Case Study (and revision): the DARPA Urban challenge

IAR Lecture 15

Barbara Webb

The DARPA challenges

- Grand Challenge: autonomous vehicles navigating desert trails and roads at ‘high’ speeds:
  - First event, 2004: all vehicles failed in first 10km of route
  - Second event, 2005: five vehicles completed 244km, 3 within 7hrs.
- Urban Challenge: autonomous vehicles driving through urban environment, obeying road laws and interacting safely with other vehicles:
  - Announced April 2006, 89 teams register, 53 first demos, 36 in qualification event, 11 in final event, 6 succeeded, 3 without human intervention in 3 missions over 97km in under 6 hours.
  - We will look at the first and second placing robots: ‘Boss’ & ‘Junior’

Boss: Team led by CMU

Boss (CMU)

Table I. Description of the sensors incorporated into Boss.

Sensor	Characteristics
Applanix POS-LV 220/420 GPS/IMU (APLX)	• Submeter accuracy with Omnistar VBS corrections • Tightly coupled inertial/GPS bridges GPS outages
SICK LMS 291-S05/S14 LIDAR (LMS)	• 180/90 deg × 0.9 deg FOV with 1/0.5-deg angular resolution • 80-m maximum range
Velodyne HDL-64 LIDAR (HDL)	• 360 × 26-deg FOV with 0.1-deg angular resolution • 70-m maximum range
Continental ISF 172 LIDAR (ISF)	• 12 × 3.2 deg FOV • 150-m maximum range
IBEO Alasca XT LIDAR (XT)	• 240 × 3.2 deg FOV • 300-m maximum range
Continental ARS 300 Radar (ARS)	• 60/17 deg × 3.2 deg FOV • 60-m/200-m maximum range
Point Grey Firefly (PGF)	• High-dynamic-range camera • 45-deg FOV

Junior: Team led by Stanford

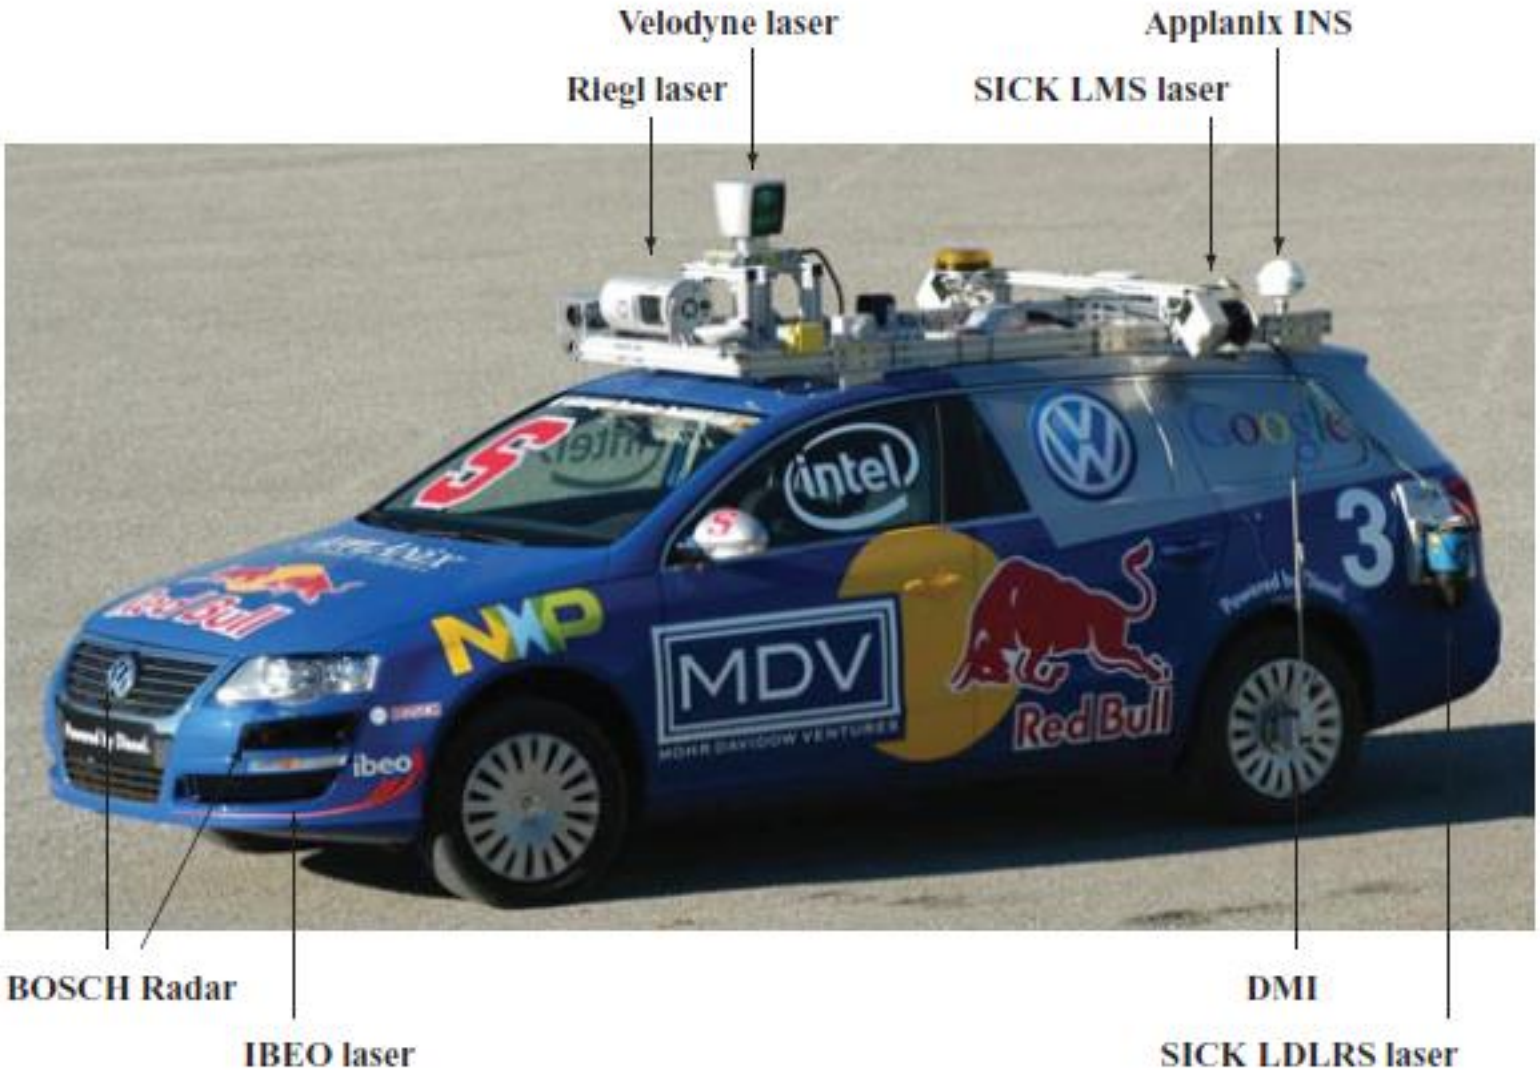

Exploiting physics?

- Using highly developed car technology as base system
- Some modifications, e.g.:
  - Boss: reduced compliance in steering, better brakes.
  - Junior: “limited-torque steering… electronic brake booster”
- Also note critical sensor technology:
  - Applanix fuses GPS and inertial and wheel encoder data for 100cm/0.1deg accuracy position estimate.
  - LIDAR uses reflected laser pulses to detect range information.
- Boss team mention criticality of endurance tests that picked up “intermittent and subtle software and mechanical defects” such as small gash in signal line causing a short circuit

Sensing for action?

- Both teams use multiple sensors, some fused, some redundant, some with specialised functions.
- Junior:
  - 2-D laser detects large, close obstacles
  - 3-D laser, use relative change in distance between rings to detect small obstacles such as kerbs

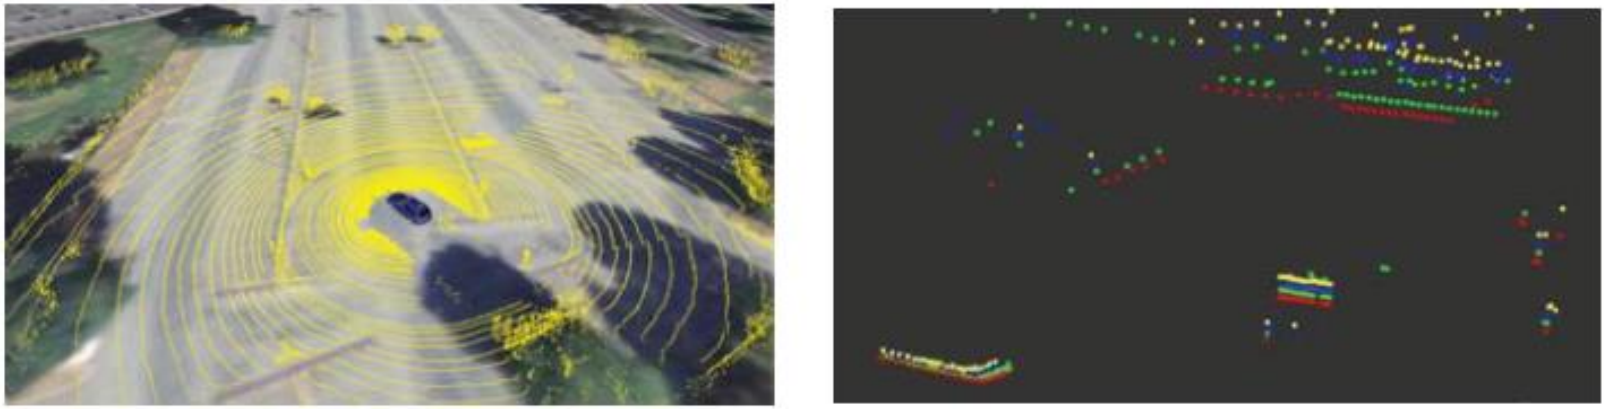

(a) (b)

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 m 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.

Exploiting dynamics?

- Boss motion controller: model-predictive control to generate dynamically feasible actions from start state to goal state.
- Control input from two parameterised functions: linear velocity and curvature
- Velocity function selected from four profiles:

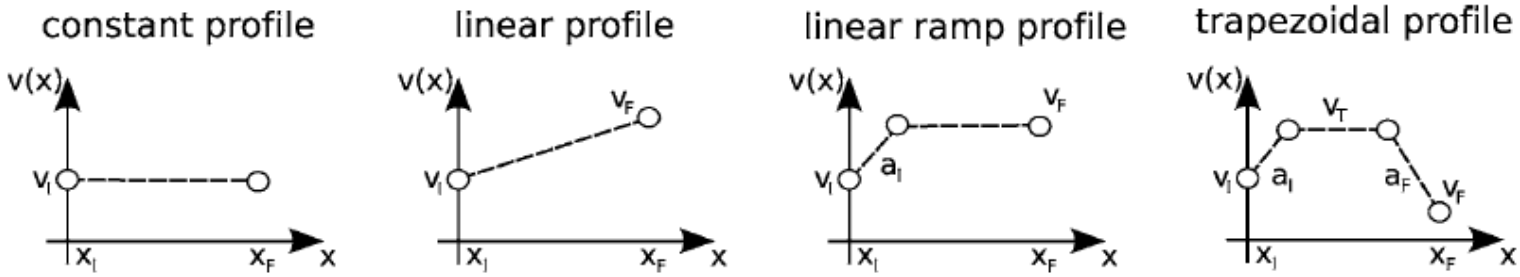

- Generate set of trajectories to goals at lateral offsets from centre-line, for each have sharp or smooth trajectory

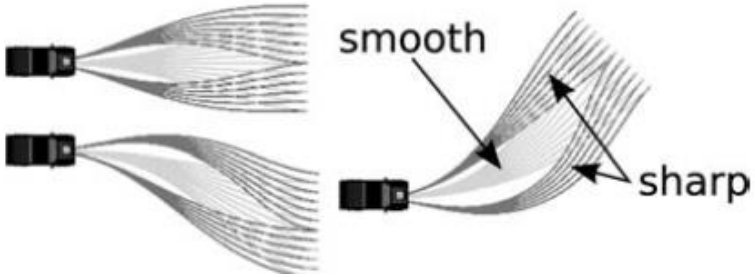

• Choose best trajectory dependent on obstacle sensing and other metrics

Combining behaviours?

Boss:

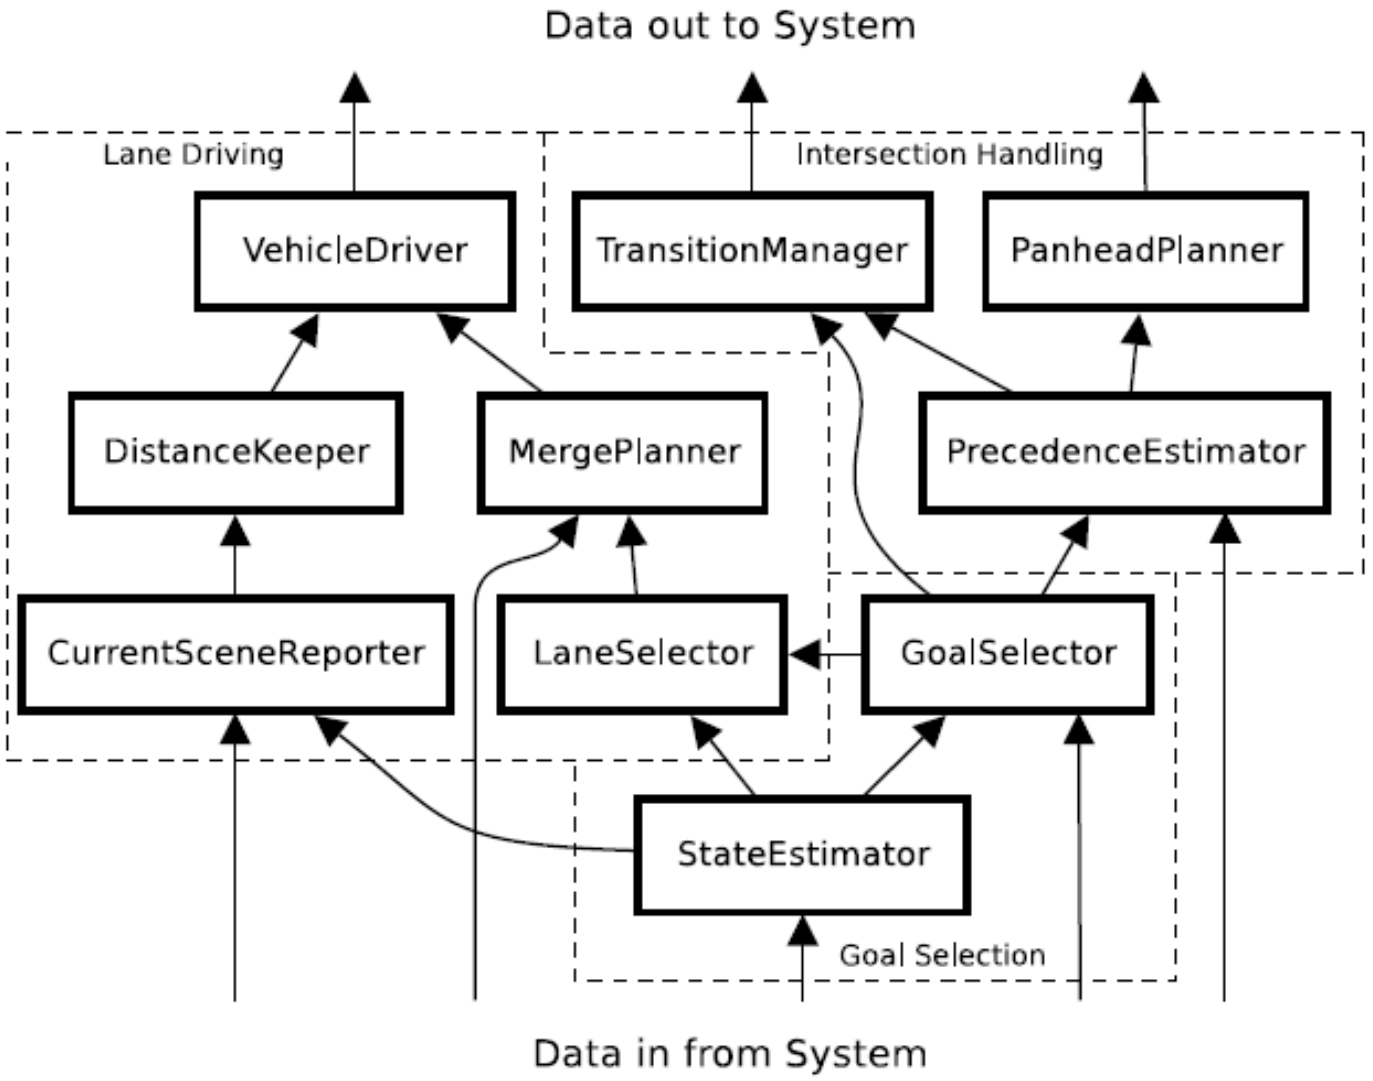

Combining behaviours?

"Junior's software architecture is designed as a data-driven pipeline in which individual modules process information asynchronously."

Finite state machine control of behaviour transitions

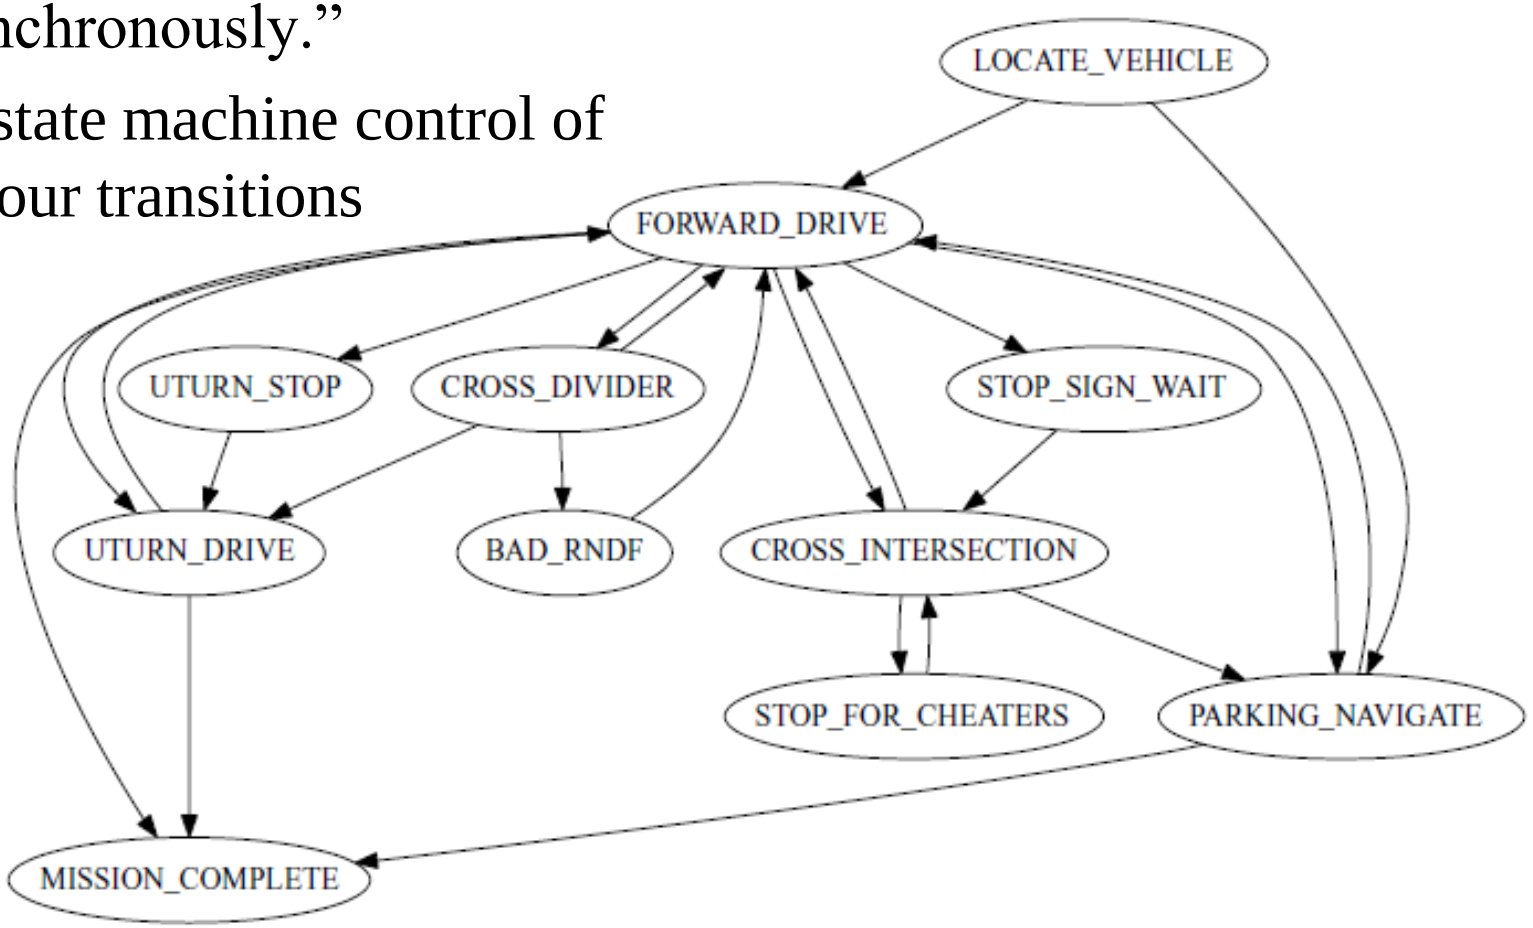

Maps?

- Competition entrants were given a detailed map in the form of a Road Network Definition File (RNDF) and high resolution aerial image
- Junior team used latter to refine former, e.g. adding way-points and smoothing trajectories

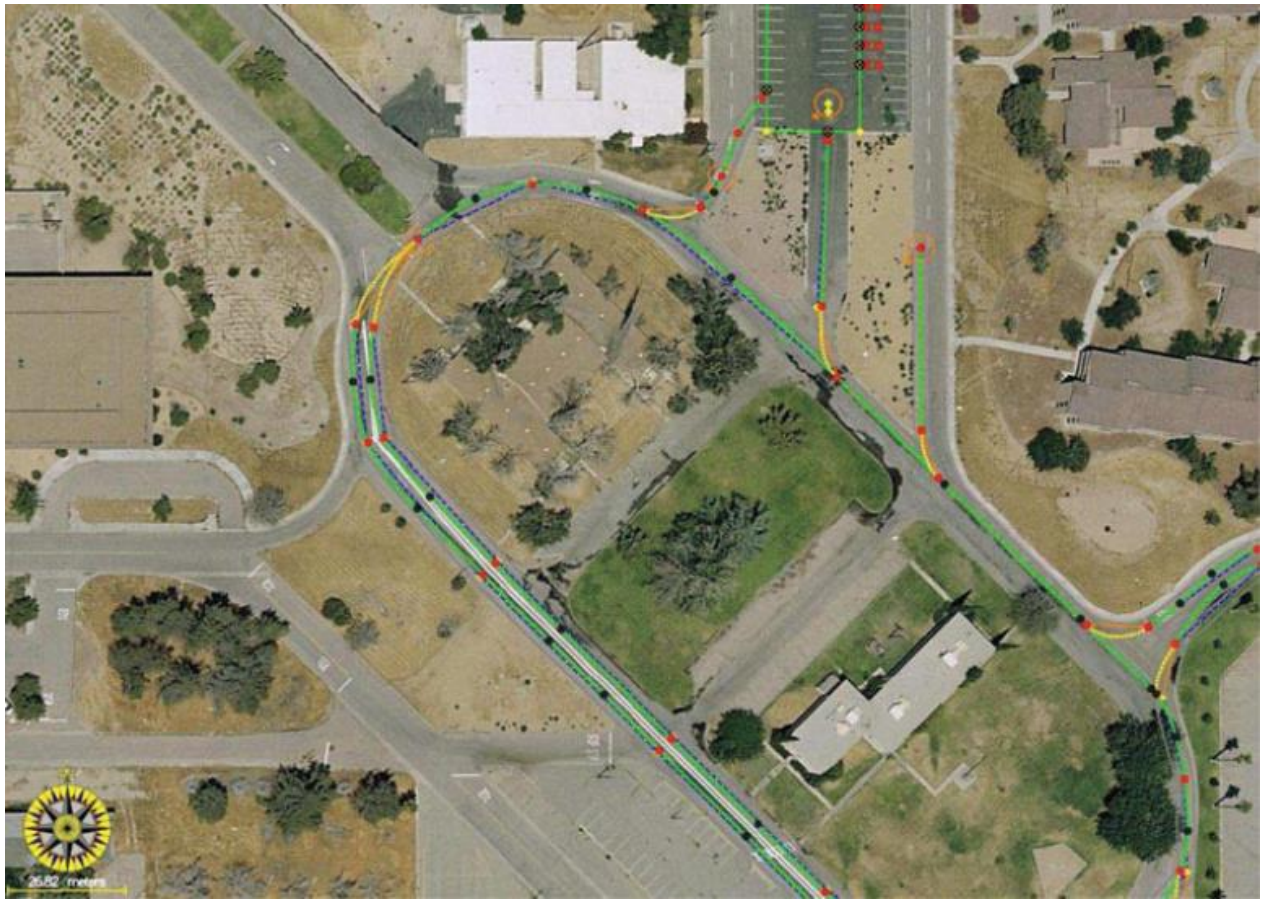

Localisation?

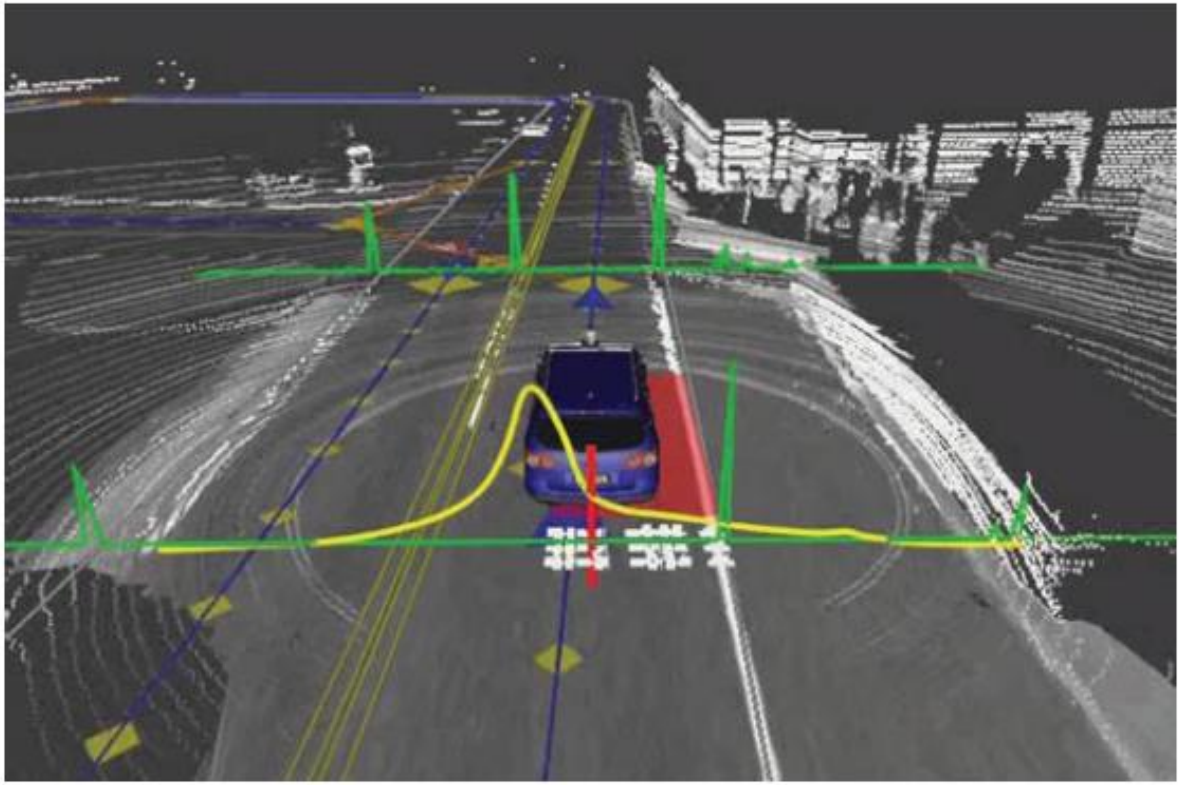

Figure 10. Typical localization result: The red bar illustrates the Applanix localization, whereas the yellow curve measures the posterior over the lateral position of the vehicle. The green line depicts the response from the lane line detector. In this case, the error is approximately 80 cm.

Filtering?

- Position filter for Boss: reject unreasonable position updates based on simple motion model for distance and heading:

$$\text{reject} = |\Delta\mathbf{x}| > v(1+\zeta)\Delta t + \varepsilon \;\vee\; \left\{ (|\Delta\mathbf{x}| > \varepsilon) \wedge \frac{\Delta\mathbf{x}}{|\Delta\mathbf{x}|} \cdot \begin{bmatrix} \cos(\theta) \\ \sin(\theta) \end{bmatrix} > \tau \right\}. \quad (12)$$

This checks that position change does not exceed that expected for velocity v with velocity error ζ and allowed jitter ε, and that any significant motion is in the expected travel direction θ with directional error τ

- Particle filter for Junior: dynamic object tracking
- Uses virtual sensor (a)– “synthetic 2-D scan” combines nearest objects from all laser data
- Any change (b) is a hypothesised to be a moving object; represent as set of particles with variable location, yaw, velocity and dimension
- Using prediction and update get particles locked onto real moving vehicles (c)

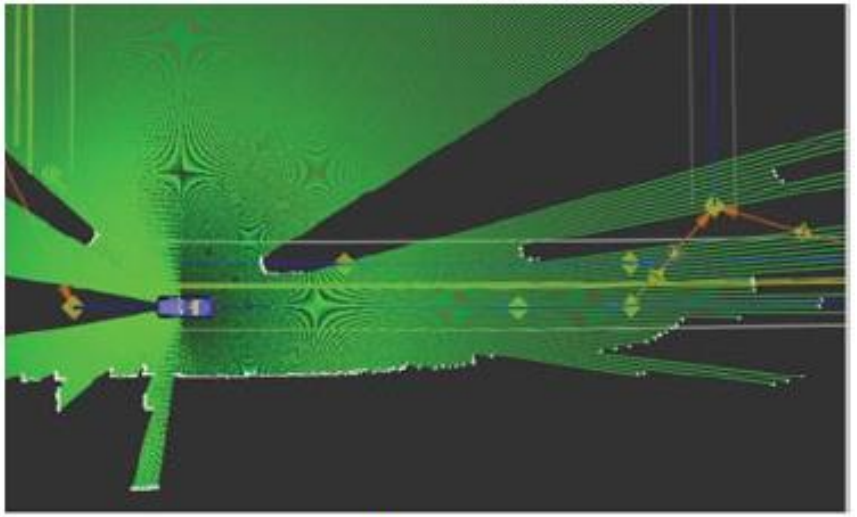

(a)

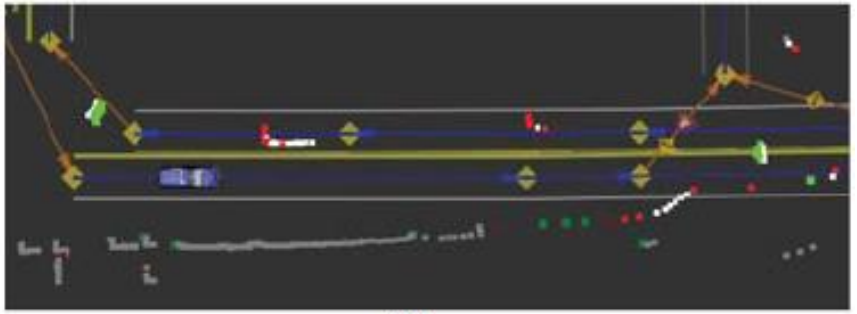

(b)

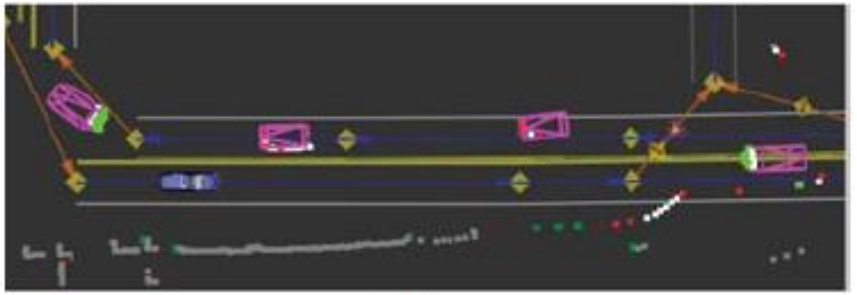

(c)

Planning?

- Junior:
  - Uses Hybrid A*
  - Standard grid-to-graph uses centre of cells as node locations, but vehicle cannot drive this path

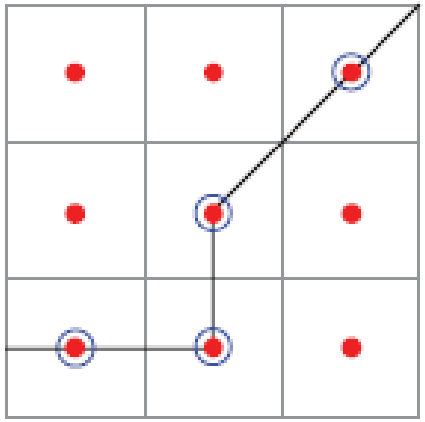

  - Instead ‘continuous’ cell co-ordinates calculated from predicted effect of control actions – trajectory that enters new cell determines associated node location

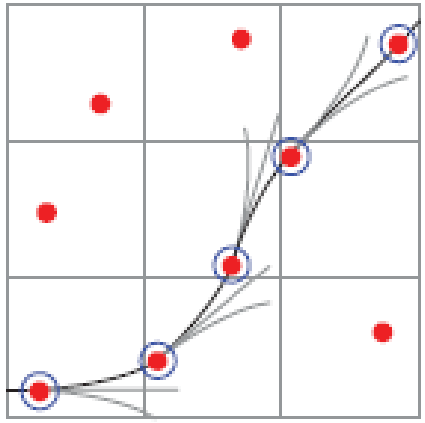

Planning?

- Boss:
  - Computes cost of all possible routes to next mission checkpoint based on connectivity graph: includes knowledge of road blockages, speed limits and time for different manouvres (e.g. left vs. right turn in traffic)
  - In unstructured environment (parking lot) uses anytime D* backward planning over state space of position, orientation and speed; variable resolution.
- "Anytime D* backward"?!?
  - Recall A* uses $f(n)=g(n)+\boldsymbol{\varepsilon}h(n)$, $\boldsymbol{\varepsilon}=1$
  - 'Anytime' uses $\boldsymbol{\varepsilon} >1$, which will run faster but give sub-optimal solution, reduce $\boldsymbol{\varepsilon}$ and replan if time allows
  - 'D' is dynamic, if map changes (e.g. detect new obstacle) recompute, but only for paths affected
  - 'backward' starts graph expansion from vehicle instead of goal, as observable changes are usually local

Hybrid architectures?

- Boss has three-layer architecture:
  - Mission planning: determines route to take to achieve high level goals
  - Behavioural: when to change lanes, give precedence at intersections, error recovery
  - Motion planning: determine trajectory that will avoid obstacles while progressing to local goals

Junior – asynchronous modular pipeline architecture

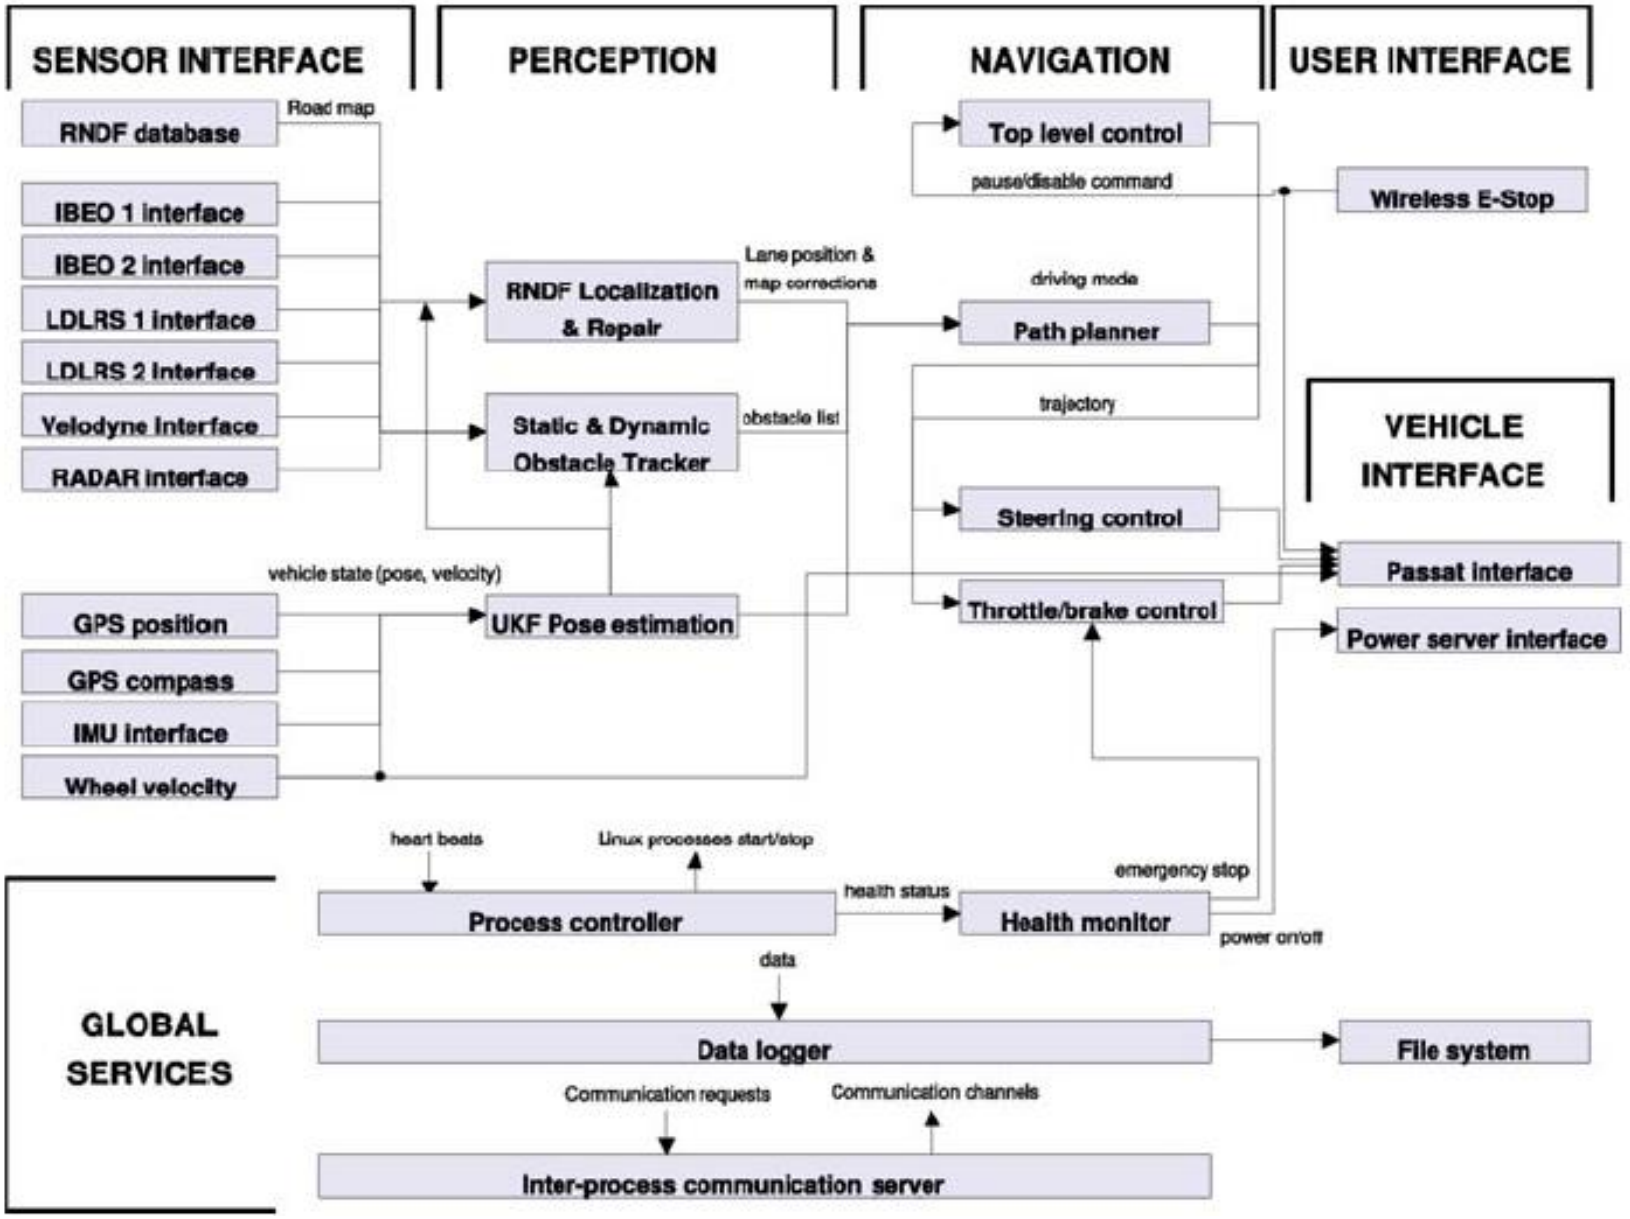

Remaining limitations?

- Sensor technology still not adequate for fully autonomous vehicles in real environments:
  - Dust raised by vehicle was then perceived as an obstacle
  - Media van jammed GPS signals
- Very limited representation of world, particularly other moving objects:
  - Boss: Mismatches between world model and reality led to assumed road blocks (another car in intersection) and long detours
  - Junior: treated car waiting at intersection as parked
- No suitable validation/verification for safety
- In real traffic, may need to be able to read social cues

The future?

- Autonomous cars coming close to commercial reality
- But most recent (2015) DARPA challenge, for humanoid robots in a rescue scenario, did not produce successful results:
  - Winner took 44 minutes to do 8 not very intelligent tasks
  - Most robots were far from autonomous

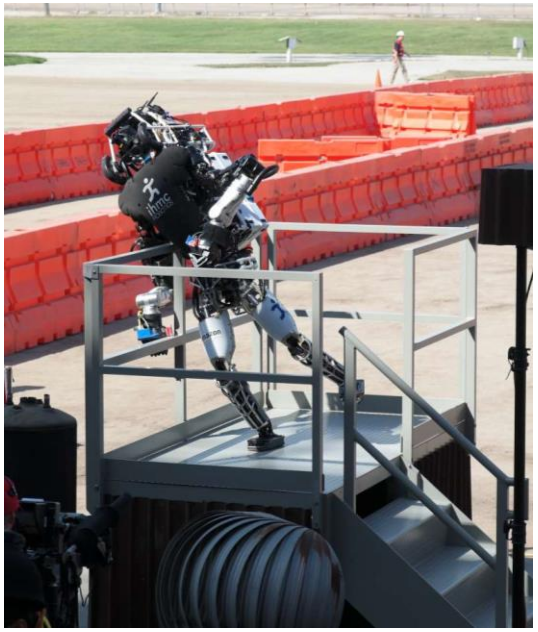

Thanks for listening and…

Please fill in an online course feedback form!

References:

http://www.darpa.mil/grandchallenge/index.asp

Urmson et al. (2008) “Autonomous Driving in Urban Environments: Boss and the Urban Challenge” Journal of Field Robotics 25(8): 425-466

Montemerlo et al. (2008) “Junior: The Stanford Entry in the Urban Challenge” Journal of Field Robotics 25(9): 569-597