Case Study

Autonomous Cars

- Sebastian Thrun’s TED Talk describing Google’s driverless car
  - www.ted.com/talks/sebastian_thrun_google_s_driverless_car.html

- Junior: The Stanford Entry in the Urban Challenge
  - http://robots.stanford.edu/papers/junior08.pdf

DARPA Urban Challenge

- November 3, 2007
- 6 hours to complete a 96 km urban area course
- multiple robotic vehicles carrying out missions on the same course at the same time
- basic rules:
  - stock vehicle
  - obey California driving laws
  - entirely autonomous
  - avoid collisions with objects typical to urban environment
  - must also be able to operate in parking lots
  - DARPA supplied environment map with information on lanes, lane markings, stop signs, parking lots, and special checkpoints

Junior

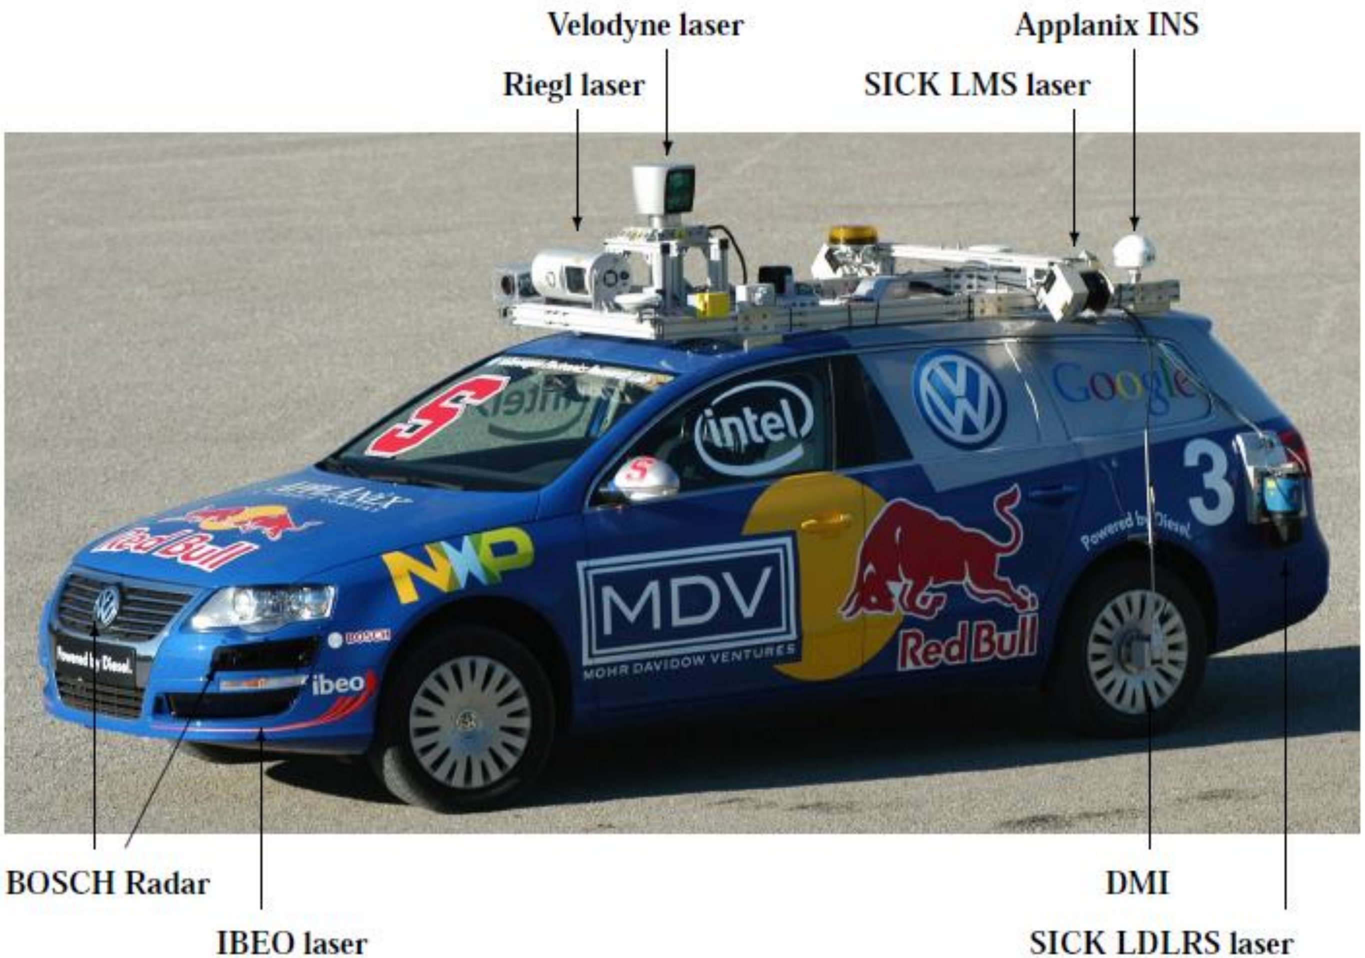

Figure 1: Junior, our entry in the DARPA Urban Challenge. Junior is equipped with five different laser measurement systems, a multi-radar assembly, and a multi-signal inertial navigation system, as shown in this figure.

Junior: Computation

Figure 2: All computing and power equipment is placed in the trunk of the vehicle. Two Intel quad core computers (bottom right) run the bulk of all vehicle software. Other modules in the trunk rack include a power server for selectively powering individual vehicle components, and various modules concerned with drive-by-wire and GPS navigation. A 6 DOF inertial measurement unit is also mounted in the trunk of the vehicle, near the rear axle.

Junior: Sensors

- Applanix POS LV 420 position and orientation system
  - multiple GPS receivers, GPS heading measurement, inertial measurement, distance measurement, Omnistar Virtual Base Station
    - position and orientation errors less than 100 cm and 0.1 degrees
- 2 side facing SICK LMS laser range finders
- 1 forward facing RIEGL LMS-Q120 laser range finder
- 1 roof mounted Velodyne HDL-64E laser range finder
- 2 rear mounted SICK LDLRS laser range finders
- 2 front mounted IBEO ALASCA XT laser range finders
- 5 BOSCH radars mounted in front grill

Laser Obstacle Detection

- challenging!
  - curbs
  - moving and static obstacles
  - overhanging obstacles (tree branches, signs, etc.) that can be safely driven under

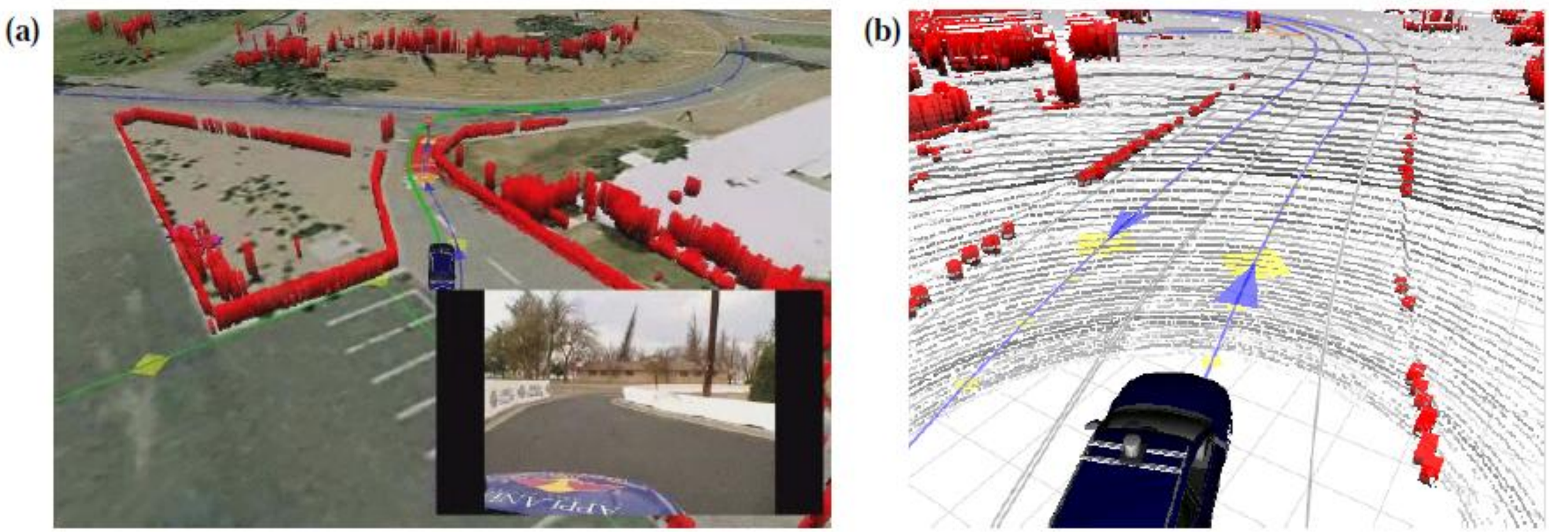

Figure 5: Obstacles detected by the vehicle are overlayed over aerial imagery (left) and Velodyne data (right). In the example on the right, the curbs along both sides of the road are detected.

Laser Obstacle Detection

- primary sensor is the Velodyne laser range finder
  - cannot reliably detect curb-sized obstacles near the vehicle because of self-occlusion

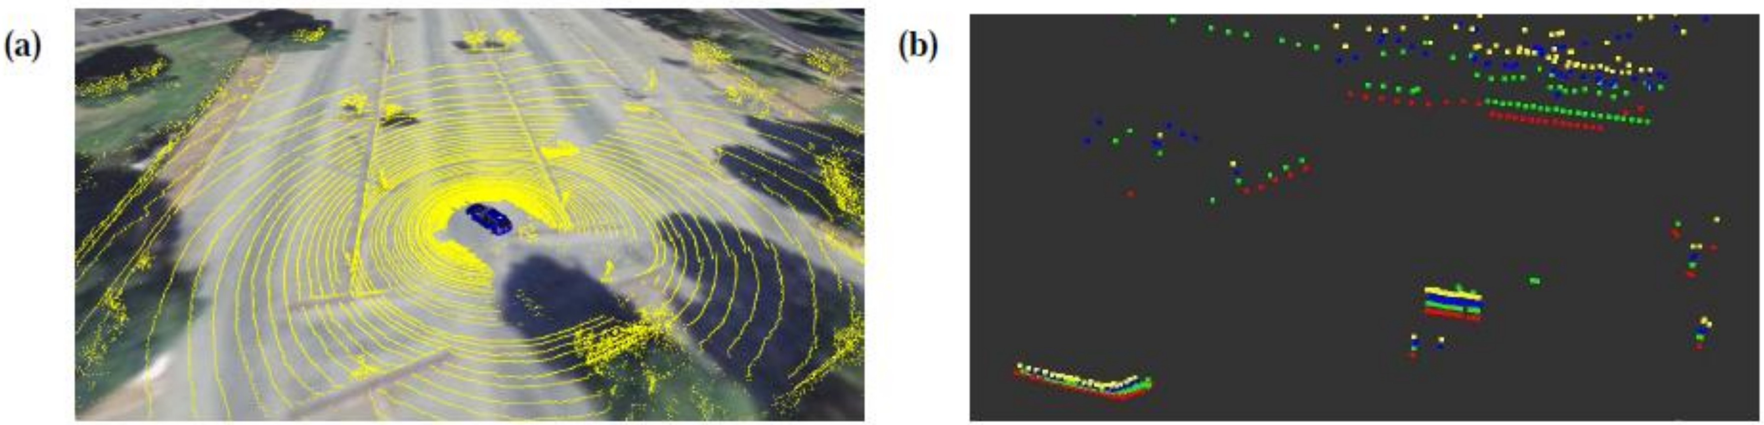

Figure 4: (a) The Velodyne contains 64 laser sensors and rotates at 10 Hz. It is able to see objects and terrain out to 60 meters in every direction. (b) The IBEO sensor possesses four scan lines which are primarily parallel to the ground. The IBEO is capable of detecting large vertical obstacles, such as cars and signposts.

Laser Obstacle Detection

- curbs near the vehicle detected using the 2 front facing IBEO laser range finders and 2 rear facing SICK laser range finders
- only obstacles close to the vehicle (5m in front and 15m in back) considered

Static Mapping

- discrete grid-based local maps based on laser scan data

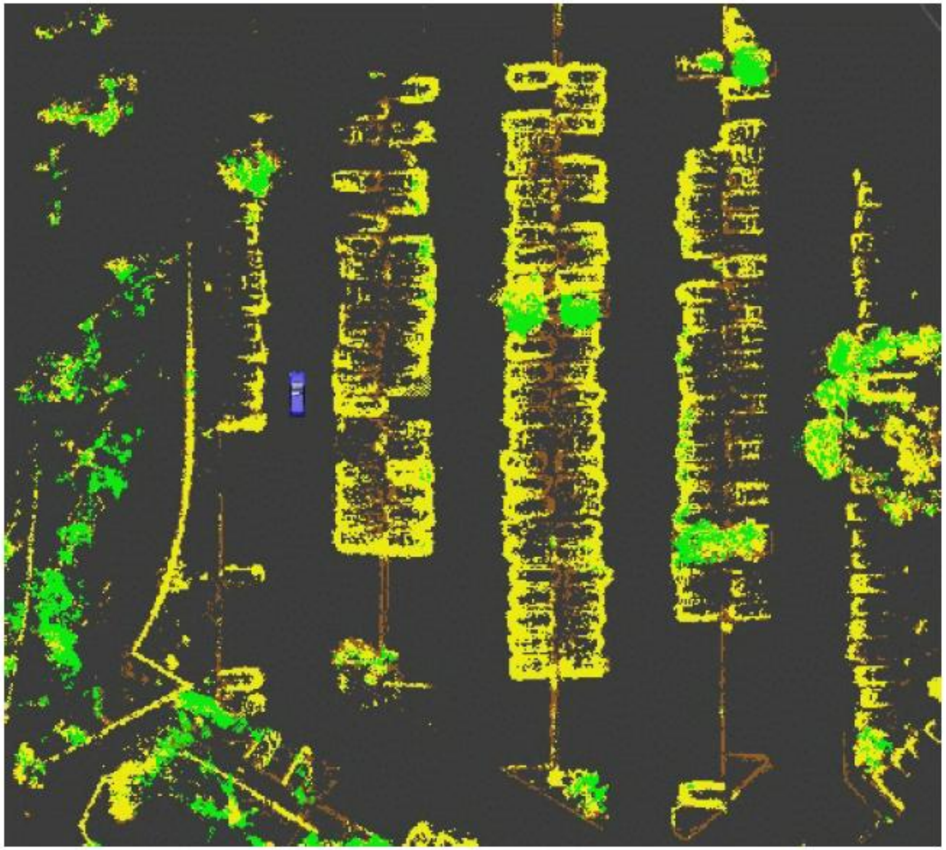

Figure 6: A map of a parking lot. Obstacles colored in yellow are tall obstacles, brown obstacles are curbs, and green obstacles are overhanging objects (e.g. tree branches) that are of no relevance to ground navigation.

Static Mapping

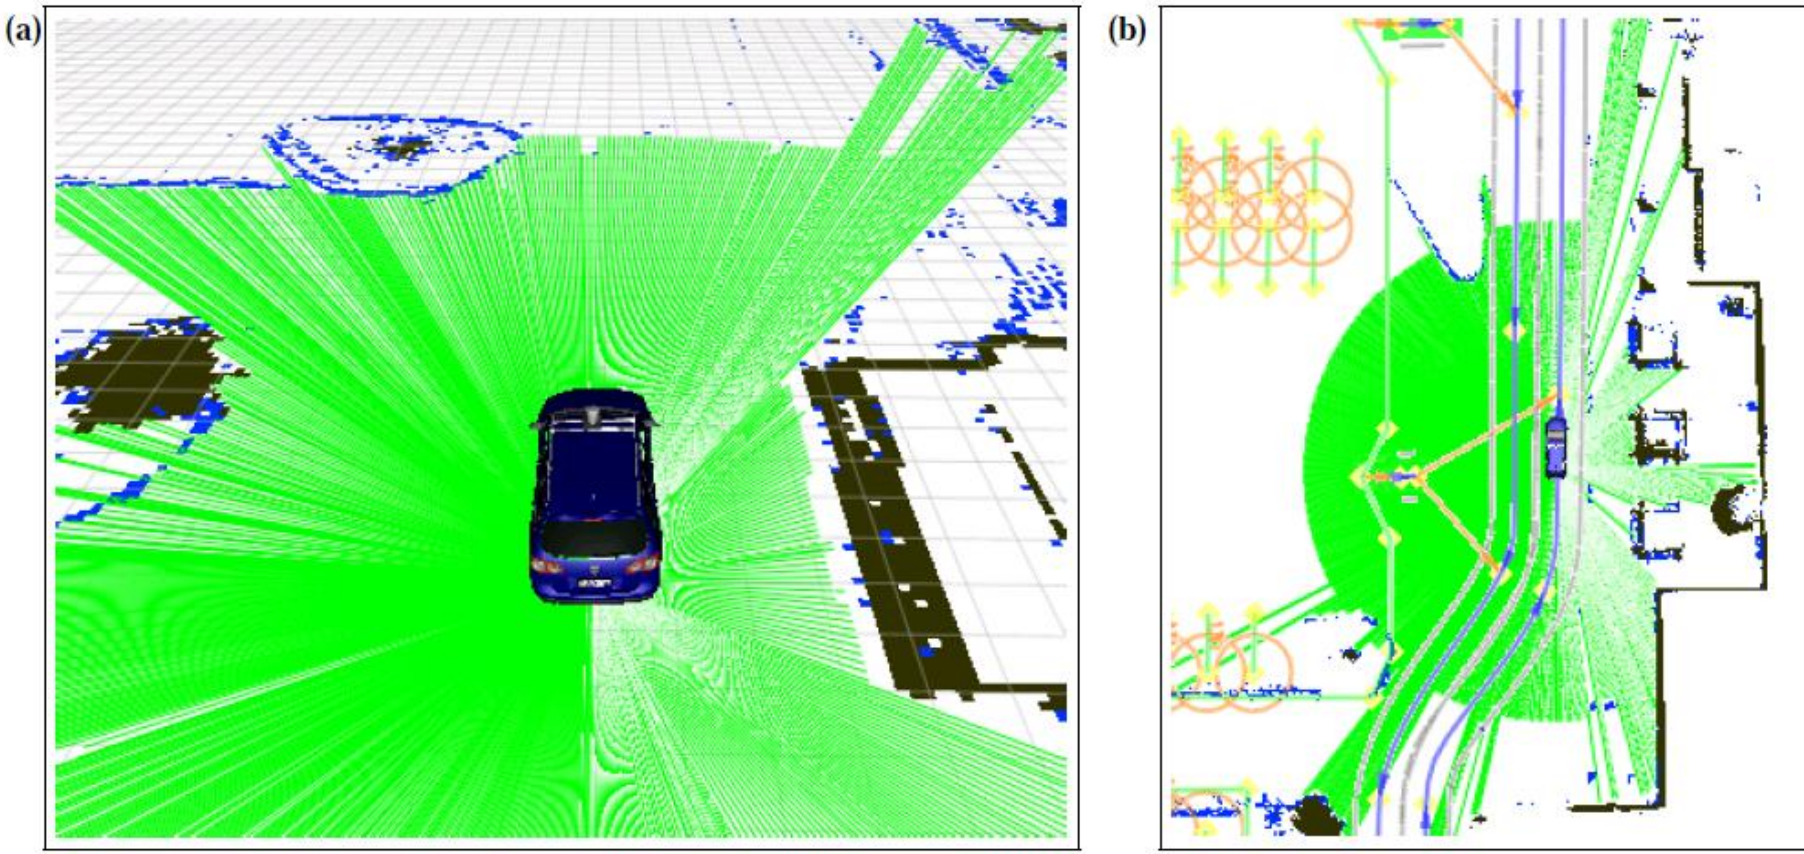

Figure 7: Examples of free space analysis for Velodyne scans. The green lines represent the area surrounding the robot that is observed to be empty. This evidence is incorporated into the static map, shown in black and blue.

Dynamic Object Detection and Tracking

- laser range scan data is mapped into a synthetic 2D scan of the environment

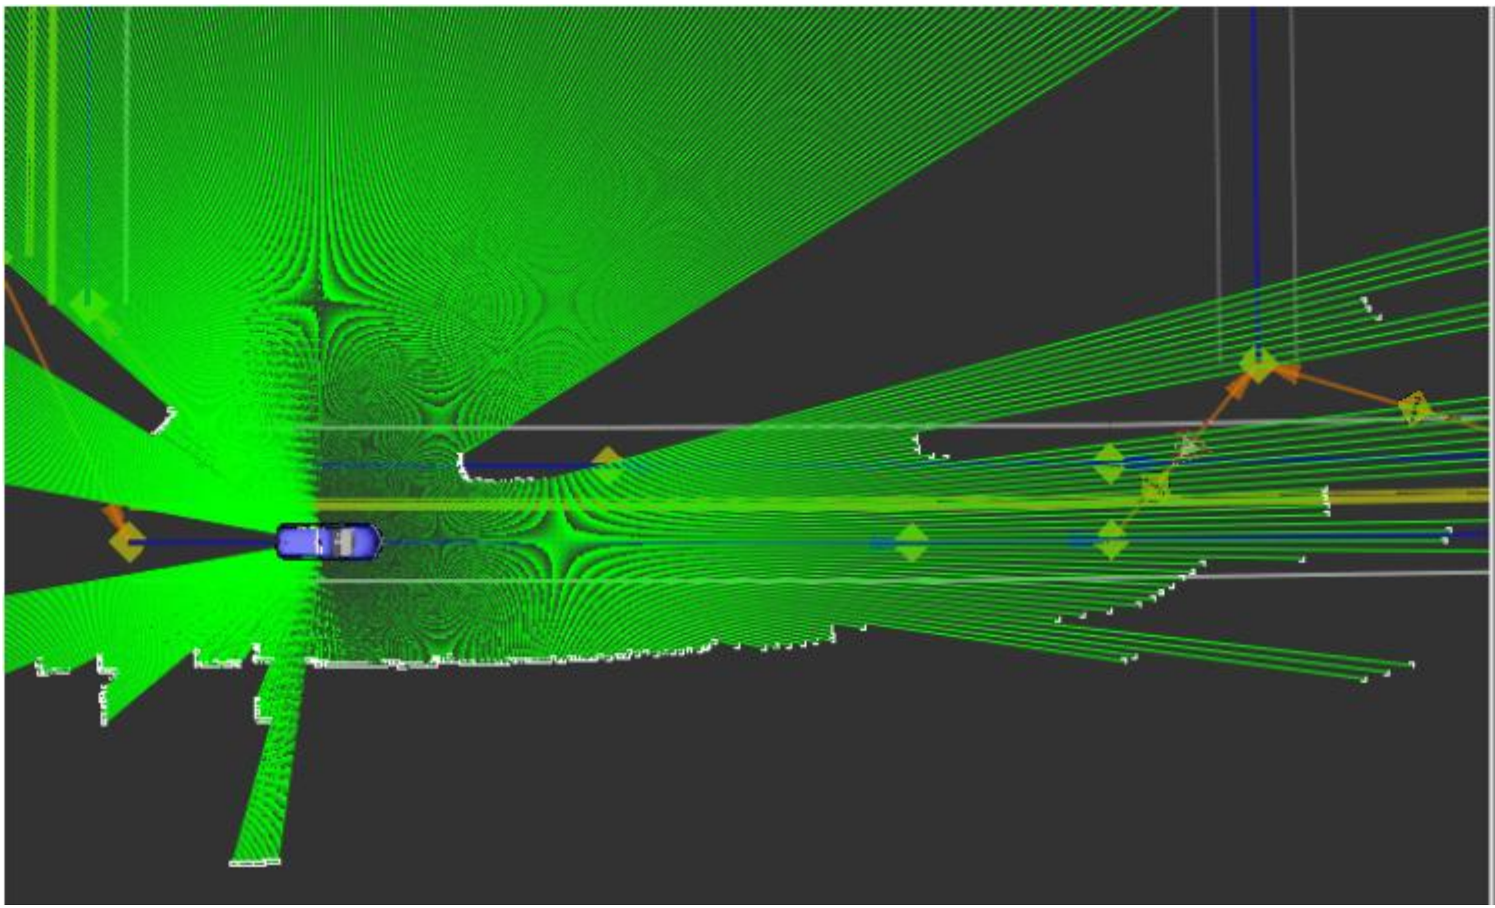

Dynamic Object Detection and Tracking

- areas of change are detected by comparing two synthetic scans taken over a short time interval
  - if an obstacle in one scan falls in the freespace of the second scan then this is evidence of motion

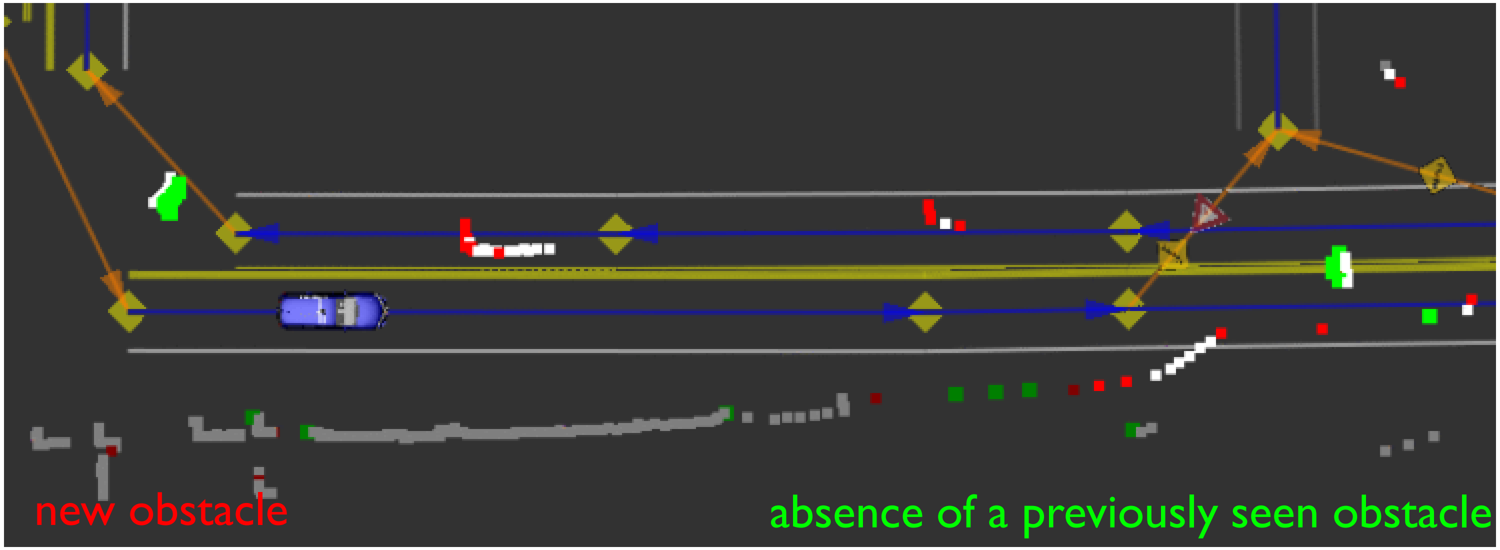

Dynamic Object Detection and Tracking

- particle filters are instantiated at each detected moving obstacle to track the object over time

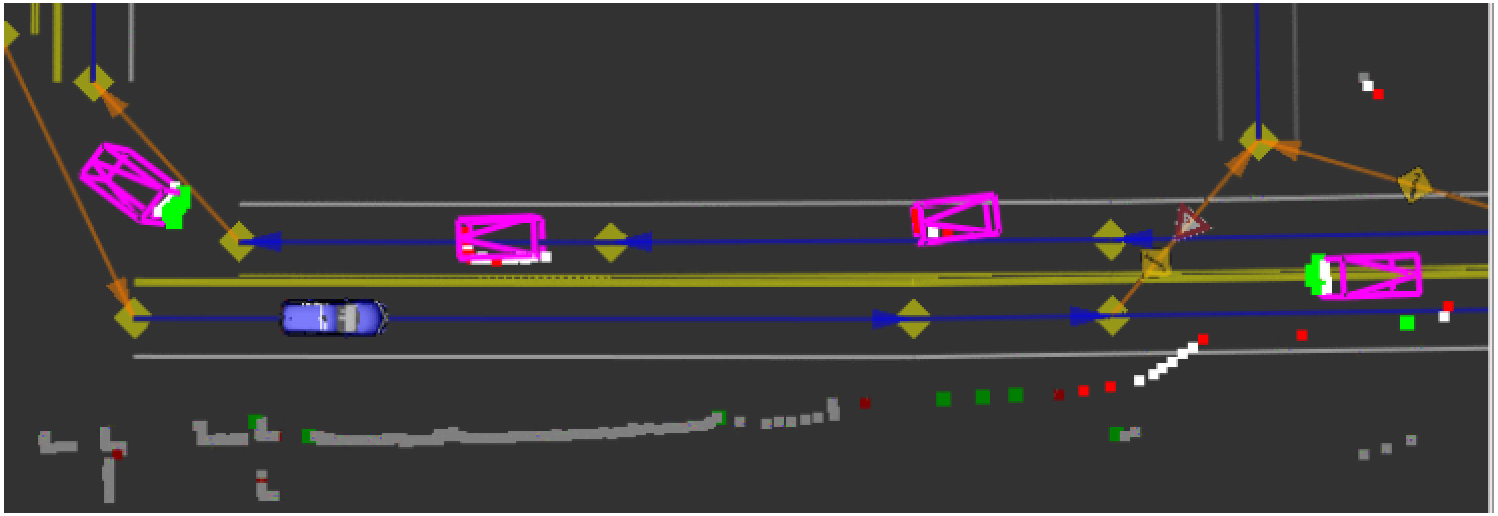

Dynamic Object Detection and Tracking

- camera view

Precision Localization

- high precision GPS and inertial measurements are not sufficiently precise to achieve single lane localization in the DARPA supplied map
- DARPA map is also inaccurate!
- high precision localization within the lane is achieved using curb measurements and road reflectivity
  - 2 side facing SICK LMS and the forward facing RIEGL LMS-Q120 laser range finders pointed downward to measure infrared reflectivity change between road surface and painted lane markers
- offsets greater than 1 m relative to GPS were common observed

Precision Localization

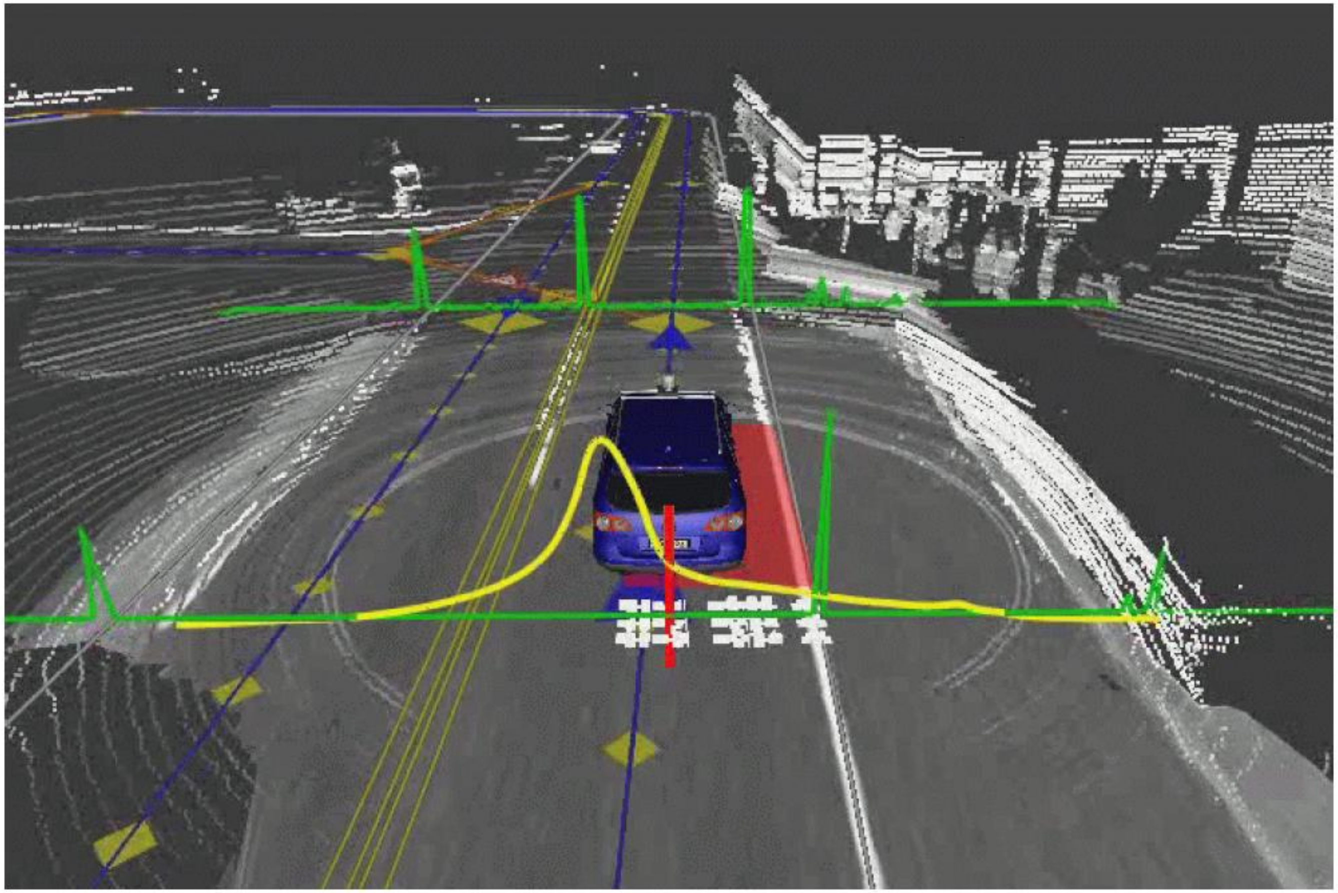

Figure 10: Typical localization result: The red bar illustrates the Applanix localization, whereas the yellow curve measures the posterior over the lateral position of the vehicle. The green line depicts the response from the lane line detector. In this case, the error is approximately 80 cm.

Global Path Planning

- performed at each checkpoint and when a permanent road blockage is detected

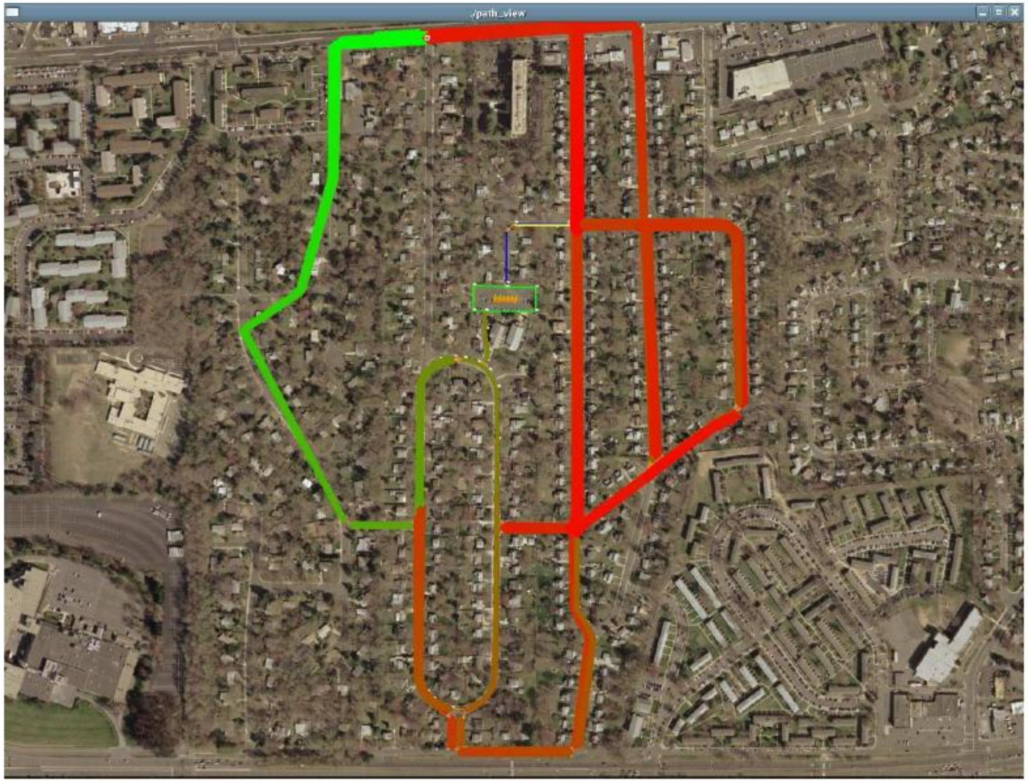

Figure 12: Global planning: Dynamic programming propagates values through a crude discrete version of the environment map. The color of the RNDF is representative of the cost to move to the goal from each position in the graph. Low costs are green and high costs are red.

Global Path Planning

- performed using dynamic programming on a discrete version of the map
- cost function weights choices (which way to turn at an intersection, lane changes, etc.) probabilistically
  - tends to perform lane changes earlier rather than later
  - will avoid passing a slow moving car in the right lane if a right hand turn is up coming
  - may decide against left hand turns at intersections if an alternate route is not much longer

Road Navigation

- global planner outputs a central path and also paths that have slight lateral shifts
  - allows the vehicle to pass slower moving vehicles

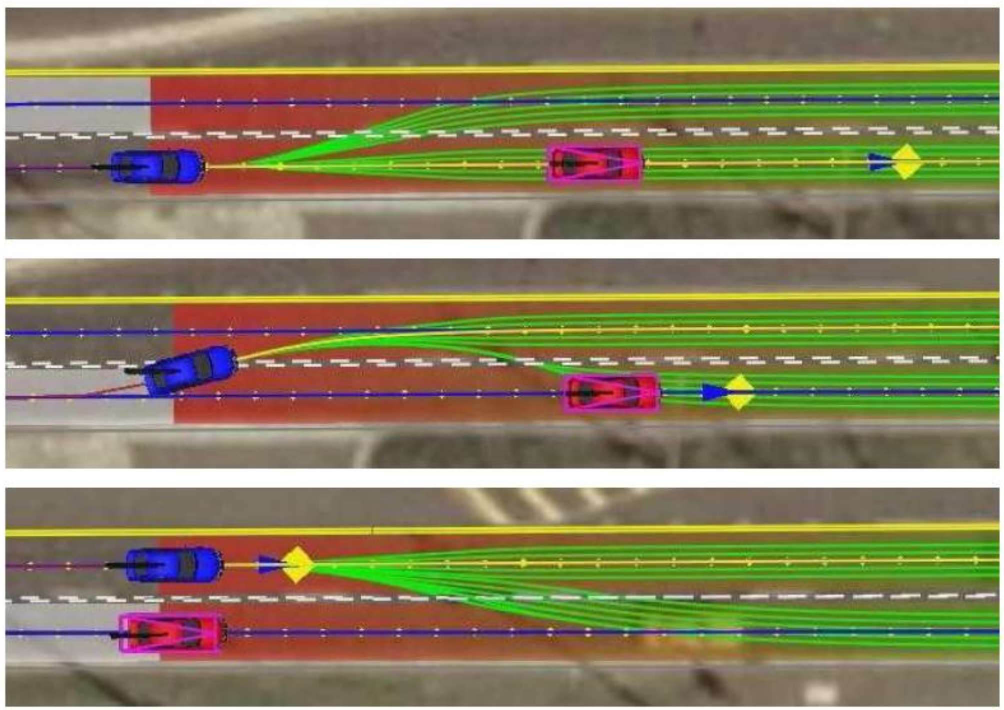

Figure 14: A passing maneuver. The additional cost of being in a slightly sub-optimal lane is overwhelmed by the cost of driving behind a slow driver, causing Junior to change lanes and pass.

Freeform Navigation

- used in parking lots and maneuvers such as U-turns
- planning performed using a variation of A* called hybrid-A*

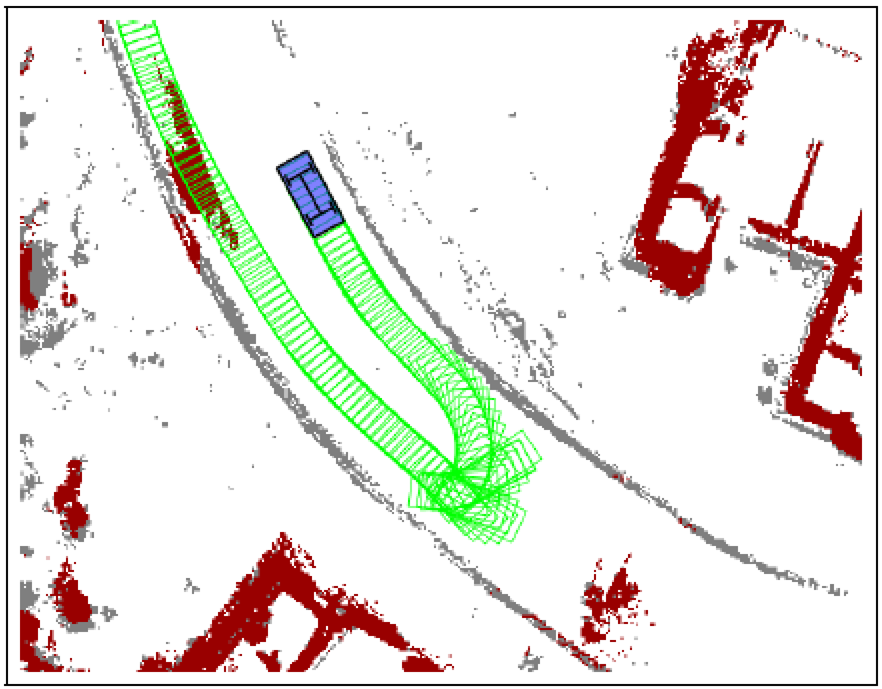

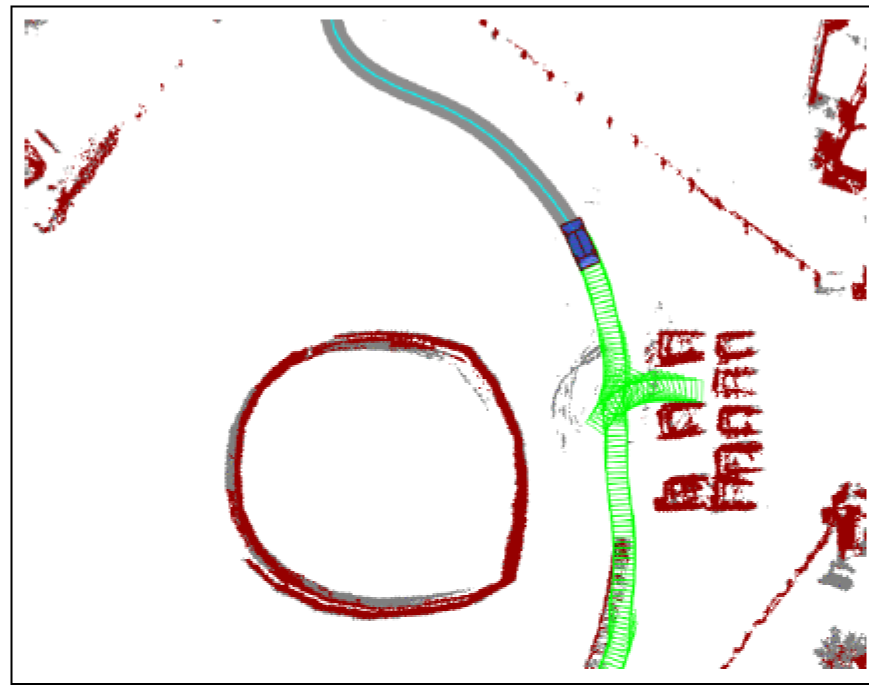

Behavior Hierarchy

- behavior governed by a finite state machine

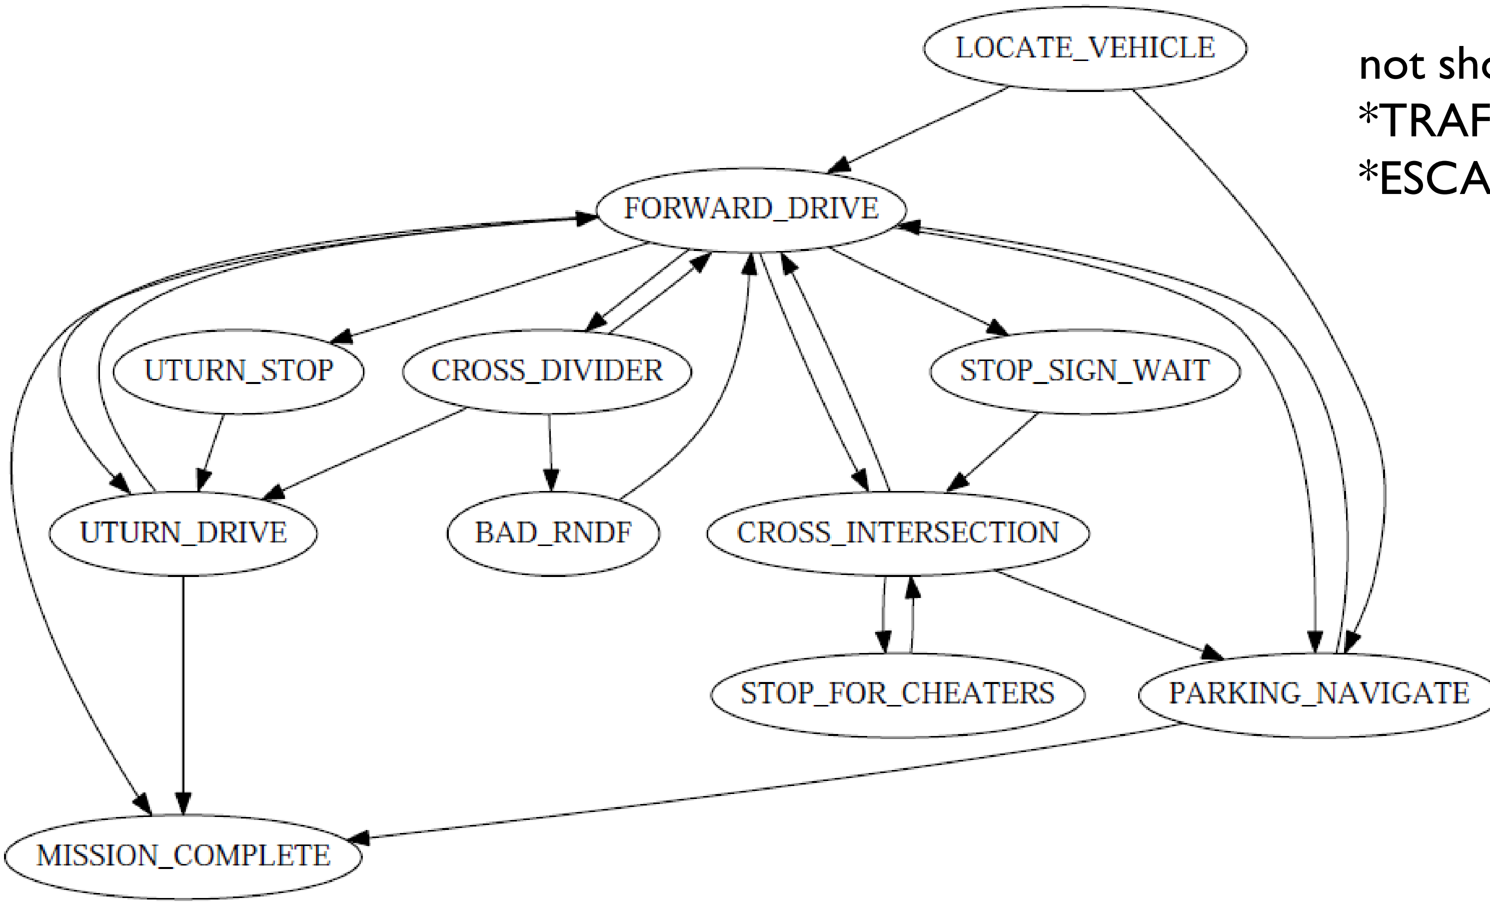

not shown
*TRAFFIC_JAM
*ESCAPE

Figure 21: Finite State Machine that governs the robot's behavior.

Escaping Blockages

- TRAFFIC_JAM and ESCAPE use hybrid-A* planner to find a path around road blockages

(a) Blocked intersection (b) Hybrid A* (c) Successful traversal

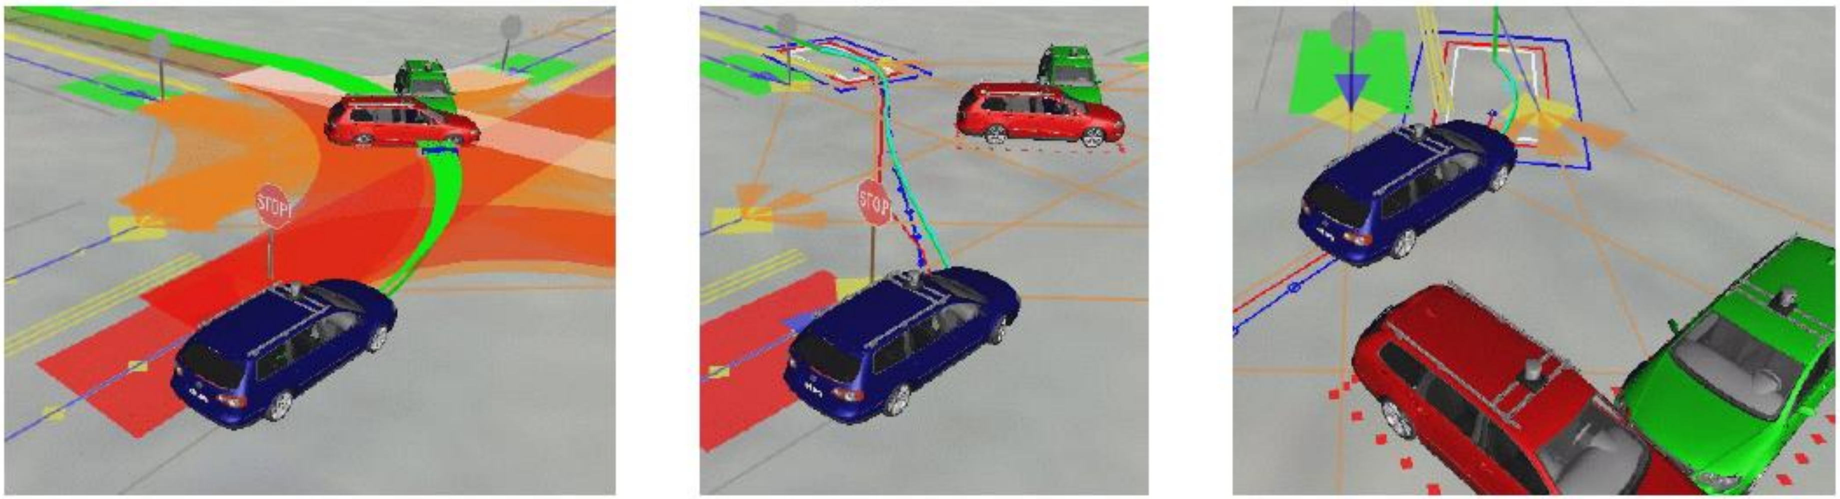

Figure 22: Navigating a simulated traffic jam: After a timeout period, the robot resorts to hybrid A* to find a feasible path across the intersection.

Results

- all 3 missions accomplished
- 55.96 miles in 4 hrs 5 mins 6 secs (22 km per hour)
  - good for 2^{nd} place
- notable race events
  - travel over a dirt road
  - passing a disabled competitor
  - avoiding a head on collision with a competitor driving partially in the wrong lane
  - braking in reaction to an aggressive merge
  - guilty of an aggressive merge
  - guilty of pulling alongside a competitor stopped at a stop sign (mistaken for a parked car)