Carnegie Mellon University

School of Computer Science

Undergraduate Honors Thesis

Non-holonomic Trajectory Planning for High Speed Navigation

Matthew Johnson-Roberson

Advisor William Whittaker

Abstract

In this thesis a new approach to high-speed navigation is presented. The limitations of traditional path tracking techniques have negative implications for the high-speed traversal of rugged terrain. Through the investigation of trajectory generation that is reactive and stable, the requirements for path tracking and obstacle avoidance at high speed can be determined. A reactive swerving approach, as implemented, allows autonomous vehicles to achieve higher speeds on rougher terrain. An analysis of the implementation of such an approach demonstrates the importance of dynamics and terrain-vehicle interactions at high speed. This approach to high speed navigation offers the possibility of broadening the application and use of autonomous vehicles in the world.

Introduction

High-speed navigation presents many challenging problems for an autonomous system. Robots currently lack the ability to operate at high speeds in semi-structured environments. As speed is increased, traditional approaches to autonomous vehicle navigation in off-road environments present many problems. The ability of an autonomous vehicle to travel quickly is dependent on its ability to assess terrain and choose a safe trajectory through that terrain until eventually reaching its goal.

The margin and tolerance of error at high speeds is small. Path following and obstacle avoidance as they stand today don't meet these strict requirements. The amount of tracking error generated by traditional techniques is insufficient to drive at racing speeds. The need for the vehicle to stay in the safe region (as determined by the perception) is of paramount importance at high speed. The safety of the vehicle is a function of the perception's ability to detect a safe path, but it is also compounded by the navigation's ability to execute that path and when infeasible react in a way that is safe. Part of the challenge of high-speed navigation is determining what is and is not safe and achievable.

While sensing technology has developed to allow for a longer sensing horizon, planning and execution must rise to meet the challenge for high speed navigation. The proposed technique addresses solutions for both. In addition, the technique suggests that improvement of short-range reactive navigation will make robots safer, faster, and more reliable. It helps to reduce tracker error and motion through terrain that is not on the path and it also allows for obstacle avoidance at high speeds in emergency situations. It purposes a framework to determine the safety of possible vehicle trajectories for such obstacle avoidance. This framework addresses not only obstacles but also dynamic hazards such as rollover and projectile motion. These techniques combine to greatly improve the performance of autonomous navigation and thus allow for greater speed and complexity of terrain.

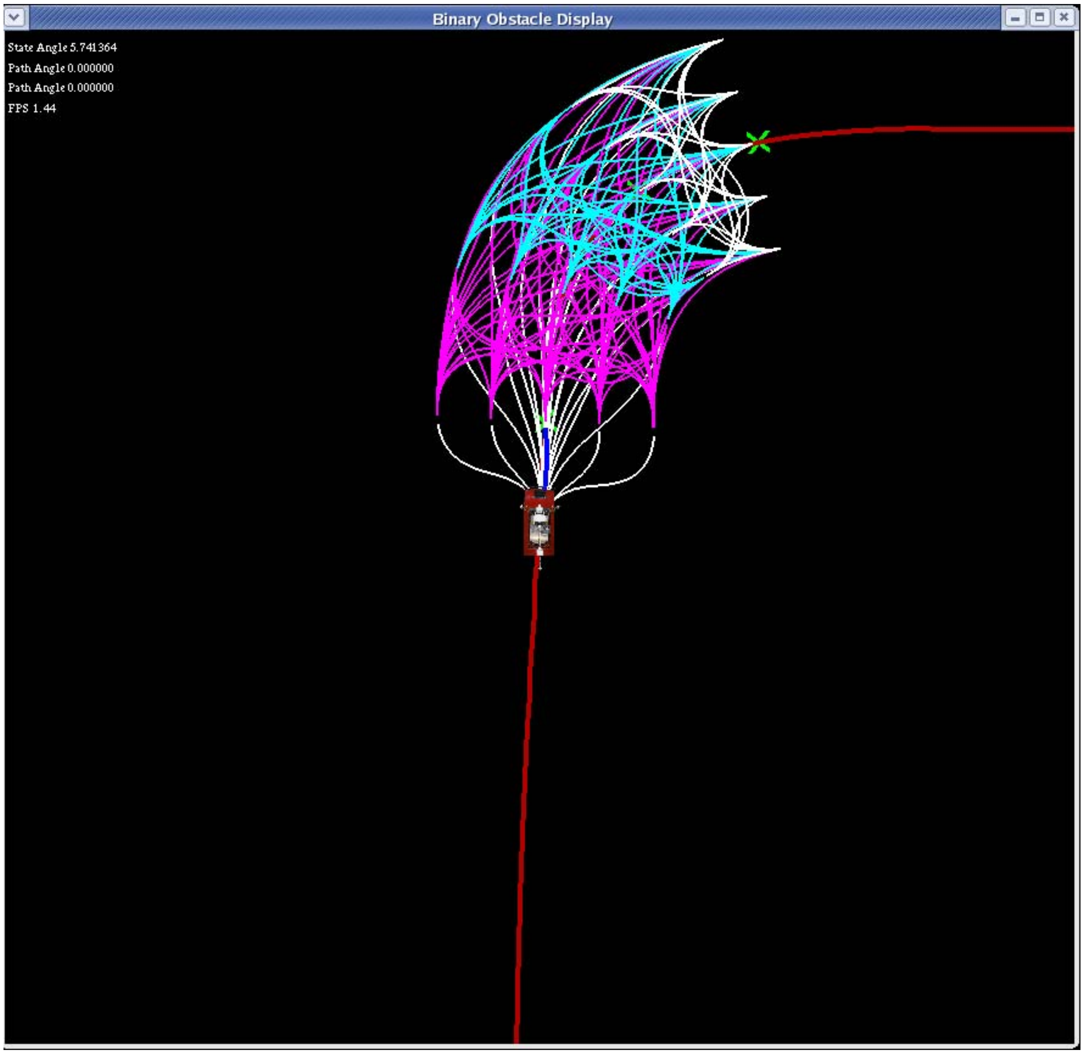

State Angle 5.741364
Path Angle 0.000000
Path Angle 0.000000
FPS 1.44

Figure 1: By considering all trajectories from the robot's current pose that end on the boundaries of the road it is possible to check the feasibility of all such plans with the sensor map of obstacles and road and course definitions.

Context

This technology is being developed to compete in the 2005 DARPA Grand Challenge. The Grand Challenge is a driverless desert competition with a $2 million dollar prize. The goal of this competition is to develop an autonomous vehicle that can traverse 175 miles of desert terrain with no human intervention in under ten hours. The exact route will be revealed to the competitors two hours before the race. After the robot has left the starting line there can be no further communication with the vehicle. Our team competed in last year’s challenge with the strongest showing, but were unable to complete the course and are returning this year to achieve victory.

The test platform is a 1986 model 998 HMMWV (High Mobility Multi-purpose Wheeled Vehicle). It has been fully upgraded to support drive-by-wire through the modification of acceleration, braking and shifting. Emergency break modification allows for safe autonomous operation through the use of a wireless electronic kill. The vehicle is equipped with six laser range finders, one of which is actively stabilized on a gimbal; these provide a comprehensive terrain model. The vehicle also has an integrated GPS/INS system to provide accurate pose information.

The Grand Challenge is expanding the frontier field of high speed navigation of unrehearsed terrain. Autonomous robots have little experience in traveling fast in unstructured environments. Successful approaches have allowed for high speed navigation of simple well defined environments like roads [8][14]. Very complex unknown terrain has also been explored by Mars Rovers and slow car-like vehicles. [2][5][10][13]. The Grand Challenge has motivated this work by pushing not only the speed of autonomous navigation but also increasing the difficulty of the terrain to be traversed.

Previous Work/Background

High-speed navigation has been previously approached as a path-following problem where many aspects are simplified to allow for the speed of computation and stability necessary. At speeds where dynamic instability becomes an issue, techniques must be safe to be effective. Early work based on the Global Positioning System opened an entire field of applications to robotics. Path following with safeguards was utilized with mining dump trucks at high speeds in simple environments. [11] Previously pure pursuit was implemented to path follow on mobile robots. [4][1]. Pure pursuit was utilized during the Grand Challenge on Sandstorm for its robustness, but it also limited its performance because of the increasingly dilatory effect of tracker error as speed increases.

Polynomial trajectories offer the promise of reducing tracker error by improving the fidelity of the model of the vehicle’s motion on a trajectory. By modeling the steering actuator they more accurately represent the way the vehicle will respond to an issued steering command. By reducing the error along the path the vehicle is safer and can achieve higher speeds. Starting with clothoid curves, which model constant curvature change, they have evolved into higher order polynomials which model constant rates of change of curvature. These evolutions result in a reduction of tracker error around turns [6] [7].

Clothoid trajectories have also been utilized to navigate vehicles [9]. This approach was adapted for high speed [3]. A pre-computed set of clothoid trajectories were generated and a voting scheme consisting of checking the path of each clothoid with an obstacle map allowed for quick cancellation of unsafe motions. This approach successfully allowed for the navigation of terrain with large obstacles at 10 m/s, but lacked the ability to go faster or deal with more complex terrain.

High speed navigation is also being attempted with a stochastic approach based on simulation. Dynamic terrain interaction based on statistical models of speed and curvatures through space is being developed [12]. Nominal trajectories are generated, and then checked against the sensor generated terrain model. If they are not viable a search is done to find a trajectory that is possible to execute on the terrain. Still in development, the navigation has only occurred in simulation but showed promising results.

No current approach has allowed for both high speeds and the safe traversal of extremely complex terrain that would expand robotic applications.

Methodology

Building upon the previous approaches to path following and obstacle avoidance to create a hybrid technique that works for the domain of high speed navigation in a semi structured environment, will allow for stable safe operation of autonomous vehicles in more hazardous environments. To achieve this end trajectories must be generated quickly and reliably. Such paths must be generated, adapted, and followed based upon real time perceptual information. These paths must be tracked with less error then previous techniques to ensure the safety necessary for high speeds.

Assumptions about the nature of the terrain allow for quick analysis from limited sensor data. Combining a priori data with a simple local model of the world it is possible to create coarse trail boundaries which roughly correspond to the safe areas for traversal. A trail based model of the world has proved effective in traversing the terrain of the open desert. Attempting to center the vehicle on the trail abstracts away many of the issues that are still open problems in exploratory off road navigation.

By generating plans emanating from the vehicle in trajectory space, it is possible to more accurately determine the path of the vehicle for a set of actions between the current and next planning cycle. These actions can be vetoed based upon their intersection with obstacles or their motion outside the boundaries of the trail. The search space is generated live based upon the constraints of the path ahead. A set of cubic trajectories is created originating from the vehicles current position and ending at a look-ahead point along the path varying based upon speed. This creates a trajectory between an initial posture (x,y,θ,κ) postion heading and curvature. And a final (x,y,θ,κ) allows for a smooth cubic function mapping steering curvature to arc length or s with $\kappa(s) = a + bs + cs^2 + ds^3$

$$\kappa(s) = \kappa_0 + as + bs^2 + cs^3$$

$$\theta(s) = \theta_0 + \int_0^S \kappa(s) = \theta_0 + \kappa_0 s + \frac{as^2}{2} + \frac{bs^3}{3} + \frac{cs^4}{4}$$

$$x(s) = x_0 + \int_0^S \cos(\theta(s)) \qquad y(s) = y_0 + \int_0^S \sin(\theta(s))$$

The initial parameters are set to the vehicle's state at the beginning of the planning cycle. The final parameters are generated by taking the course road model and extracting a curvature for the segment of interest.

Distance equation:

$$d(x_i, y_i) = \sqrt{(x_i - x)^2 + (y_i - y)^2} - r$$

Objective function:

$$J(x, y, r) = \sum \left(\sqrt{(x_i - x)^2 + (y_i - y)^2} - r \right)^2$$

Using circle fitting a curvature is calculated for the preplanned path ahead. The terminal heading is set to be facing the next way point.

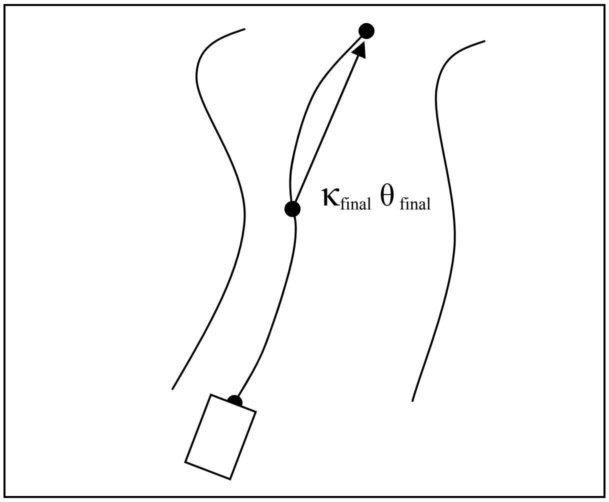

Figure 2: Path Heading Calculation

From this initial curve swerving options are generated which allow for short range reactive obstacle avoidance. These are generated by using the road boundaries based upon the tangent to the direction of the path

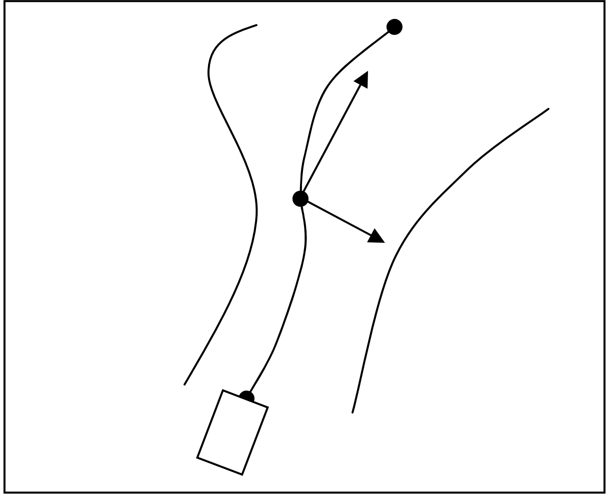

Figure 3: Path Tangent Calculation

With this array of possible motions several metrics are used to select the one to be followed. Prior to any selection the curves are laid into an obstacle map and canceled based upon their intersection with anything the perception system has deemed impassible.

A traditional arc voter then integrates the cost along the curve by stepping through the points of the arc and summing the cost per cell. Curve selection uses a simple heuristic of distance to the preplanned path. The technique utilizes distance from the road boundary and distance to the path as determined by the high-level planner. We begin by stepping along the curve and the path based upon arc length, so as to compare the most relevant distances between the curves and the path. These distances are summed giving a total value for the curve. Because only a portion of the curve is followed, the region of interest for path distance is a fraction of total curve length. Therefore we only compare the curve to the path in a limited region. Additionally, costs along the curve are ignored if they are under the obstacle threshold. This is done because part of the methodology of this high speed approach is that small differences in cost as unimportant when it comes to choosing a safe path. Traditional arc voters will compare and aggregate the costs along the path, but such methods afford little advantage in this situation. By sacrificing true optimality this tracker adds stability and thus increases the speed achievable. This allows for greater freedom to move along the road and avoid the edges as well as obstacles.

The preplanned path and speeds are assumed to be safe if kinematically achievable. Since the high-level planner only has a coarse knowledge of the vehicles motion some smoothing must be done to actually steer the vehicle. Traditional path-following approaches such as pure pursuit assume that such smoothing will not result in a true path that collides with an obstacle. By searching the trajectory space and mapping that to obstacle-laden map space that assumption can be avoided.

Upon selection of a valid curve, sample the $\kappa(s)$ along the curve based upon distance traveled. That distance is computed after each sampling $\kappa(s)$. The steering controller is asked to track the curve by updating its curvature at its maximum rate.

Results

Both pure pursuit and the current implementation aboard the robot were run in simulation and compared at various speeds on a Z course as shown in Figure 4 to note tracker error over significant distances and the repeatability of path following. This course, with high curvature turns, was chosen to examine the new technique's performance in minimizing overshoot when going around tight corners. The course, along with the old tracker's results, appear in Figure 6. The new tracker's results are shown in Figure 5. The course was run at 4 m/s, 5 m/s, and 6 m/s.

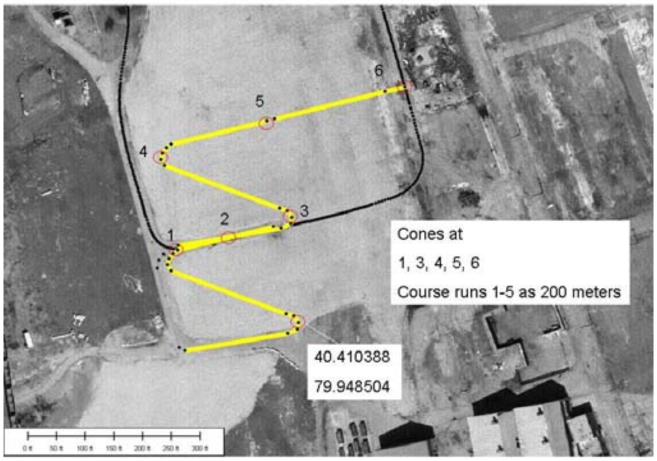

Figure 4: Z Course at LTV Test Site in Pittsburgh

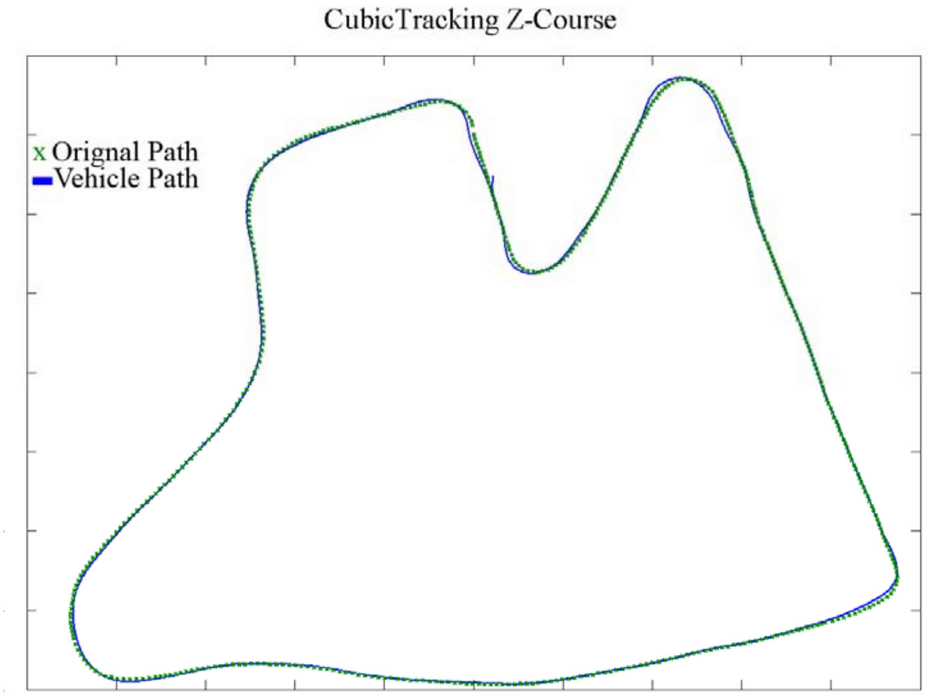

Figure 5: Cubic Tracking Results on Z Course

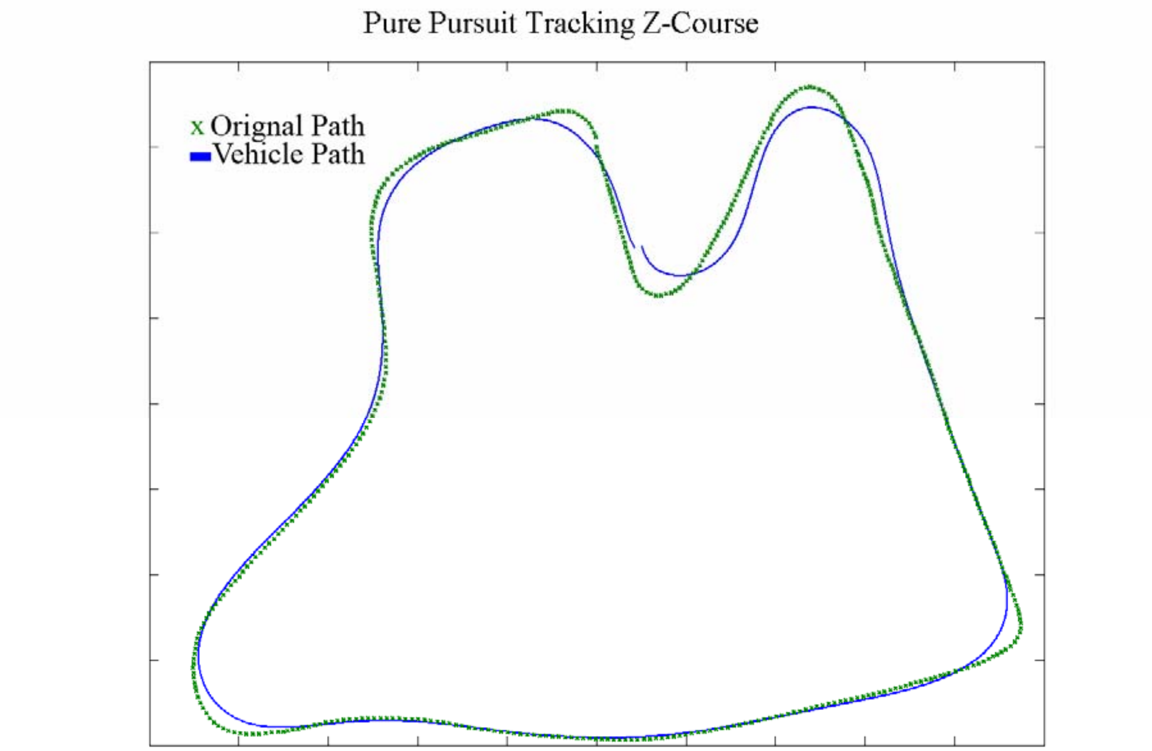

Figure 6: Pure Pursuit Results on Z Course

As indicated in the graphs, along the turns pure pursuit begins to cut inside the turn as it is approaching the middle of the curve. This creates significant tracker error in the corners of turns that would be catastrophic if the path were on a cliff edge or similar geometry. The new tracker minimizes this error and centers the vehicle on the path for the entire turn.

The next battery of tests was serving obstacle avoidance. These runs are characteristic of detecting an obstacle very close to the vehicle at high speeds and having to react to it quickly and safely. To magnify the results the sensor horizon has be artificially brought in to simulate detecting the obstacle later. The tests were run at 10 m/s and the sensor horizon was tested at 5m and 15m.

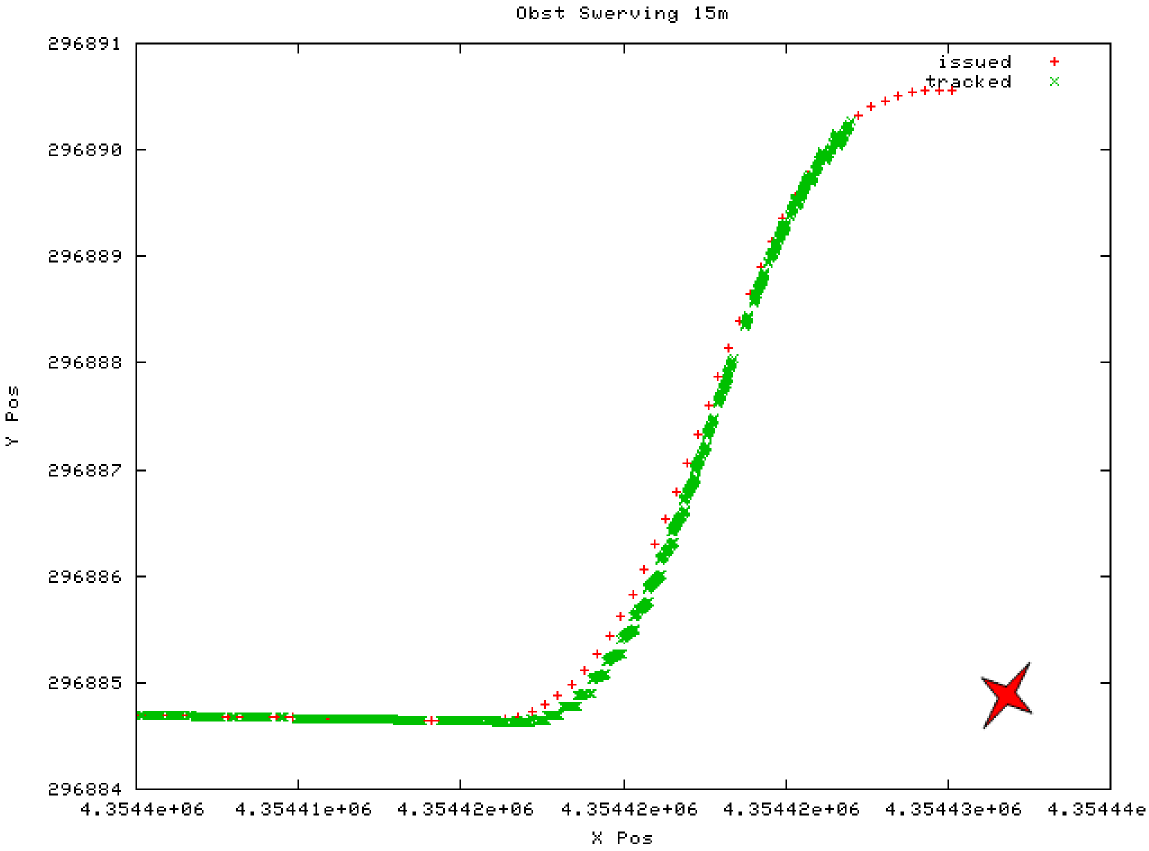

Figure 7: Obstacle detected 15m in front of the vehicle

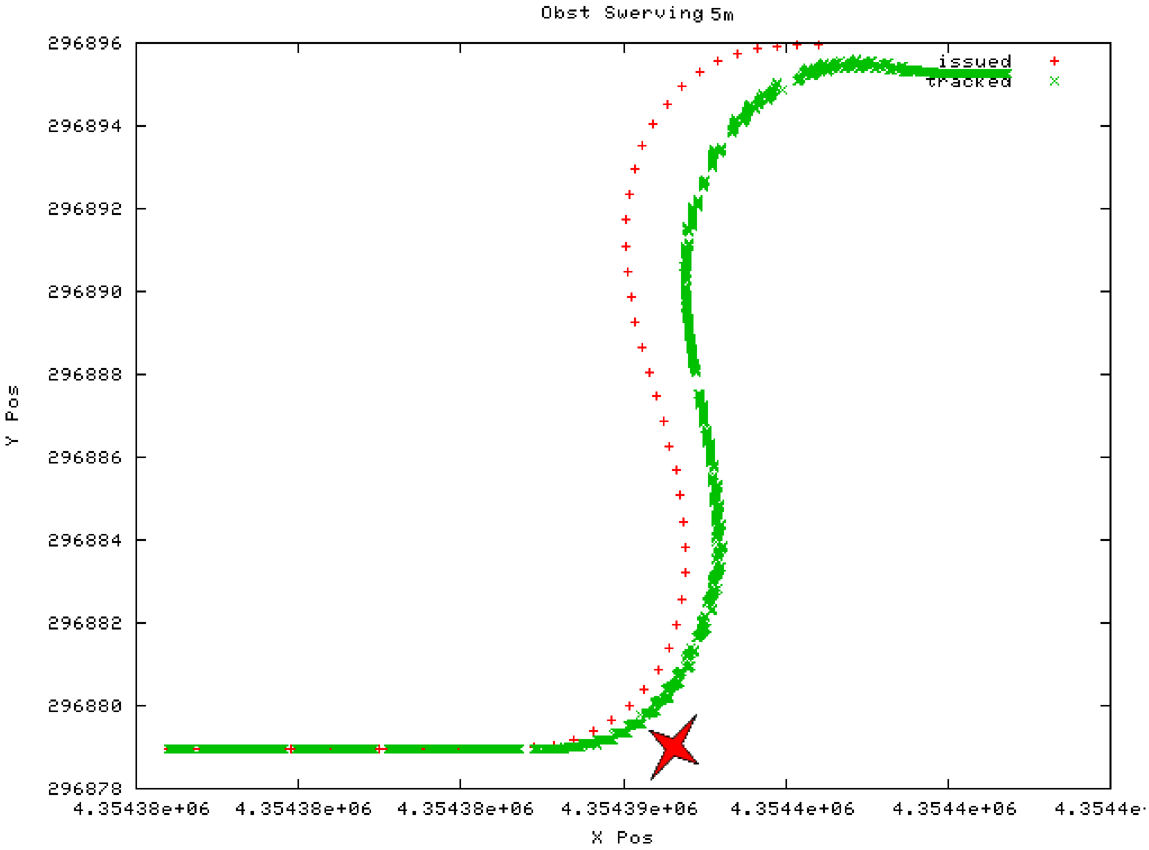

Figure 8: Obstacle detected 5m in front of the vehicle

The vehicle is able to successfully dodge the obstacle with only 5m of space to issue and execute the reactive trajectory. The point at which the obstacle was detected and the path of the vehicle are shown in Figures 7 and 8. With a 5m sensor horizon the vehicle has .5 seconds to determine the safe path and execute it. The longer sensor horizon shows that the technique is able to generate a smoother obstacle avoidance trajectory when it has more time to plan. If the speed was increased to above 10 m/s some of the trajectories would make the vehicle unstable indicating the need for simulation of trajectories to determine their dynamic feasibility.

Future Work

Dynamic Simulation

Dynamic Simulation of the trajectories at each planning step offers the possibility of determining dynamic effects such as rollover, slipping/skidding, and projectile motion. Work is currently proceeding to simulate the robot over LIDAR-generated terrain maps. While computationally very expensive, the possible benefits are great. Based upon the fidelity of the simulation much of the tracker error could be reduced. As speed increases, the need for modeling of dynamic effects becomes increasingly important. Because the vehicle is constantly very close to instability, safeguarding must be implemented to combat the possibility of the vehicle losing control because of terrain or commanded trajectory.

3D Trajectories

Work has been done to expand the polynomial curves to account for z parameters and therefore stretch the curve across a 3d terrain. This could aid the tracking and following of curves over rough terrain, again reducing tracker error. The ease of utilizing such curves in the current framework makes it very viable to extend this method to account for terrain height changes. This technique, much like dynamic simulation, will only be limited by the ability to generate height maps that are complete and accurate from sensors at high speeds. As that ability is explored, the viability of the aforementioned techniques will become clearer.

Conclusion

Robust path tracking is essential to the success of high-speed autonomous ground vehicles. The results from the swerving tests reveal the need for longer sensor horizons. The robot is much more likely to be successful if it can generate safe trajectories that are far from the vehicle’s zone of instability. This ability is directly proportional to the distance the obstacle is detected in front of the vehicle. The trajectories that are generated when the obstacle is detected very close to the vehicle are naturally unstable. Out of this research two things have become evident. Sensors and algorithms for detecting obstacles are a limiting factor for high speed navigation. Even the state of the art in sensing technology cannot always detect obstacles in time to make a smooth avoidance trajectory. This result stresses the importance of a nimble platform which can make sharp evasive maneuvers at speed and a detailed model of the vehicle rollover and slippage characteristics. These two advances will allow for a true step forward in the speeds achievable by autonomous vehicles.

References

[1] O. Amidi "Integrated Mobile Robot Control", Technical Report CMU-RI-TR-90-17, Robotics Institute, Carnegie Mellon University, Pittsburgh, PA, May 1990.

[2] J. Biesiadecki, M. Maimone & J. Morrison. "The Athena SDM Rover: a Testbed for Mars Rover Mobility", Proc. i-SAIRAS 2001, St-Hubert, Canada, June, 2001.

[3] D. Coombs, K. Murphy, A. Lacaze & S. Legowik. "Driving Autonomously Offroad up to 35km/h", Proc. IEEE Intelligent Vehicles Symposium, Dearborn USA, 2000.

[4] R. Coulter "Implementation of the Pure Pursuit Path Tracking Algorithm", Technical Report CMU-RI-TR-92-01, Robotics Institute, Carnegie Mellon University, Pittsburgh, PA, May 1992.

[5] S. Golberg, M. Maimone & L. Matthies. "Stereo Vision and Rover Navigation Software for Planetary Exploration", In Proceedings of the IEEE Aerospace Conference, Big Sky, USA, March 2002.

[6] Kelly, A., Nagy, B. "Reactive Nonholonomic Trajectory Generation via Parametric Optimal Control", to appear International Journal of Robotics Research", International Journal of Robotics Research, Vol (No), 2002..

[7] Nagy, B., Kelly, A., "Trajectory Generation for Car-Like Robots Using Cubic Curvature Polynomials", Field and Service Robots 2001 (FSR 01), Helsinki, Finland - June 11, 2001.

[8] D. Pomerleau, "RALPH: Rapidly Adapting Lateral Position Handler", IEEE Symposium on Intelligent Vehicles, September, 1995, pp. 506 - 511.

[9] D.H. Shin and S. Singh tech. report CMU-RI-TR-90-31, Robotics Institute, Carnegie Mellon University, December, 1990.

[10] R. Simmons, E. Krotkov, L. Chrisman, F. Cozman, R. Goodwin, M. Hebert, L. Katragadda, S. Koenig, G. Krishnaswamy, Y. Shinoda, W. Whittaker, & P. Klarer. "Experience with Rover Navigation for Lunar-Like Terrains", Proc. IEEE IROS, 1995.

[11] S. Singh, D. Feng, P. Keller, G. Shaffer, W. Shi, D.H. Shin, J. West, and B.X. Wu, "A System for Fast Navigation of Autonomous Vehicles", Technical Report CMU-RI-TR-91-20, Robotics Institute, Carnegie Mellon University, September, 1991.

[12] Spenko, M., Iagnemma, K., and Dubowsky, S "High Speed Hazard Avoidance for Mobile Robots in Rough Terrain", Proc. SPIE Conference on Unmanned Ground Vehicle Technology VI, Orlando, FL, Vol. 5422, April 2004.

[13] A. Stentz & M. Hebert. "A Complete Navigation System for Goal Acquisition in Unknown Environments", IEEE IROS, 1995.

[14] C. Thorpe, T. Jochem, and D. Pomerleau. "The 1997 Automated Highway Free Agent Demonstration", IEEE Conference on Intelligent Transportation Systems, November, 1997, pp. 496 - 501.