoided</i> | <i>5 of 5</i> | <i>0 of 5</i> | <i>5 of 5, nipped 2 tire stacks</i> | <i>2 of 5</i> | <i>5 of 5 perfect</i> |
| <i>Gates passed</i> | <i>48 of 50</i> | <i>0 of 50</i> | <i>50 of 50</i> | <i>8 of 50</i> | <i>50 of 50</i> |
| <i>Competed course</i> | <i>Yes, ~14 minutes</i> | <i>No, killed after striking 1st obstacle movie1, movie2,</i> | <i>Yes, ~13minutes</i> | <i>No, killed after steering instability before tunnel</i> | <i>Yes, 12m, 11s movie,</i> |
| <i>Faults</i> | <i>GPS miss-alignment by 2 degrees; discovered after 3rd run</i> | <i>GPS remained miss-aligned, sluggish brake response</i> | <i>GPS remained miss-aligned, brake response enhanced</i> | <i>Several processes inadvertently left running on vision computer</i> | <i>Essentially perfect</i> |

Table 4.3. Summary of runs during 2nd site visit

4.3 Grand Challenge Event (GCE) Twenty-three (23) bots were invited to participate in the Grand Challenge Event on October 8, 2005. Prospect Eleven was seeded 10th. The event involved the traversal of 132 mile course in the desert outside Primm, Nevada. The fastest bot completing the course in less than 10 hours would claim the \$2 million prize. The actual coordinates of the course were not divulged to the teams until 4:00 am on October, 8; 2 hours before the start of the event. Provided was a standard file containing waypoints, lane widths and speed limits. No information was provided on lane roughness nor the location of obstacles, if any. In preparation for the receipt of this file, Josh Herbach'08 had written an analysis program that allowed us to view, evaluate and modify the course. This allowed Josh, Andrew Saxe'08 and I to modify the constraints and to estimate some of the impacts of the modification. While we had conducted one high speed test of Prospect Eleven, we were not confident about its performance at high speed. Since a 35 mph global speed limit allowed for completion in just under 9 hours, we imposed this conservative constraint. Also, we were more confident in Prospect Eleven's stereo vision system than we were in the unbiased accuracy of GPS. Consequently we increased the narrowest lanes to be at least 9 feet wide, thus relying on vision to keep Prospect Eleven on a collision-free path ahead.

Prospect Eleven performed admirably in the Challenge. It launched without difficulty ([see video](#)) and reappeared on schedule at the 8 mile mark ([see video](#)) and passed over the Union Pacific mainline at the 9 mile mark. Unfortunately, the steering became unstable shortly thereafter at about 9.4 miles. A "Personal Best"! After recovery of the vehicle, it was apparent that there was a bug in the obstacle detection code that grew to consume essentially all of the computing resources. One segment of the code failed to dispose of obstacles that had been passed. By the 9th mile, the process was analyzing thousands of obstacles in its effort to determine a collision-free way ahead; thus, it could only update the steering system less than once a second.

In the end, Volkswagen (a.k.a. Stanford) won the Challenge by finishing in just under 7 hours. Carnegie Mellon's two entries, Highlander and Sandstorm, finished 2nd and 3rd in just over 7 hours

and two others, Grey Team and TerraMax completed the course in about 10 hours, but the students of Prospect Eleven were the real winners. The learning and experience were “Priceless”. They knew in their hearts that they had created a phenomenal autonomous vehicle. The only disappointment was in not proving it.

4.4 Unfinished Business Shortly after the team’s return to campus and the placement of Prospect Eleven on in its van “back to Nassau Hall” Bryan Cattle’08 called Wm. Culbreth, Dean of Engineering at UNLV inquiring about temporary storage for Prospect Eleven. The resulting enthusiastic offer of assistance made it clear that Prospect Eleven needed to be turned around in order to give this Prospect Eleven yet another “second chance” to prove its worth. There was unfinished business to take of and the upcoming Fall Break was the time to “Just Do It”. So, on Saturday evening, October 30, 5 team members, Andrew Saxe’08, Gordon Franken’08, Bryan Cattle’07, Anand Atraye’07 and Scott Schiffres’06 and Prof. Kornhauser returned to Las Vegas. Joined by Ben Essenburg’05 on Sunday morning, they found Prospect Eleven at the [Center for Energy Research](#). Using the [shade cast](#) by the Amonix Integrated High Concentration Photovoltaic panel, Anand Atraye set out to find the bug in the code. In a couple of places in the code, he changed one line:

“Everywhere I saw something like:

```
currentObstacles.RemoveAt(i);
```

I preceded it with:

```
State.RelativeFrameUpdated -=  
currentObstacles[i].relativeFrameEventHandler;
```

The problem was that each time an obstacle was removed from the list of current obstacles, it was not being unhooked from the `RelativeFrameUpdated` event. Thus, the garbage collector never determined that the old obstacle was no longer needed, and as a result it was never cleared.” In other words, Prospect Eleven never really “forgot” about an obstacle that it had seen and continued to determine if it needed to avoid it. Thus, after 9+ miles, it was still evaluating obstacles it had detected at the start. The list had grown to thousands of obstacles.

With the code change and a recalibration of the GPS/INU, Prospect Eleven was driven the 35 miles down I-15 to Primm in position for a second chance at the 2005 DARPA Grand Challenge on Monday morning. Table 4.4 below summarizes Prospect Eleven’s autonomous performance over the next three (3) days.

4.4.1 Assault on the 2005 Grand Challenge Course.

Early Monday morning, October 31, 2005, ironically Halloween, we set out to run the 2005 Grand Challenge course exactly as we did during the actual Grand Challenge. Prospect Eleven was using the same RDDF (file of GPS waypoints that define the course) and the same global constraints and control coefficients. The only substantive difference was the change in the “one line of code”. Since we had limited support personnel, comprising of two support vehicles, it was necessary to have someone ride inside Prospect Eleven in case an emergency stop condition was encountered. It



simply wasn't practical to monitor both the environment ahead and the stability of the vehicle entirely from the support vehicles. While possibly dangerous, Professor Kornhauser rode the entire distance behind the wheel of Prospect Eleven, prepared to instantly take over command of the steering and brakes, while in no other way interfering with its autonomous operation. While the plan was for Professor Kornhauser to be alone, Andrew Saxe was permitted to ride "shotgun" during the first stage which traversed a dry lake bed. Since Prospect Eleven had successfully driven this section during the actual Challenge, it was expected to pose few risks. Andrew's job was to toggle the "kill" switches that would enable Prof. Kornhauser to more easily take over control of the vehicle in case of emergency.

Launch came at 7:53:30PST and was uneventful. Everything was perfect until just a few miles into the course when a mirage seemed to appear in the distance. Not to worry, it's the desert; however, it quickly became apparent that the "dry" lake was not so dry. It had rained since the Grand Challenge and the course was not traversable in a non-amphibious vehicle. The decision was to cease autonomous operation in order to not lose the vehicle. A precise autonomous run of the 2005 GC course was infeasible because of the rain. With the current condition, no Grand Challenge vehicle could have made it beyond this point. In fact, if this condition would have existed during the Grand Challenge, DARPA would necessarily have had to alter the course. It now became evident why, during the Grand Challenge, the course was not divulged earlier than 2 hours before the race. It was to ensure that the course was a fair one and that some environmental condition had not made a part of the course impassable.

Rather than go home, the decision was to continue to uncover Prospect Eleven's autonomous operational limits by continuing on the traversable portions of the 2005 GC course. The first autonomous navigation limit had been established: it can't traverse lakes and isn't smart enough to figure out a way around them, if the "desired" course is through them. This has become one of the major goals of our future research.

After a brief diversion around the lake, autonomous operation was reinitiated at reemergence of the 2005 GC course. This incident made it apparent that two people were needed inside the vehicle to properly monitor the road ahead. Other than the lake situation (which occurred at 2 other points), the only non-autonomous diversions were due to

1. places where the "road" had been "bulldozed" probably to discourage exactly what we were trying to do. These places existed at 14:16 and 16:31 (See Figure 6.1), and
2. on NV 604, a public road, where we pulled over to let a cement truck pass us (if this situation would have occurred during the Challenge, DARPA would have paused the vehicle and instructed the cement truck to carefully pass the vehicle).

These two incidents define the operational limits of the current system. Specifically, Prospect Eleven needs the autonomous ability to violate non-critical route constraints and set out to find any feasible path ahead, and it needs to be able to deal with overtaking objects. At present, it does not have these capabilities.

Also, Prospect Eleven was paused several times, much the same way that DARPA may have legitimately paused the vehicle during the Grand Challenge. Pauses were instituted prior to crossing public roads, the Union Pacific at-grade crossing, upon encountering closed gates, that



once opened, were negotiated autonomously and for preparing the onboard camera to record the traverse of Beer Bottle Pass at night.

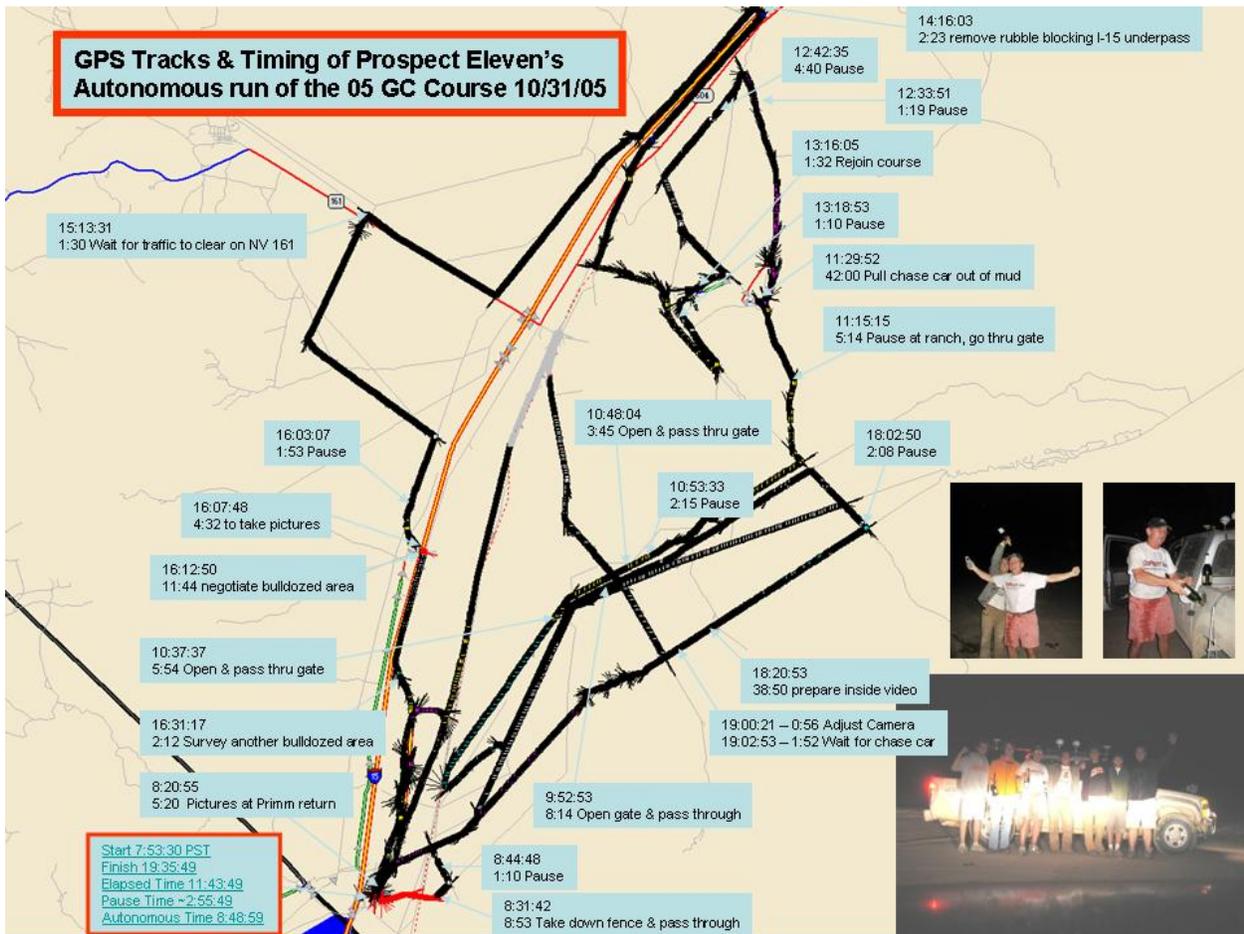


Figure 4.1 GPS tracks and timings of Post Grand Challenge run of 2005 Course

Except for the above constraints, none of which existed during the Grand Challenge, Prospect Eleven autonomously traversed the course. No changes, corrections or alterations were made to any of Prospect Eleven's autonomous systems. Prospect Eleven completed to course in 8 hours 48 minutes and 59 seconds of autonomous time, 11:43:49 of elapsed time. GPS tracks and timings for the run are presented in Figure 4.1. It can be argued that Prospect Eleven autonomously traversed an even more challenging course than that of the 2005 Grand Challenge. Except for the two lakes and the two "bulldozed" areas, Prospect Eleven ran autonomously, including places where the road was significantly rougher than what existed in early October.

What may be argued as being Prospect Eleven's highest achievement during the 2005 run was its 8-minute autonomous descent of Beer Bottle Pass at night beginning at 19:05, See Figure 4.2 below. Only its headlights were illuminating the road ahead for its stereo camera to detected the edges and steer the vehicle down the treacherous terrain. It performed marvelously.

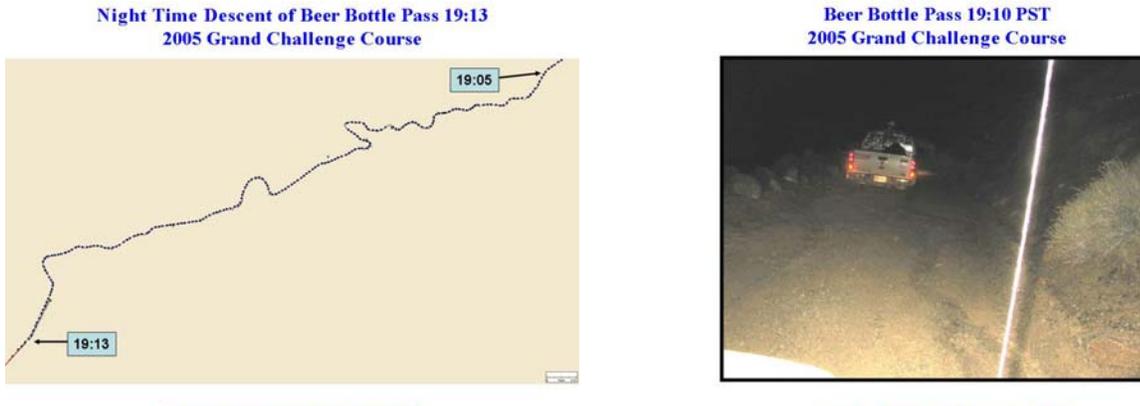


Figure 4.2 GPS tracks and image of night time descent of Beer Bottle Pass on Oct. 31, 2005

6.4.2 Return to Beer Bottle Pass

Because it was so dark the night before, no appreciation was had for what Prospect Eleven had accomplished during its final phases of its run of the 2005 course. We decided to go back to Beer Bottle Pass and rerun it in the daylight. Below are depicted the GPS tracks and an image of our autonomous return to Beer Bottle Pass on Tuesday, November 1. Prospect Eleven first ran up, then back down the Pass autonomously. See Figure 4.3 below. Since we had gotten such a late start only the run up the pass was done in daylight (twilight), the run down had again to be done at night.



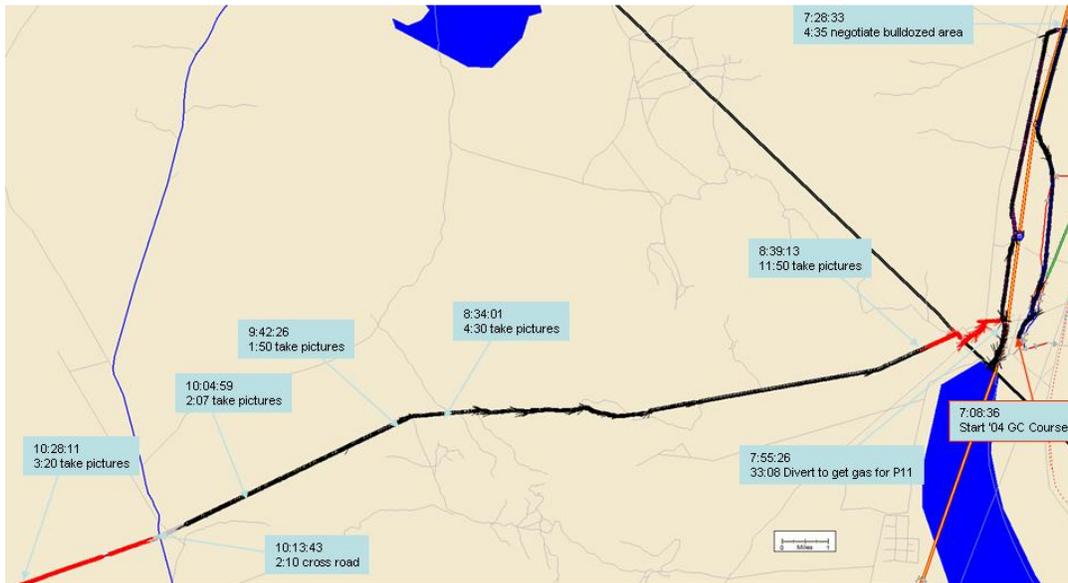
Figure 4.3 GPS tracks and an image of return to Beer Bottle Pass on November 1, 2005

4.4.3 Assault on the 2004 Grand Challenge Course

Given the success of the previous two days, it was decided to make an assault on the 2004 Grand Challenge Course on Wednesday, November 2. Since the vehicle was based in Primm and there was no knowledge on the state of the 2004 course, it was decided to make the attempt in reverse order, from Primm to Barstow by reversing the RDDF of the 2004 Grand Challenge Course and see how far we could go. Prospect Eleven was launched at 7:08:36PST and arrived at the Slash X Ranch at 18:00:00, for an elapsed time of 11h37m43s. It traversed the entire 2004 course in 8:09:00 of autonomous travel. See Figure 4.4 below. While it did not encounter any lakes, manual control did need to be used to negotiate three places where the “road” had been completely washed put and was totally impassable by any vehicle and, at 17:32, to divert around an underpass that had been filled with silt such that there was less than 4 feet of clearance. Also, Prospect Eleven need to be repaired on two occasions. Once at when the steering wheel encoder froze at 13:06 and when the front left tire blew out at 17:16. Finally, it had to be pulled over to the side of the road on two occasions (16:33 and 18:41) to let traffic pass by. Otherwise, Prospect Eleven completed the entire 2004 course in 8:44:190 of autonomous travel. Most impressively, it autonomously traversed

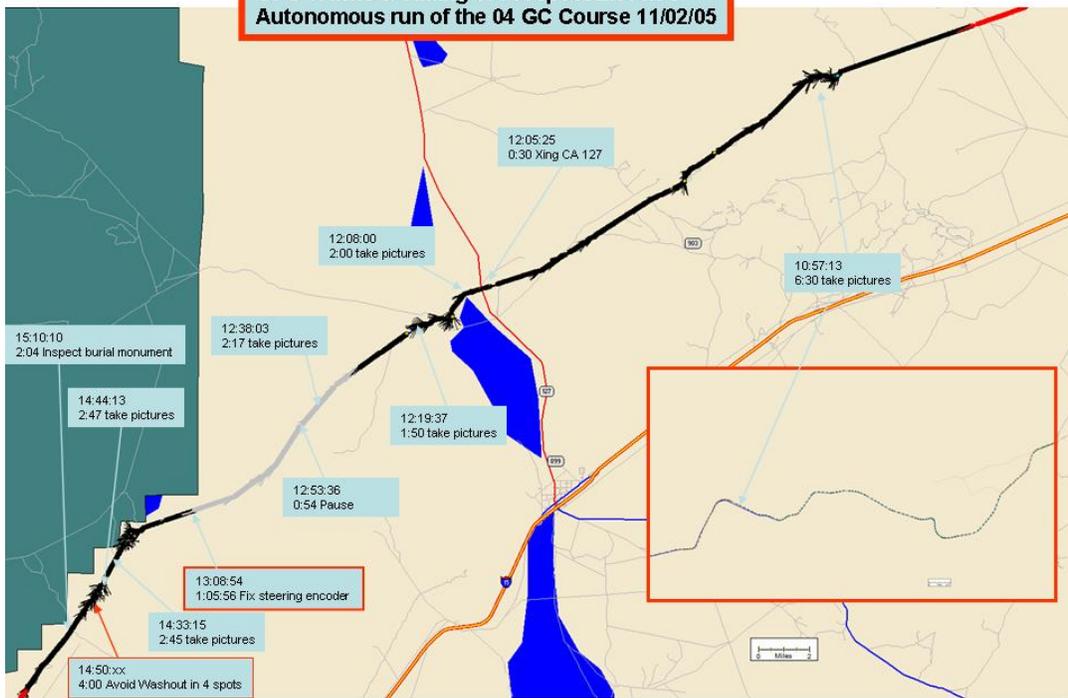
Daggett Pass at night beginning at 18:18, just as it had negotiated Beer Bottle Pass on the previous day.

**GPS Tracks & Timing of Prospect Eleven's
Autonomous run of the 04 GC Course 11/02/05**



DAPRPA Grand Challenge, Unfinished Business, November 2, 2005

**GPS Tracks & Timing of Prospect Eleven's
Autonomous run of the 04 GC Course 11/02/05**



DAPRPA Grand Challenge, Unfinished Business, November 2, 2005

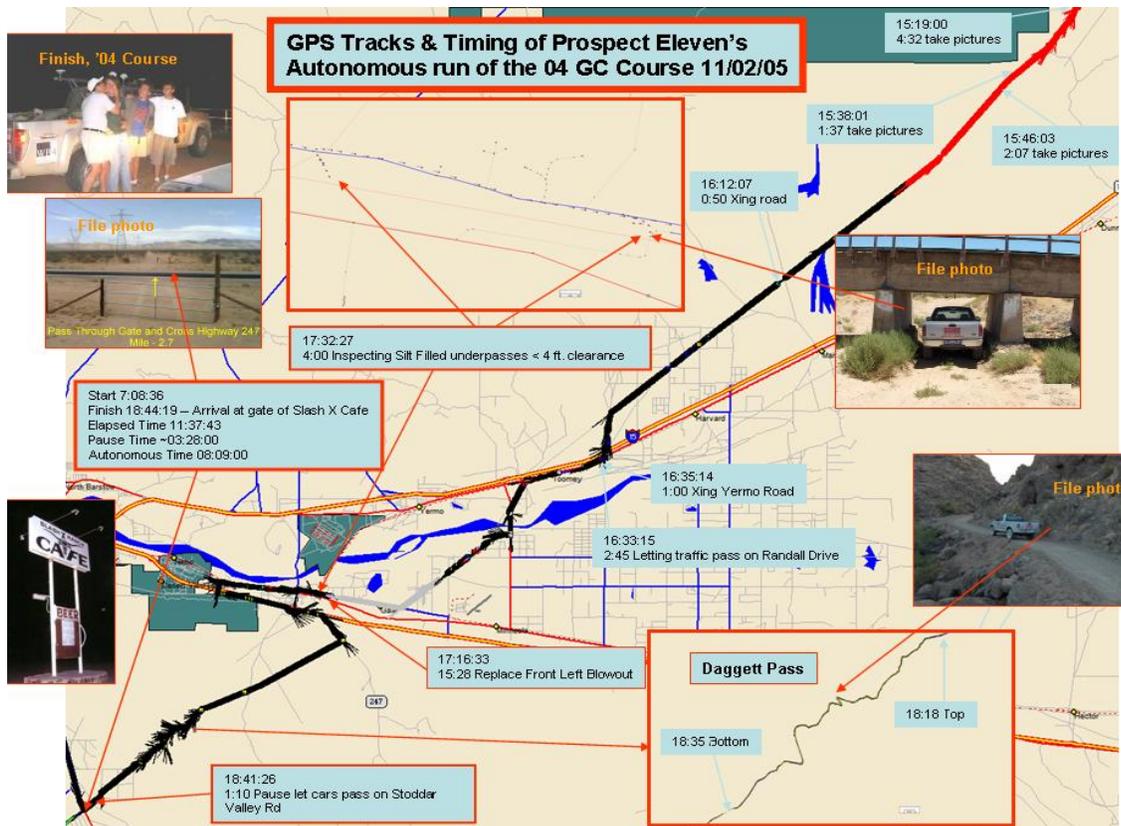


Figure 4.4 GPS tracks and timings of Post Grand Challenge run of 2004 Course

| Unfinished Business, Oct 31 – Nov 2, 2005 , Mojave Desert | | | |
|---|--|-----------------------------|------------------------------|
| Course | 05 Grand Challenge | Return to Beer Bottle Pass | 04 Grand Challenge (reverse) |
| | | | |
| Date | October 31, 2005 | November 1, 2005 | November 2, 2005 |
| Distance | 132 miles | Up & down BB Pass | 143 miles |
| Start time | 7:53:30 PST | 16:54:27 (up) 17:10:17 (dn) | 7:08:36 PST |
| Finish | 19:36:49 PST | 17:04:48 (up) 17:18:04 (dn) | 18:44:19 |
| Elapsed time | 11:43:49 | 0:10:21 (up) 0:7:47 (dn) | 11:37:43 |
| Autonomous time | 8:48:59 | same | 8:09:07 |
| Note on 2005 run | Completely autonomous with no failures except avoidance of “wet” dry lake & bulldozed segments | | |
| Note on Return 2BB | Completely autonomous with no failures except struck outcropping bolder with left front wheel after decent was completed. Caused tire blowout and substantial front wheel misalignment | | |
| Note on 20054 Run | Two failures: steering wheel encoder failure (repaired) and blowout of left front wheel caused due to misalignment. Manual avoidance of complete washouts of road, diversion around silt-filled underpass and pauses to let motorized traffic pass | | |
| Movies | PGC05Gate1 , PGC05Gate2 , PGC05Crusin' , PGCBB+04 , PGC04 | | |

Table 4.4 Summary of runs during Unfinished Business

5. Final thoughts

Prospect Eleven turned out to be a phenomenal autonomous vehicle. It performed enormously better than we had any right to expect. The stereo vision system performed though a wide range of illuminations from bright daytime conditions through dark nighttime. It was particularly well suited for collision avoidance and where the appropriate travel speed was between 7 miles per hour and 25 miles per hour. It was not tuned to perform well in really treacherous regions such as heavily washed out segments where a much slower speed is required for a vehicle to pick its way through. The limited range of its vision also did not allow it to confidently travel at higher speeds although the steering system was tested and found to be stable to at least 65 miles per hour, the highest speed at which it was tested. Future work on the vehicle will focus on enhancing its performance in both the low speed and high speed regimes.

6. References

1. Talukder, A., Manduchi, R., Matthies, L., & Rankin, A. (2002) "Fast and Reliable Obstacle Detection and Segmentation for cross country navigation", in IEEE Intelligent Vehicles 2002.
2. van der Mark, W., Groen, F.C.A., & van den Heuvel, J. (2001) "Stereo based navigation in unstructured environments", in Proc. IEEE Instrumentation and Measurement Conference, 2001.
3. Broggi, A., Caraffi, C., Fedriga, R., & Grisleri, P. (2005) "Obstacle Detection with Stereo Vision for off-road Vehicle Navigation", in Procs. Intl. IEEE Wks. on Machine Vision for Intelligent Vehicles, 2005.
4. Trucco, E. & Verri, A., (1998) *Introductory Techniques for 3-D Computer Vision*. Prentice Hall PTR, 1998.
5. Forsyth, D. & Ponce, J., (2003) "Computer Vision – A Modern Approach," Prentice Hall 2003
6. Kornhauser, A. et al, (2006) "Prospect Eleven, Princeton University's entry in the 2005 DARPA grand Challenge", Submitted for publication in Journal of Field Robotics, Jan. 2006