interactIVe

Accident avoidance by active intervention for Intelligent Vehicles

Environment perception for automated vehicles

Dr. Angelos Amditis

interactIVe Summer School

4-6 July, 2012

Contents

- Introduction
- History of automated driving
- Data fusion and challenges
- Perception in HAVEit
- Talos vehicle – MIT
- Boss vehicle – Carnegie Mellon University
- Conclusions & Outlook
- References

Automation in vehicles

Problems

- Over 90% of accidents are driver related (burdensome driving, fatigue driving, non-professional drivers etc.)
- Traffic congestion
- Pollution of the environment (CO_2 emissions)

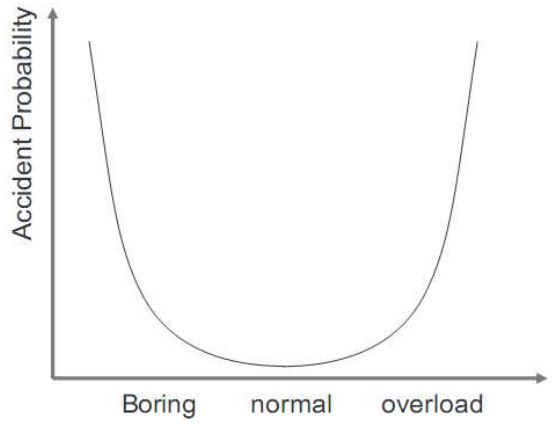

Solution-automation

- Assist the driver in monotonous and demanding driving tasks
- Benefits:
    - Safety
    - Comfort
    - Efficiency
    - Sustainability

Role of perception in automated driving

- Accurate & real-time representation of the surrounding environment (world model)
  - Static/moving vehicles
  - Road modeling
  - Other obstacles

- Use of multiple sensors
  - **Active sensors:** radars, lidars, ultrasonic etc.
  - **Passive sensors:** infrared and visual cameras
  - **Virtual sensors:** V2X communication, digital maps

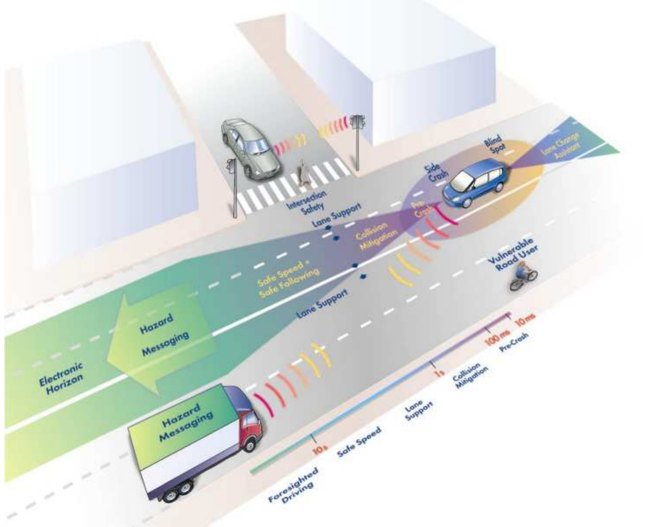

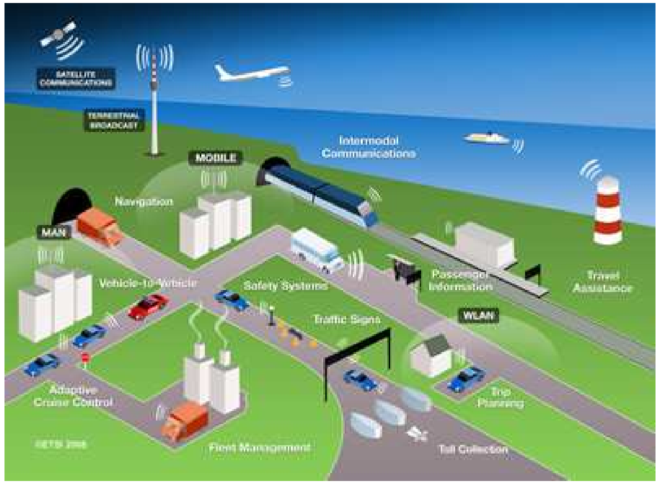

History – Early attempts

- **1939:** The idea of autonomous vehicles gained widespread public exposure at GM's Futurama exhibit at the 1939 World’s Fair, where the automaker envisioned "abundant sunshine, fresh air and fine green parkways" upon which cars would drive themselves.
- **1953:** GM and Radio Corporation of America (RCA) had developed a scale model automated highway system, which allowed them to begin experimenting with how electronics could be used to steer and maintain proper following distance.
- **1958:** GM tested a Chevrolet with a front-end featuring "pick-up coils" that could "sense the alternating current of a wire embedded in the road and would adjust the steering wheel accordingly.

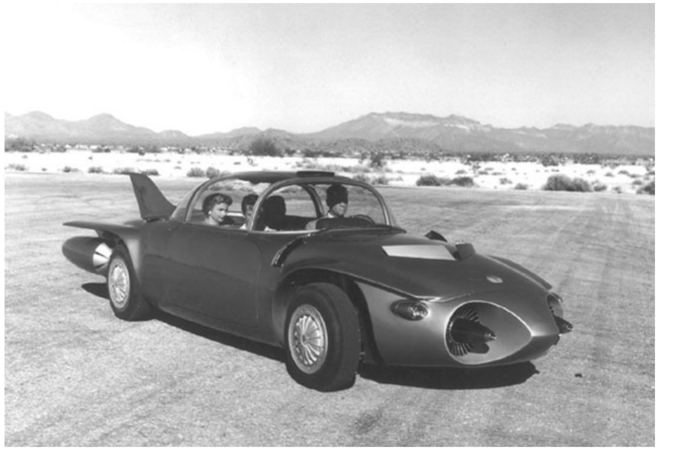

History – First worthy attempts

- **1977:** Tsukuba Mechanical Engineering Lab in Japan creates the first autonomous, intelligent, vehicle. It tracked white street markers and achieved speeds up to 30 kilometers per hour.
- **<u>1980 (breakthrough)</u>:** Ernst Dickmanns and his group at Bundeswehr University Munich (UniBW) build robot cars using saccadic vision, estimated approaches like Kalman filters, and parallel computers. They went up to 96 km/h on an empty street.

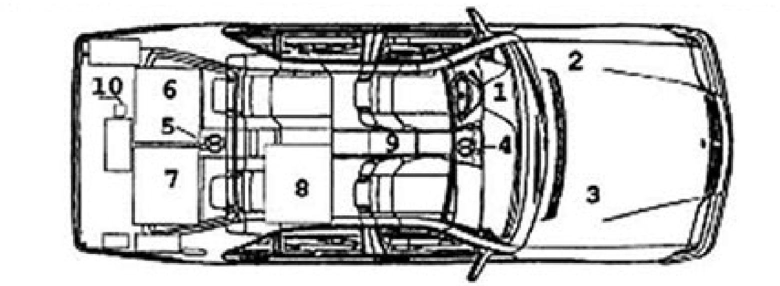

1 Torque motor for steering
2 brake system
3 electric throttle control
4 front platform } for 2 CCD
5 rear platform } cameras
6, 8 Transputer system, image processing
7 processors for gaze & locomotion control
8 user interface
9 linear accelerometers
10 angular rate sensors

History – 1990's

- **1987-1995:** The pan-European Prometheus project, also known as the EUREKA Prometheus Project, the largest autonomous vehicle project so far, is funded by the EC.
- **1995:** Throttle and brakes needs human intervention, but a Mercedes-Benz model created by UniBW drives from Munich to Copenhagen and back, more than 1000 autonomous miles on a highway in traffic, and exceed speeds of 177 km/h. It completes the journey with 95% autonomous driving.
- **1995:** CMU Navlab "*No Hands Across America Project*". The car made almost 3000 miles 98.2% autonomously needing a bit of help with obstacle avoidance. Throttle and brakes needed human control.
- **1997:** AHS (Automated Highway System) revolutionary demonstration made in 1997 included more than 20 fully automated cars (US activity).

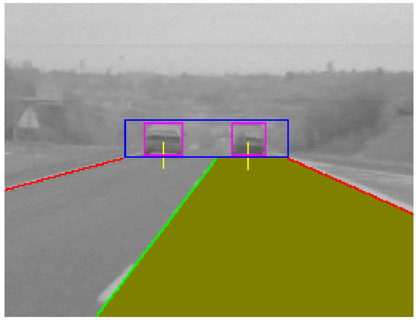

History – DARPA’s Challenges

- **2005:** DARPA’s American “grand challenge” begins with no traffic and a few road markers, if necessary, in the desert. The course has 2935 GPS points and is revealed in advance. The top car, with a max speed of 40 km/h, to complete the 211 kilometer desert course is the VW of Stanford, which finished the course in 6 hours and 54 minutes.
- **2007:** DARPA's "Urban Challenge" won by Carnegie Mellon University. Sensor systems become more elegant and semi-autonomous features begin to hit the mainstream with manufacturers from Audi and Volvo, to GM and Mercedes incorporating features like collision avoidance, lane recognition, and driver attention assist into their new vehicle lines.

History – Latest activities

- **University of Parma's Road Trip:** Parma’s VISLAB undertook the most geographically daunting autonomous car journey in 2010, driving from Parma to Shanghai. The trip took them 16,000 km through 9 countries in 100 days. *The first autonomous vehicle to be ticketed by a traffic cop* (in Russia).
- **Shelley Climbs the Mountain (Audi TTS):** The car conquered the 12.42-mile sprint to the summit of Pike’s Peak in 27 minutes. Sure, that’s 17 minutes off the human record, but considering the first human-guided, steam-powered car took more than nine hours to make the ascent in 1901, it is an auspicious debut for a computer-controlled car.
- **Google Driverless Car:** Google's fleet of seven autonomous Toyota Prius hybrids has racked up more than 140,000 miles with only occasional human intervention since hitting the road in 2010.

History – Latest European projects

- **CityMobil** (2006-2011): Major research, development & demonstration project. It addressed the integration of automated transport systems in the urban environment. Demos realized at Heathrow airport, Rome, La Rochelle and Valencia (www.citymobil-project.eu).
- **HAVEit** (2008-2011): At the HAVEit Final Event, 17 partners from the European automotive industry and scientific community demonstrated the highly automated future of driving (http://www.haveit-eu.org).
- **SARTRE** (2012): Stepping into the domain of platoons where the lead vehicle is driven by a professional driver while following vehicles will enter a semi-autonomous control mode (www.sartre-project.eu).

- **interactIVe** (2010-2013): Next generation ADAS for safer & more efficient driving based on active intervention. Safety systems that brake and steer autonomously (http://interactive-ip.eu/).

Role of data fusion

- Combination of *multiple heterogeneous sensors* and *information sources* to create an **accurate environment representation** and identify **relations between road entities**
- Multi-sensor data fusion
  - Make the best out of each sensor (minimize uncertainty, combine partial/complementary information)
  - Maximize data quality and availability
  - Increase reliability and robustness
  - Reasoning of present situation and prediction of future risks
- **Joint Directors of Laboratories – JDL model**
  - *Object refinement*: Object recognition and tracking
  - *Situation refinement*: Maneuver identification, objects' relations etc.

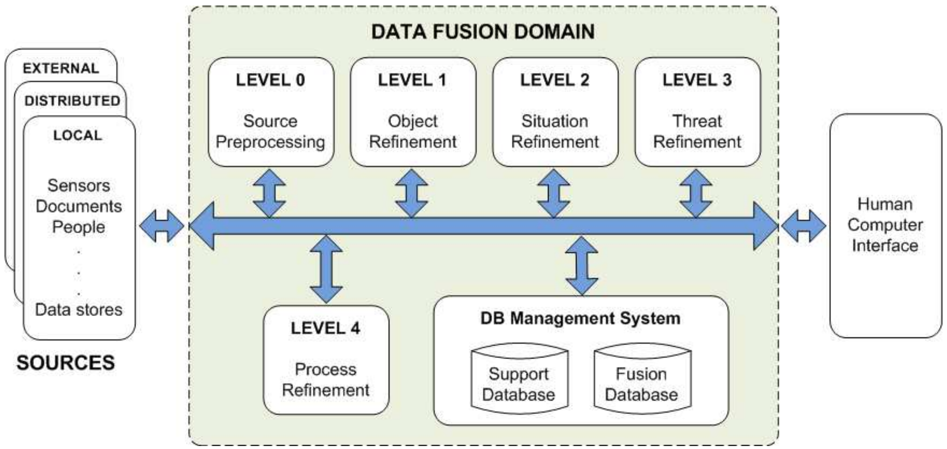

Challenges for data fusion in automated driving

- Integrated perception approach
- High performance requirements (compared to other ADAS)
  - Nearly 0% false alarms
  - Improved accuracy
  - Real time operation
  - Longer range (driver is not directly “into the loop”)
- Complexity
  - Large sensor network for 360º coverage
  - Many sensor fusion modules
- Integration of wireless messages (handle delays, latencies etc.)
- Lower the cost in the vehicle’s side
  - Need for low cost sensors (performance?)
  - Infrastructure investments

Sensor data fusion in the HAVEit project

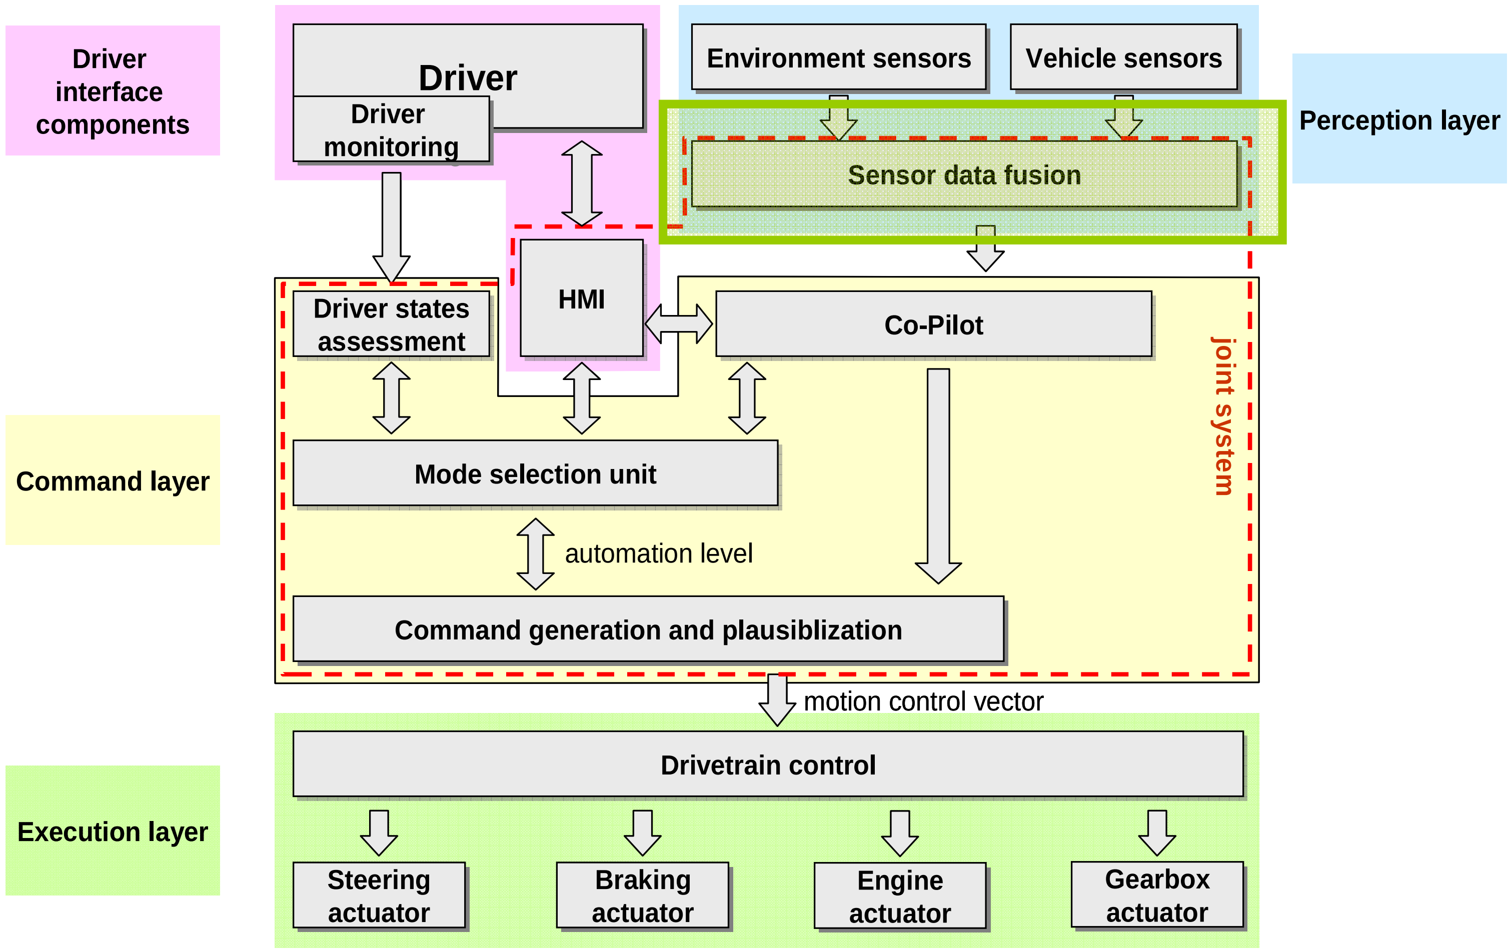

Sensor data fusion overview

- **Perception layer**
  - Ego vehicle state
    - Kinematic
    - Relative to the road
  - Road Environment
    - Lanes
    - Objects
  - Additional information

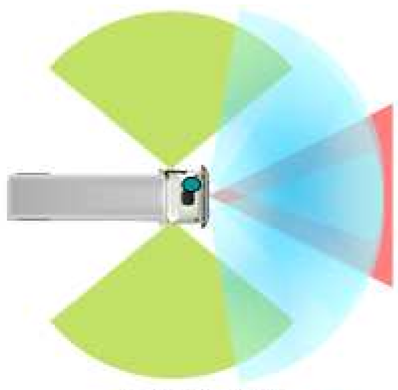

- **The Generic data fusion concept**
  - 2 levels of processing hierarchy
  - Implementation of the same algorithms for different demos
  - Implementation of software modules applicable to many H/W platforms

Track level fusion architecture

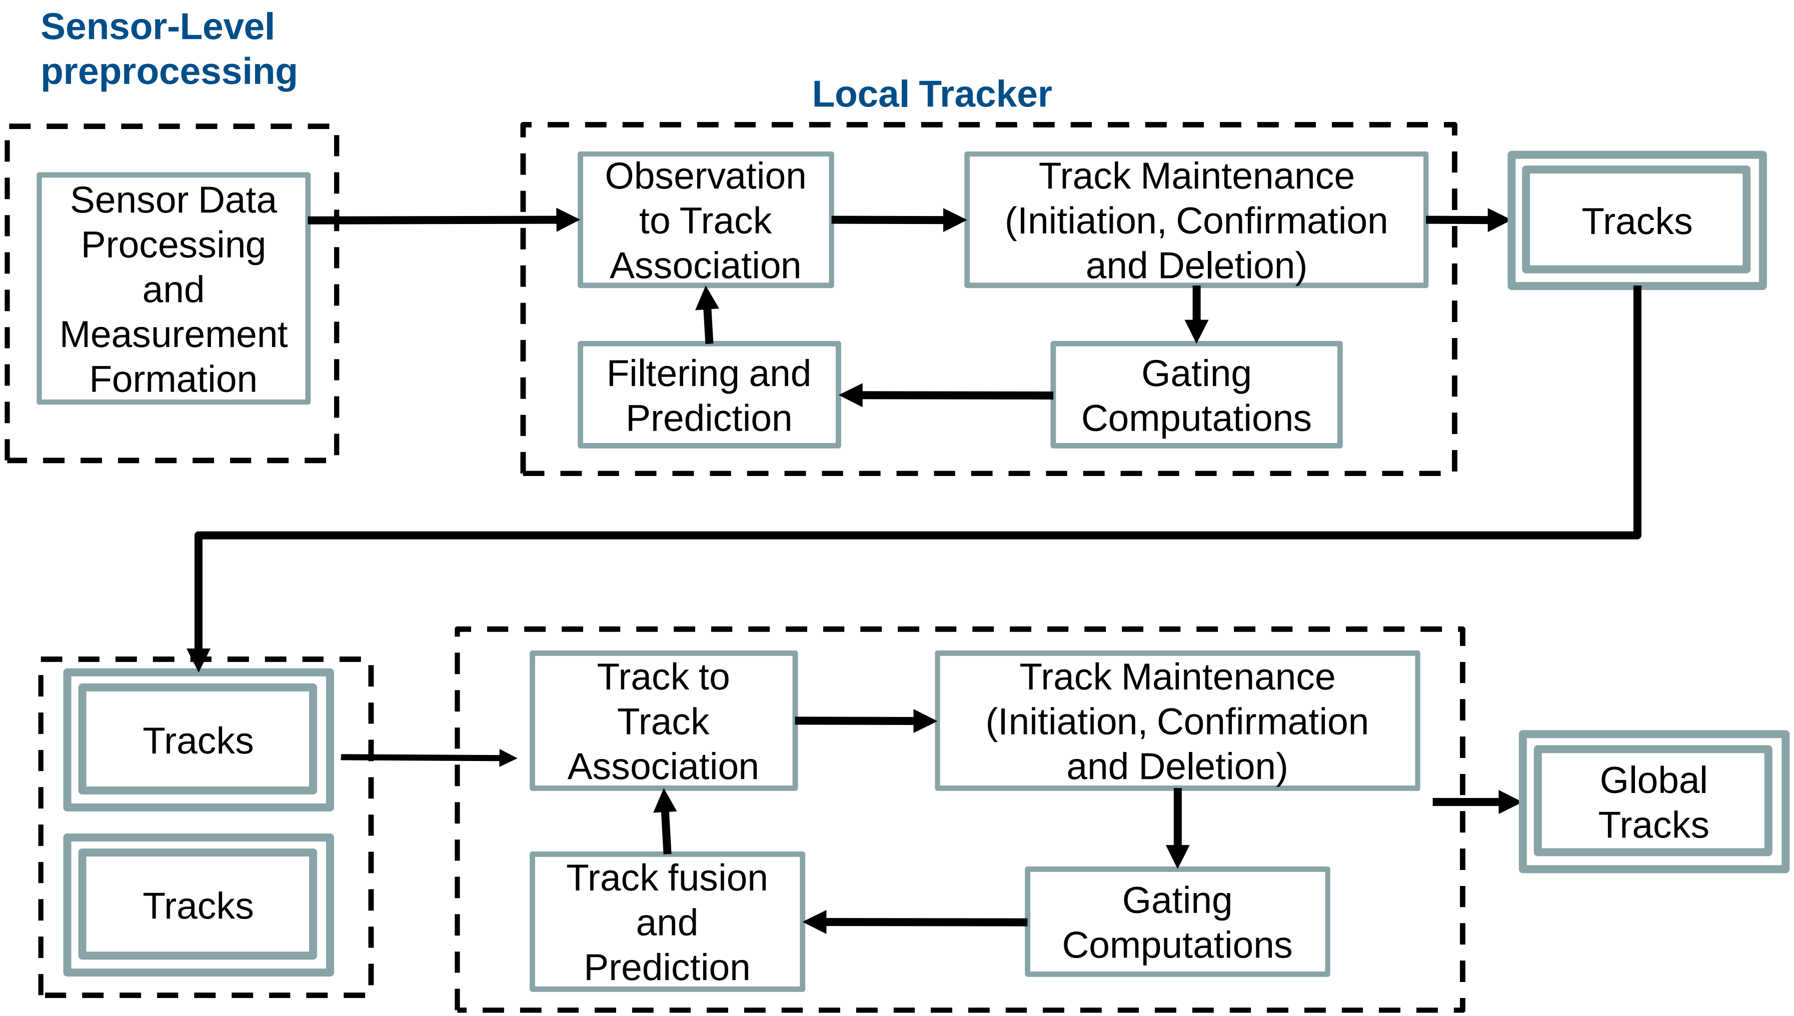

Target tracking

- **Sensor level tracking**
  - Tracking is carried inside each sensor
  - Measurement to track assignment using *auction algorithm*
  - *Track management* (confirmation & deletion) is done using “hit” and “miss” based rules
  - Track state update is done using the standard *Kalman filter*

- **Central level tracking**
  - Identify local tracks that represent the same object
  - Fuse local track estimates
  - Track ID maintenance in track transitions between sensor FOVs
  - Object management using probabilistic or rule based methods

Track fusion

- Takes as input the track lists of the local trackers and gives a single track list in the output

- The track-to-track association module identifies which tracks from different tracks list represent the same object

- The Mahalanobis distance of the two tracks is calculated as follows: $$d_{ij}^2 = \tilde{x}_{ij}'\left(P_i + P_j - P_{ij} - P_{ij}'\right)^{-1}\tilde{x}_{ij} = \tilde{x}_{ij}' S_{ij}^{-1}\tilde{x}_{ij}$$

- The fused estimate of the two independent estimates is

$$\tilde{x} = \tilde{x}_i + \left(P_i - P_{ij}\right)\left(P_i + P_j - P_{ij} - P_{ji}\right)^{-1}\left(\tilde{x}_j - \tilde{x}_i\right) = \tilde{x}_{ij}' S_{ij}^{-1}\tilde{x}_{ij}$$

Lane estimation

- Lane geometry
  - Kalman filtering
  - Clothoid model

$$y(l) = y_0 + \int_0^l \sin(c_o\tau + \frac{c_1\tau^2}{2})d\tau \qquad x(l) = \int_0^l \cos(c_o\tau + \frac{c_1\tau^2}{2})d\tau$$

$$y(x) = y_0 + \tan(h)\cdot x + c_0 \cdot \frac{x^2}{2} + c_1 \cdot \frac{x^3}{6}$$

- Lane description

$$\begin{bmatrix} Curv \\ Curv_{rate} \\ y_{offs} \\ head \\ width \end{bmatrix}$$

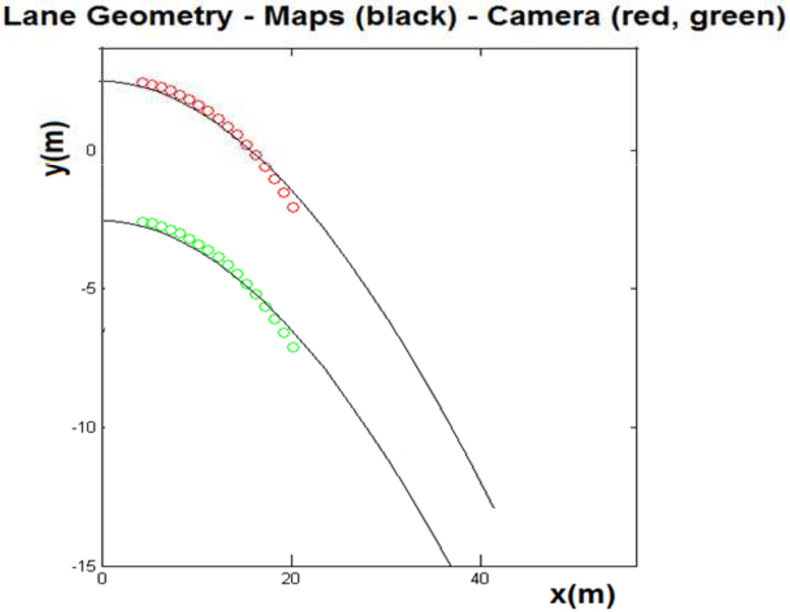

- Lane estimation is based on the camera sensor (proved to be more reliable)
- Lane estimation based on laserscanner measurements was used as a back-up solution

Object Perception – Occupancy Grid Fusion

- A stochastic tessellated representation of spatial information
- Probabilistic estimates of the occupancy state of each cell
- Estimation of the free space around the vehicle
- Objects are extracted from the occupancy grid

$$P(Z_i \mid occ) \qquad P(Z_i \mid empty)$$

$$P(O_{x,y} \mid Z_1 \ldots Z_N) = \frac{P(O_{x,y}) \times \prod_{i=1}^{N} P(Z_i \mid O_{x,y})}{P(O_{x,y} = occ) \times \prod_{i=1}^{N} P(Z_i \mid O_{x,y} = occ) + P(O_{x,y} = empty) \times \prod_{i=1}^{N} P(Z_i \mid O_{x,y} = empty)}$$

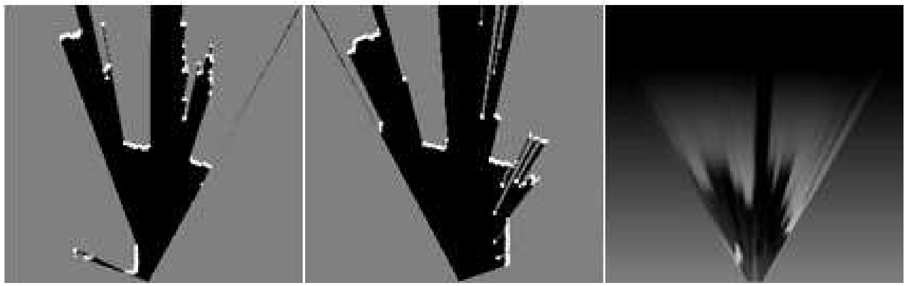

Tracking Estimation Techniques – Multiple Hypothesis Tracking

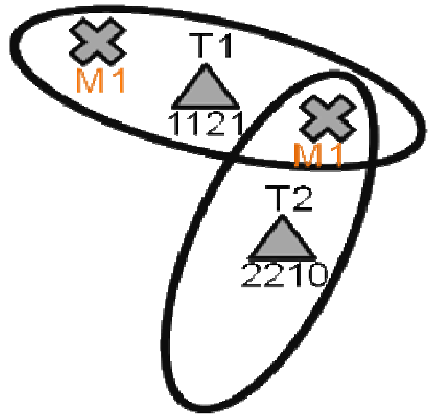

- Handling conflicting assignment Situation

- Propagation of all possible track measurement pairs

- Future observations resolve ambiguities in the past

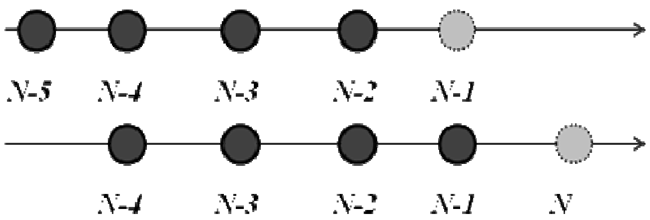

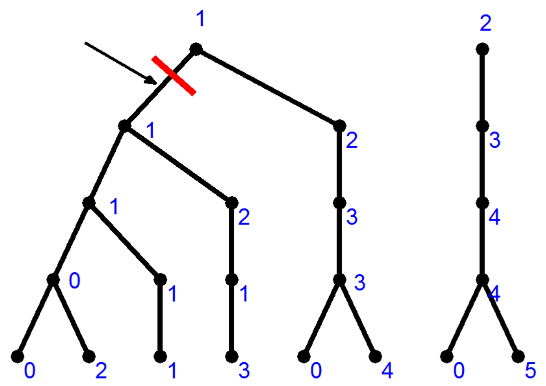

Tracking Estimation Techniques – Filtering/State Prediction

- *Motion models*
  - Simple kinematic models (CV, CA, CTR…)
  - Bicycle/4 wheel models
  - Goal & Motion models – Typical behavior and motion patterns
- *IMM (Interacting Multiple Model)*
  - Parallel use of multiple motion models
- *VS (Variable Structure)-IMM*
  - Road state constraints

Situation awareness

- Predict future state and relations among road entities
- Reasoning theories
  - Fuzzy systems
  - Bayesian probability theory
  - Dempster-Shafer
- Situation assessment modules
  - Path prediction
  - Maneuver detection
  - Driver intention
  - High level events

Data fusion in cooperative systems

- Increased perception
  - Range
  - Accuracy
  - High level information (e.g. traffic)
- A new challenge for data fusion
  - Synchronization
  - Association
  - Localization is crucial for correct perception
  - Spatial and temporal alignment (requires accurate clock synchronization)
  - Enhanced situation awareness

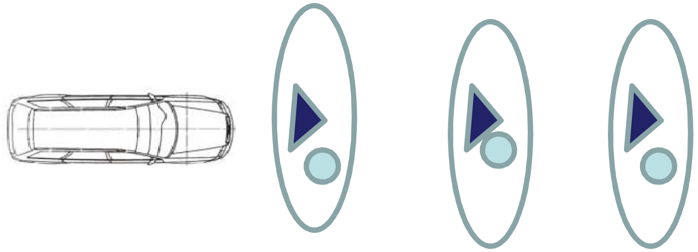

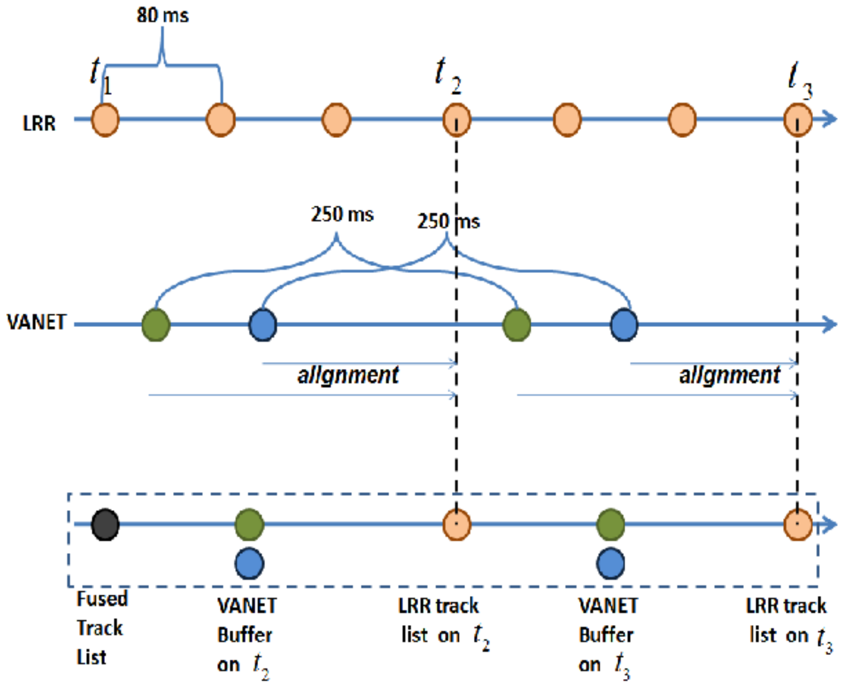

Talos vehicle – The MIT approach

- Sensor-rich design
  - 7 planar LIDARs
  - 1 roof-mounted 3D LIDAR unit
  - 15 automotive radars
- Perception modules
  - Local frame: Euclidean coordinate system with arbitrary origin where sensor information is projected
  - **Obstacle detector:** Concurrent processing of LIDAR and radar data for obstacle detection and tracking
  - Hazard detector: Identification of objects that the vehicle shouldn't drive over (i.e. potholes, curbs)
  - Lane tracking modules: Detection, filtering, tracking and fusion of lanes

Obstacle detector

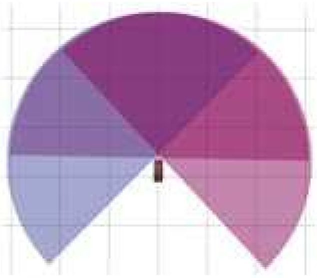
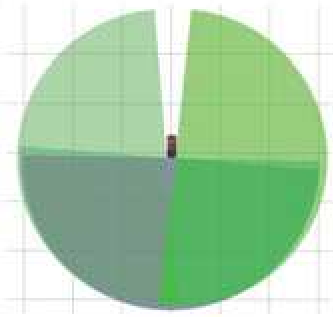
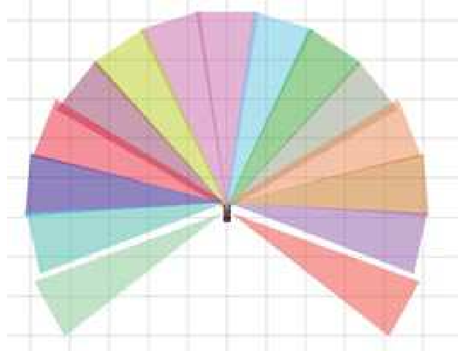

- Large number of sensors
  - comprehensive FOV
  - redundancy both within and across sensor modalities
- LIDARs: near-field obstacle detection
- Radars: moving vehicles in the far field
- The obstacle tracking system was decoupled into two subsystems:
  - one using LIDAR data
  - the other using radar data
- Each subsystem was tuned individually for a low false-positive rate;
- The output of the high-level system was the union of the subsystems' output.
- A simple data fusion scheme allowed each subsystem to be developed in a decoupled and parallel fashion.
- From a reliability perspective, this strategy could prevent a fault in one subsystem from affecting another.

LIDAR-based obstacle detection

- 7 planar LIDARs & a 3D LIDAR
- 3D LIDAR cloud much more dense
- Benefits from the use of planar LIDARs
  - Avoid large blind areas immediately around the vehicle
  - Fault tolerance in case of 3D LIDAR failure
  - Faster update rates (75 Hz) compared to the 3D LIDAR (15 Hz)
- Each LIDAR produces a stream of range and angle tuples which in turn are projected into the local frame
- The LIDAR returns contain observations of the ground and of obstacles
- The first phase of data processing is to classify each return as “ground,” “obstacle,” or “outlier” which is performed by a “front-end” module
- The planar LIDARs all share a single frontend module, whereas the 3D LIDAR has its own specialized front-end module
- In either case, their task is the same: to output a stream of points thought to correspond only to obstacles

LIDAR front-ends

Planar LIDAR	3D LIDAR
A single planar LIDAR cannot reliably differentiate between obstacles and nonflat terrain	More sophisticated obstacle-ground classifier because of high-density of data
Use of many planar LIDARs, with overlapping FOVs but different mounting heights, to ensure that nearby objects can be detected and discriminated from nonflat terrain	Identification of possible ground points → Creation of a nonparametric ground model → Identification of ppossible obstacle detections
Increased system fault tolerance	Outlier rejection challenging due to numerous outlier returns

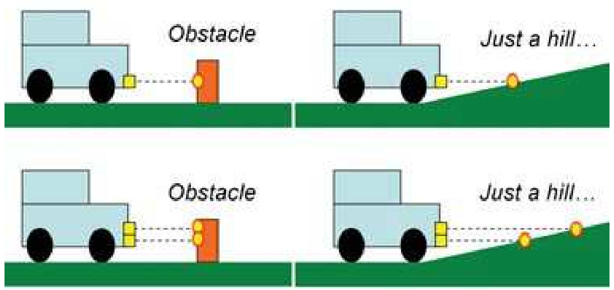

LIDAR clustering & tracking

- Tracking individual hits over time is computationally prohibitive
- Creation of *chunks* (records of multiple, spatially close-rage samples)
  - a first step towards data reduction
  - serves as a mechanism for fusing planar & 3D LIDAR data (obstacle detections from both front-ends are used to create & update chunks)
- A physical object is typically represented by more than one chunk
- Clustering chunks into groups is the input to the tracking procedure
- At each time step new groups are associate with previous ones
- Comparison of the bounding boxes of the associated groups yield velocity estimates which are noisy
- Trivial Kalman filtering takes place inside every chunk
- Strength: Production of velocity estimates for rapidly moving objects with very low latency
- Weakness: Estimating the velocity of slow moving obstacles (<3 m/s)

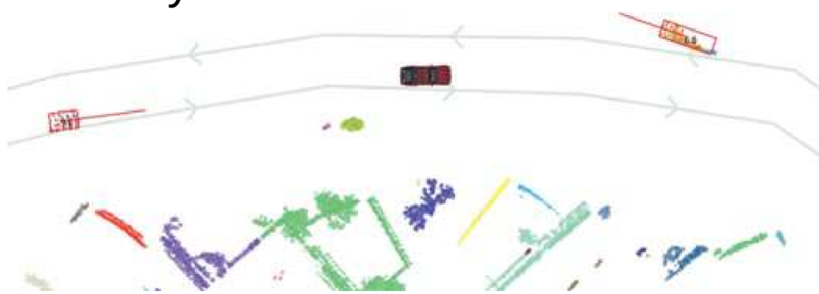

Radar-based vehicle detection

- Complements the LIDAR subsystem
- Detection of moving objects beyond the reliable detection range of LIDARs
- Range, bearing and relative velocity measurements (Doppler)
- Association of radar detections to active tracks based also on accurate velocity measurements
- Track update based on constant velocity (CV) model
- Radars cannot distinguish easily small and large objects
- To avoid false positives radars where used only for the detection of moving objects

Boss vehicle – The Carnegie Mellon approach

- Many sensors to provide the necessary redundancy and coverage
  - 11 LIDARs (incl. one 3D LIDAR)
  - 5 radars
  - 2 cameras
- Novel aspect: A pair of pointable sensor pods
- Winner of the 2007 DARPA Urban Challenge

- Perception modules
  - **<u>Moving obstacle detection and tracking</u>**: Provides a list of object hypotheses and their characteristics
  - **<u>Static obstacle detection and mapping</u>**: Generation of obstacle maps from numerous scanning lasers
  - <u>Roadmap localization</u>: Self-localization to roads with known geometry and estimation of the shape of dirt roads
  - <u>Road shape estimation</u>: Estimation of the road geometry (curvature, position, heading) when it is not known a priori

Moving obstacle detection and tracking (1/3)

- 2-layered architecture
  - For each sensor type a specific sensor layer is implemented
  - New sensor types can be added with minimal effort
  - Prediction of current object hypotheses takes place in the fusion layer
- Two tracking models
  (a) simplified bicycle model with fixed shape
  (b) point model without shape information

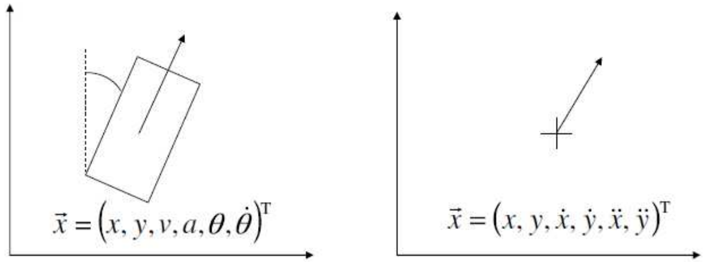

(a) (b)

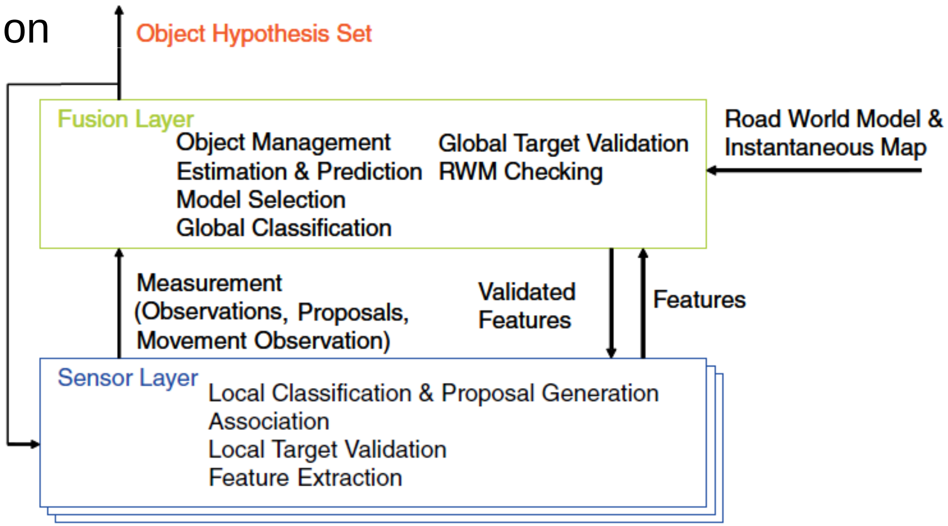

- Object hypotheses classification
  - *moving* or *not moving*
  - *observed moving* or *not observed moving*

Moving obstacle detection and tracking (2/3)

- Goal: Find all vehicles around the ego-vehicle
- Two-step validation process to remove irrelevant data from raw sensor data
  - *Sensor-specific* (sensor layer): use of velocity measurements from the radars to distinguish static ground from moving vehicles etc.
  - *Non-sensor-specific* (fusion layer): checks against road geometry and instantaneous obstacle map (list of untracked objects in 3D)
- Result: A list of validated features that potentially originate from vehicles
- Association between validated features and predicted object hypotheses
- For each extracted feature (associated or not) multiple possible interpretations are generated based on sensor characteristics
- Each generated interpretation is compared with its associated prediction & if differs significantly, or if it not associated, a new object hypothesis is formed
- In a different case the new hypothesis can replace the current one
- For each feature multiple new hypotheses can be generated
- A set of new object hypotheses is called a proposal

Moving obstacle detection and tracking (3/3)

- For each associated feature the interpretation that best fits the prediction is used to generate an observation
- An observation is used to update the state estimation for the associated object hypothesis in the fusion layer
- If no observation is generated then only the proposal exists
- A movement observation can be provided as additional info is available
- Proposals, observations and movement observations are used in the fusion layer to update the objects hypotheses list and the estimated object states
  - First the best tracking model is selected by a voting algorithm which is based on the number and type of proposals from the different sensors
  - Then the state estimate is either updated with the observation by the sensor layer or the model for the object hypothesis is switched to the best alternative
  - Finally a classification of the movement state of for each object hypothesis is carried out
  - Result: An **updated list of object hypotheses** that are accompanied by the classification of the movement state

Static obstacle detection and mapping

- Numerous scanning lasers
  - *Instantaneous obstacle map*: Used in the validation of moving obstacle hypotheses
  - *Temporally filtered obstacle maps*: Removal of moving obstacles, reduction of spurious obstacles
- Several algorithms are used to generate obstacle maps
- The curb detection algorithm is presented here
- Geometric features (i.e. curbs, bushes) are information sources for determining road shape in urban and off-road environments
- Dense LIDAR data provide sufficient information

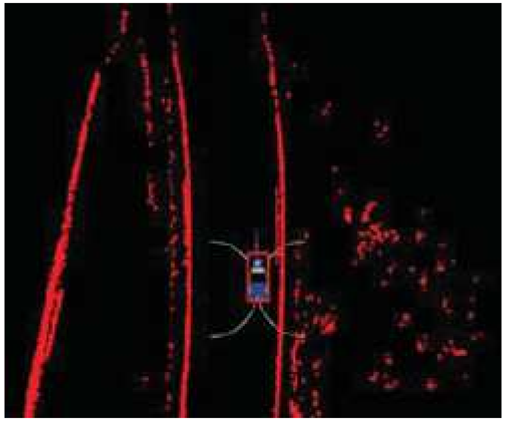

The curb detection algorithm

- Main principles
  - Road surface is assumed to be relatively flat and slow changing
  - Each LIDAR is processed independently (simplifies the algorithm)
- Three main steps
  - *Pre-processing*: mitigation of false positives due to occlusions and sparse data, formatting the data for feature extraction
  - *Wavelet-based feature extraction*: Analysis of height data through a discrete wavelet transform using the Haar wavelet, classification of points as *road* or *nonroad*
  - *Post-processing*: Extra heuristics to eliminate false positives and detect some additional nonroad points

The Haar wavelet (mother wavelet & scaling function)

$$\Psi(t) = \begin{cases} 1 & \text{if } 0 \le t < \frac{1}{2}, \\ -1 & \text{if } \frac{1}{2} < t < 1, \\ 0 & \text{otherwise,} \end{cases}$$

$$\varphi(2^j t - i) = \begin{cases} 1 \text{ if } 0 \le t < 1, \\ 0 \text{ otherwise,} \end{cases} \quad j > 0 \wedge 0 \le i \le 2^j - 1.$$

Conclusions & outlook of perception in automated driving

- Open Urban environments are too complex to handle by today's State of the Art
- First steps
  - Limited geographic extend (e.g. dedicated lanes)
  - Low speeds
  - Infrastructure support
- More robust object perception
  - Target Identification
  - Deal with clutter/occlusion
- Improved reasoning
- Generic robust and real-time perception platforms
- Well defined sensor/information sources interfaces
- Central fusion architectures

Future research needed

Perception

- Advanced fusion techniques and all around perception (360 degrees)
- Common and fault tolerant perception architecture
- Real-time perception platform (incl. plug & play concept)
- Reliable object recognition and accurate road representation
- Free space detection & object classification
- …

General

- Innovative, low-cost and reliable sensors and actuators
- X-by-wire technologies
- V2X communication (incl. standardization activities)
- Accurate positioning & enhanced digital maps
- Human machine interaction and mode (automation level) transition
- …

References

- ETSI, Intelligent Transport Systems (ITS) http://www.etsi.org/website/technologies/intelligenttransportsystems.aspx
- Highly Automated Vehicles for Intelligent Transport (HAVEit project), website: http://www.haveit-eu.org
- Hong Cheng, *Autonomous Intelligent Vehicles: Theory, Algorithms, and Implementation*, Springer-Verlag London Limited 2011
- Chris Urmson et al., "Autonomous Driving in Urban Environments: Boss and the Urban Challenge", *Journal of Field Robotics* 25(8), 425–466 (2008), 2008 Wiley Periodicals, Inc.
- John Leonard et al., "A Perception-Driven Autonomous Urban Vehicle", *Journal of Field Robotics*, 1-48 (2008), 2008 Wiley Periodicals, Inc.
- Tom Vanderbilt, "Autonomous Cars Through the Ages", *AUTOPIA - Road to the Future*, February 2012
- Jameson M. Wetmore, Driving the Dream - The History and Motivations Behind 60 Years of Automated Highway Systems in America, Automotive History Review, Summer 2003, pp. 4-19

interactIVe

Accident avoidance by active intervention for Intelligent Vehicles

Thank you.

Co-funded and supported by the European Commission

SEVENTH FRAMEWORK PROGRAMME

Dr. Angelos Amditis
Research Director @ ICCS
e-mail: a.amditis@iccs.gr
phone: +30 210 772 2398