Challenges for Time-Critical Transportation Cyber-Physical Systems

Raj Rajkumar
ECE & Robotics
Carnegie Mellon University

Carnegie Mellon

Society and Automobiles

- More than 1 million people die every year in automotive accidents globally
  - Largest global killer of ages 10-24
  - Tens of millions of injuries
  - Global annual cost of road injuries in medical care, disability and property damage is $518 billion.
- Traffic delays are expensive.
  - The average US driver spends a week stuck in traffic per year.
  - In the EU, € 80 billion wasted per year due to traffic congestion.
- Loss of independence and self-esteem of senior and disabled citizens

Carnegie Mellon

Transportation Vision

Grand Challenge: Zero fatalities

- Fewer accidents and injuries
- Higher throughput and much lower delays
- Smart traffic infrastructure
- Real-time information delivery and system-level feedback

Metrics

1. **Safety**
2. Throughput
3. Environmental Impact
4. Quality of Life

Vision: "Vehicles that do not crash"

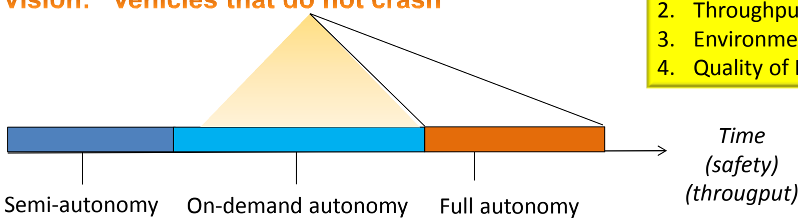

Carnegie Mellon

What must Intelligent Vehicles do?

Carnegie Mellon

Sensors

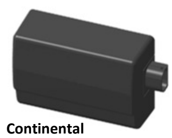

Vehicle Tracking

Vehicle Tracking

- Raw Data

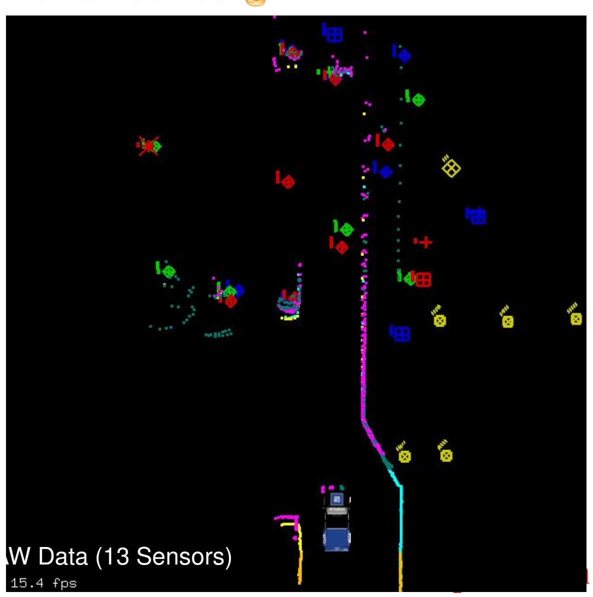

Vehicle Tracking

- Raw Data
- Feature Extraction

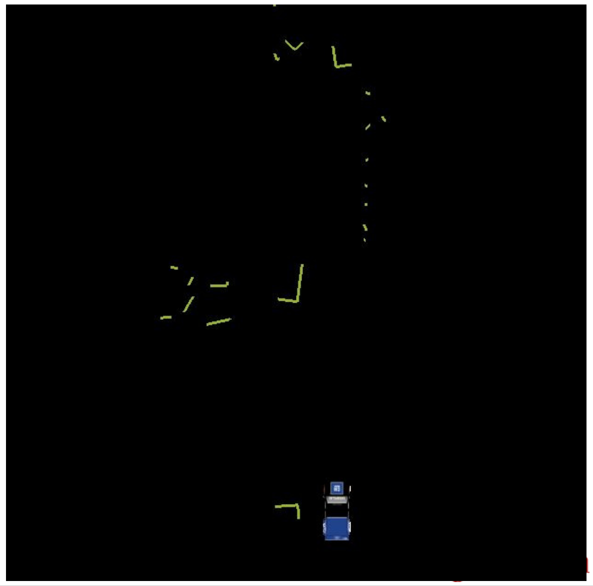

Vehicle Tracking

- Raw Data
- Feature Extraction
- Measurement Validation

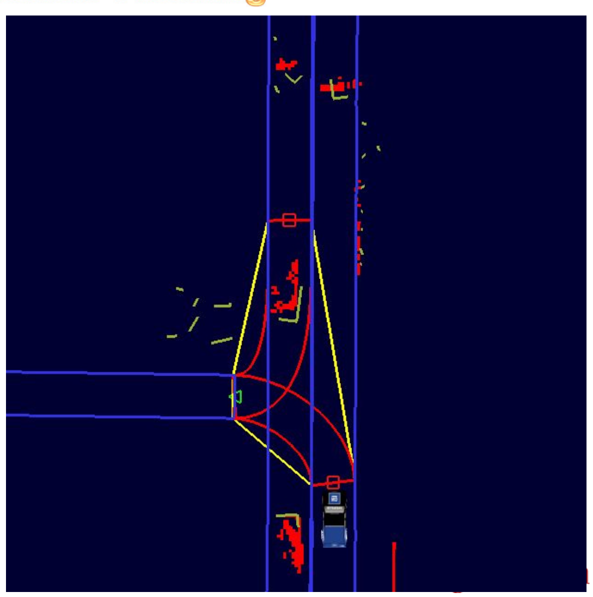

Vehicle Tracking

- Raw Data
- Feature Extraction
- Measurement Validation
- Data Association

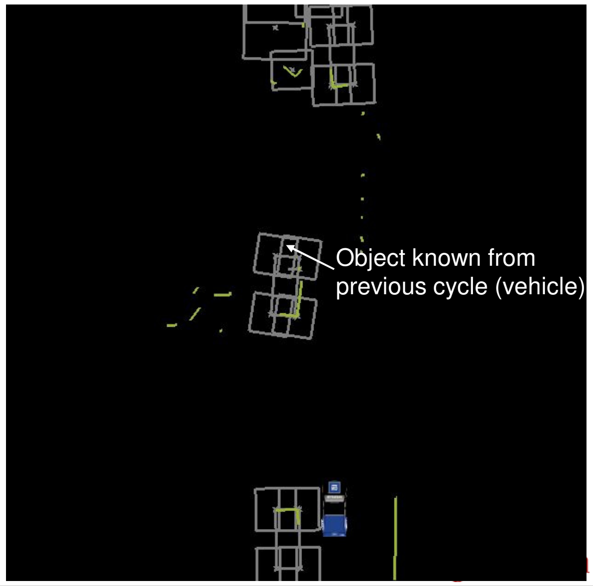

Vehicle Tracking

- Raw Data
- Feature Extraction
- Measurement Validation
- Data Association
- Proposals & Observation

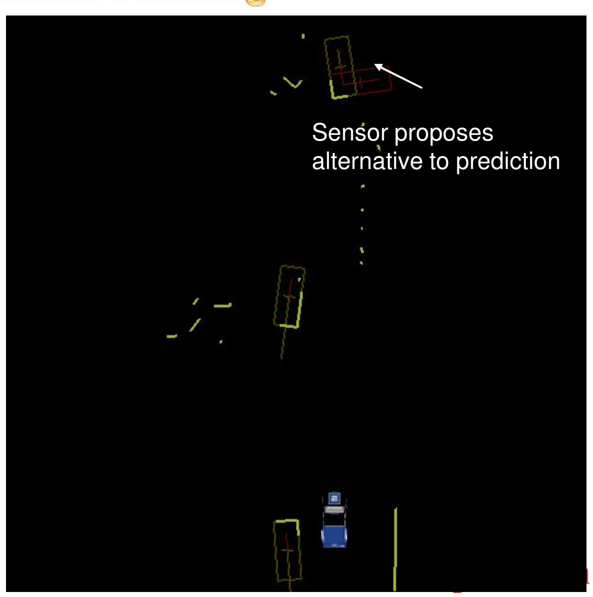

Vehicle Tracking

- Raw Data
- Feature Extraction
- Measurement Validation
- Data Association
- Proposals & Observation
- Model Voting
- Estimation
- Statistics

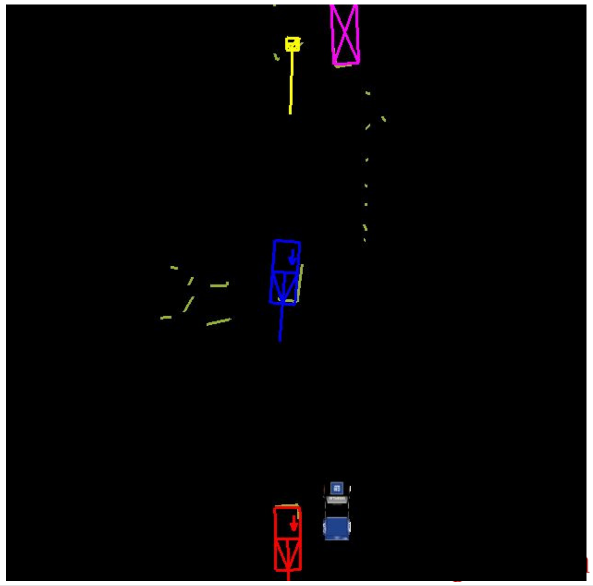

Vehicle Tracking

- Raw Data
- Feature Extraction
- Measurement Validation
- Data Association
- Proposals & Observation
- Model Voting
- Estimation
- Statistics
- Prediction

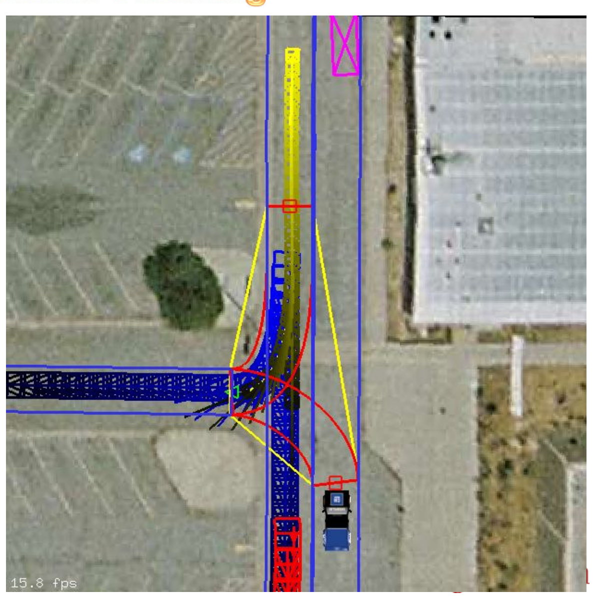

Operating in the Real World

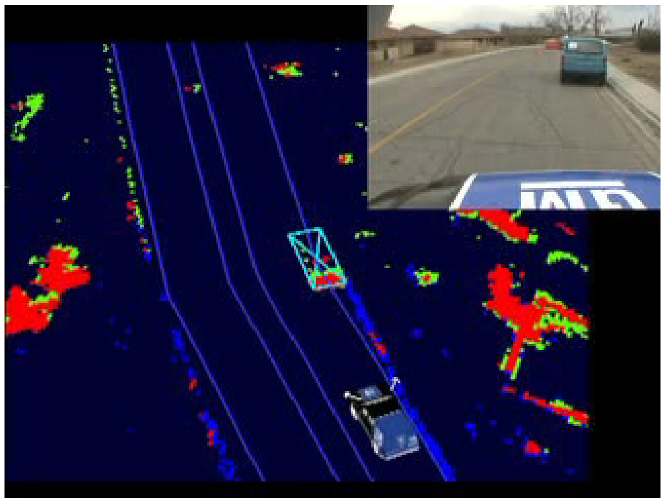

Carnegie Mellon

Logical Architecture for Autonomy

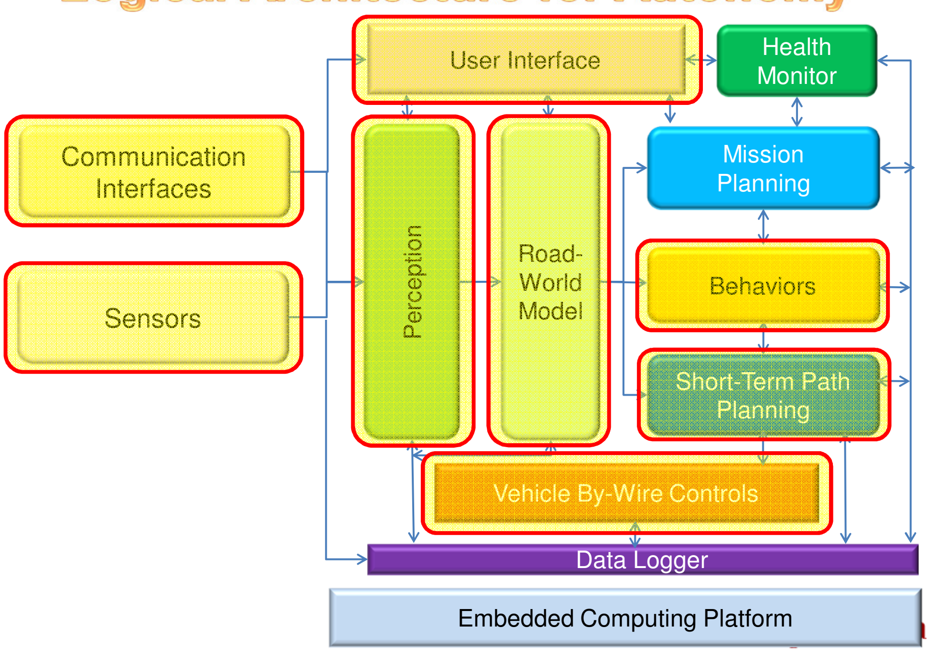

No Scarcity of Problems

Sensor Fusion & Perception

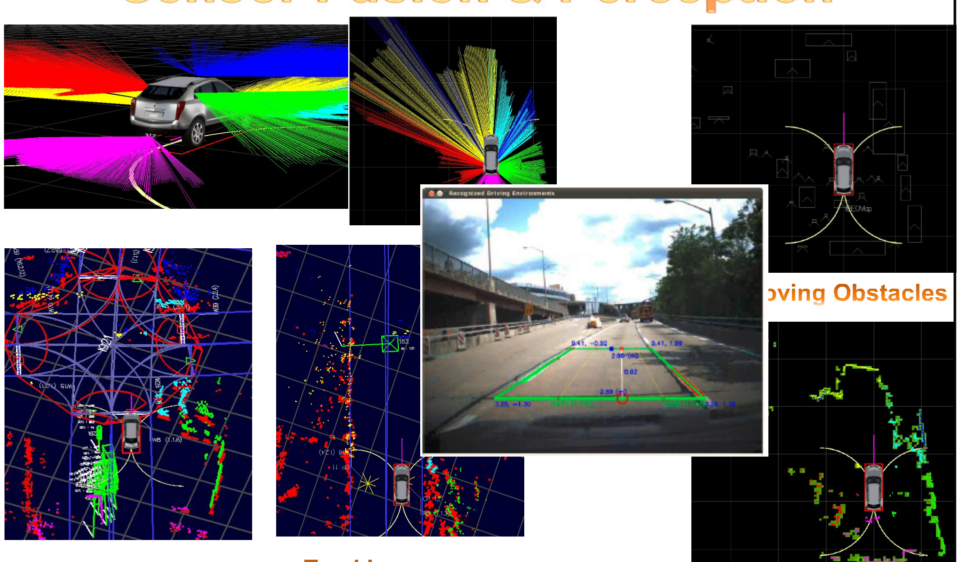

Tracking cars at intersections

Tracking cars on roadways

Static Obstacles

Encounters of the *Daily* Kind

More Encounters

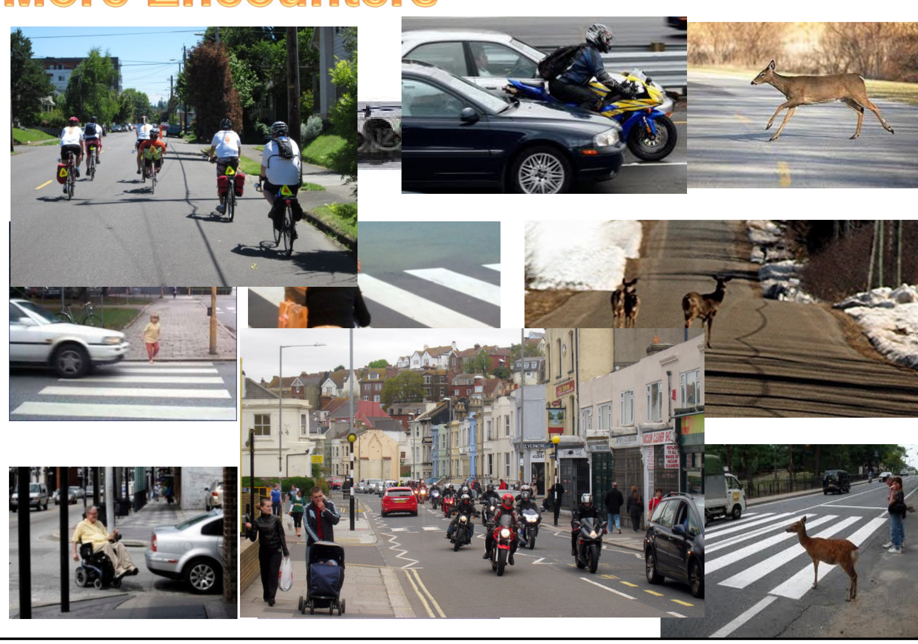

Solutions?

Carnegie Mellon

Pedestrian & Bicycle Detection

Dr. Paul Rybski

Carnegie Mellon

Workzone Detection

Prof. David Wettergreen

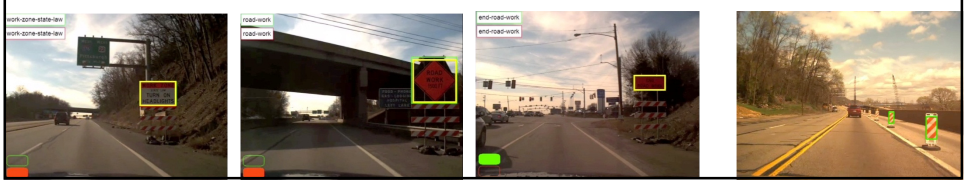

Planning

Dr. John Dolan

- Need a **planner for avoiding obstacles, over-steering/under-steering on curves, and roadside structures**.
- Computationally efficient (CPU) high-performance planner

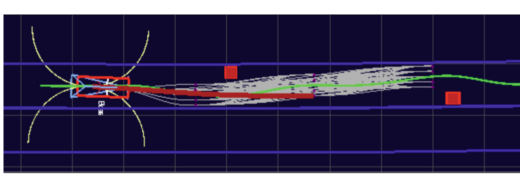

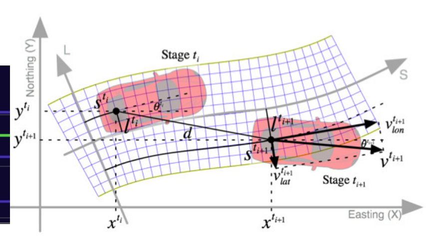

Carnegie Mellon

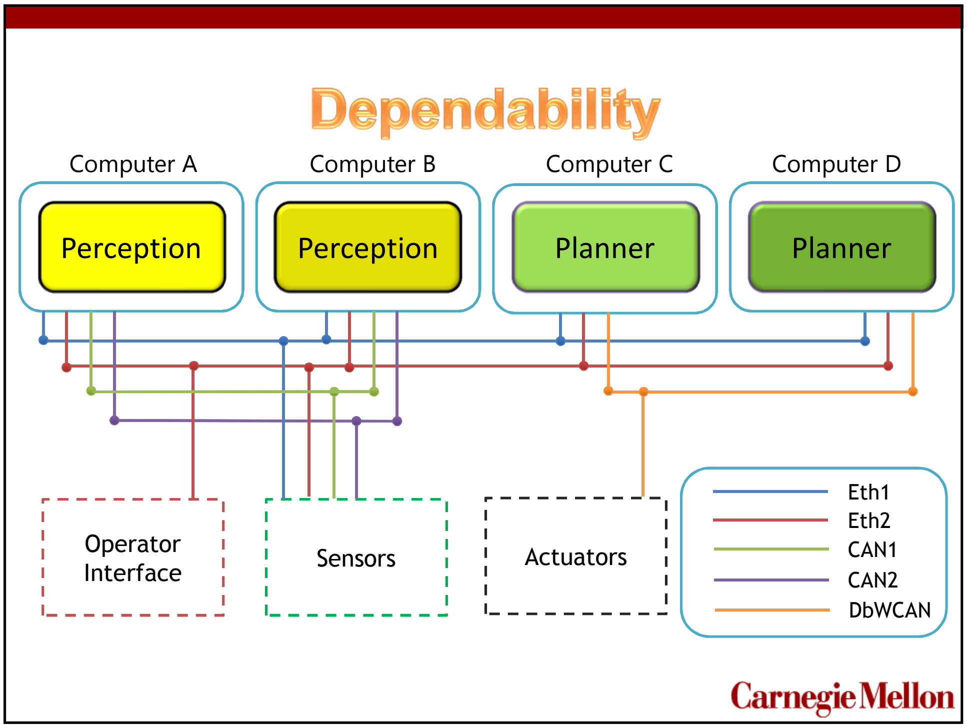

Dependability
Computer A
Computer B
Computer C
Computer D
Perception
Perception
Planner
Planner
Operator Interface
Sensors
Actuators
Eth1
Eth2
CAN1
CAN2
DbWCAN
Carnegie Mellon

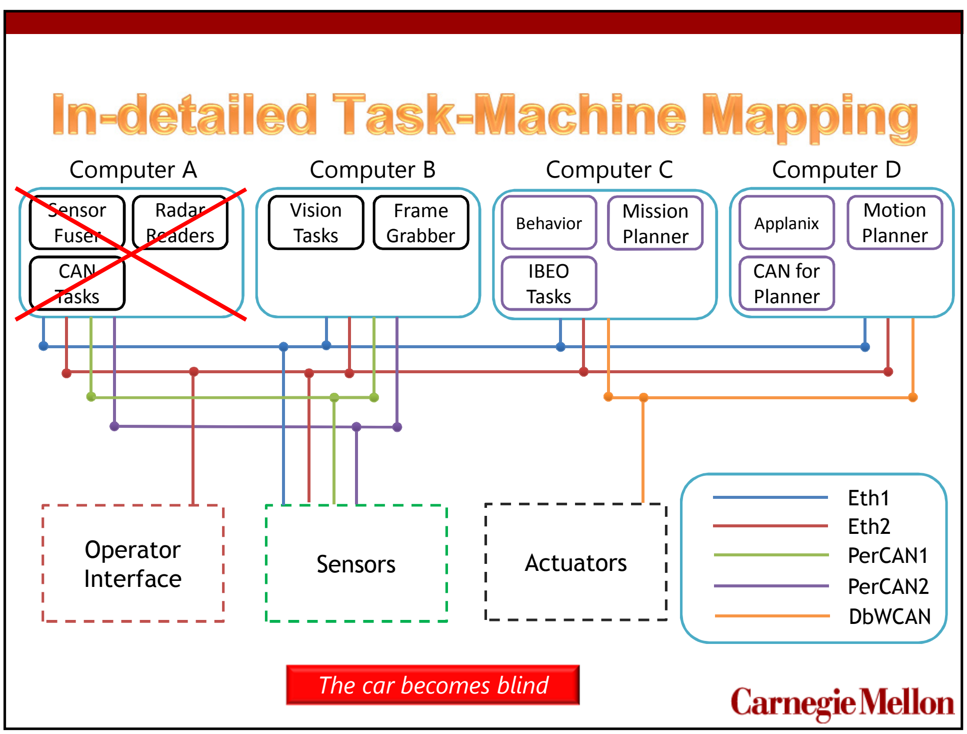

In-detailed Task-Machine Mapping
Computer A
Computer B
Computer C
Computer D
Sensor Fuser
Radar Readers
CAN Tasks
Vision Tasks
Frame Grabber
Behavior
Mission Planner
IBEO Tasks
Applanix
Motion Planner
CAN for Planner
Operator Interface
Sensors
Actuators
Eth1
Eth2
PerCAN1
PerCAN2
DbWCAN
The car becomes blind
Carnegie Mellon

Real-Time Dependability Architecture

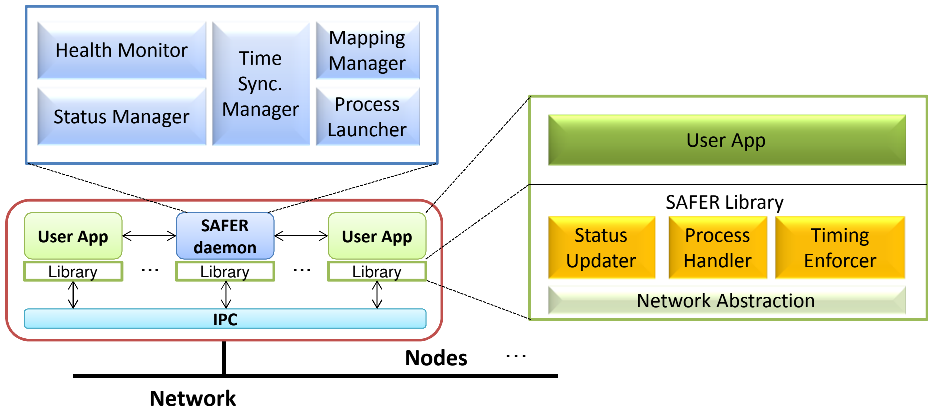

How/what do we gracefully degrade?
How do we re-deploy existing resources?

Carnegie Mellon

Coordination and Cooperation

Carnegie Mellon

V2V Protocols

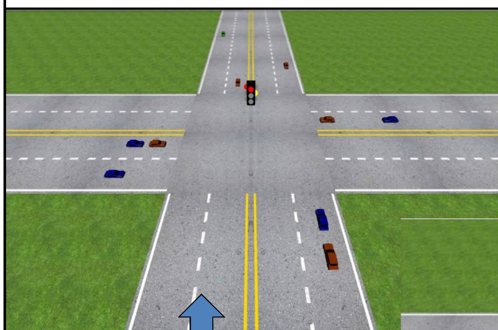

With
Traffic Lights

Using V2V

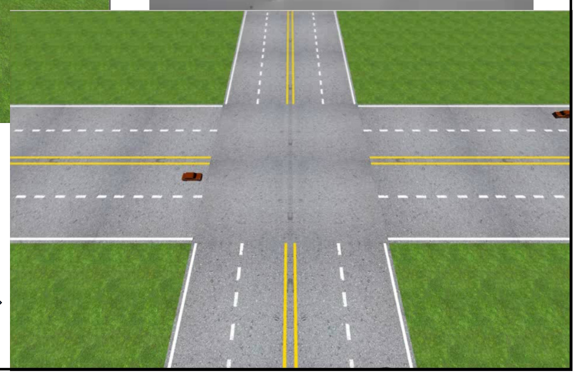

If all vehicles could communicate…

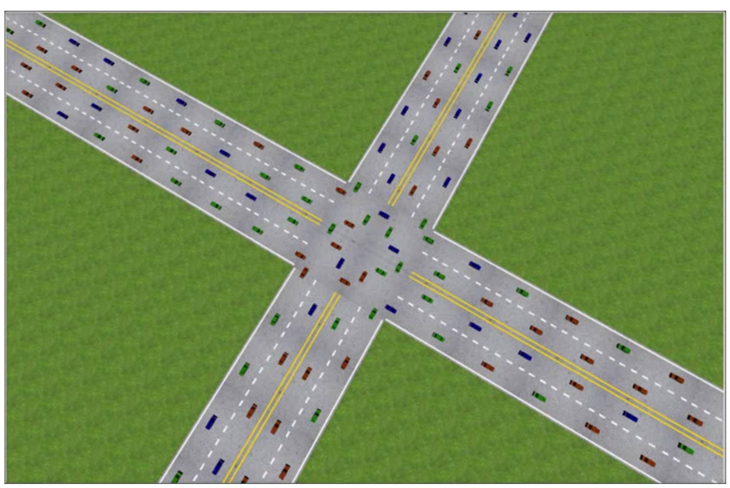

Carnegie Mellon

Challenges for Time-Critical Cyber-Physical Systems

Carnegie Mellon

Real-time Processing Challenges

- Worst-case execution time? May not be characterizable at all.
- Resource allocation and multicore allocation on multiple cores
- Parallel processing of multiple real-time tasks
- Coordination of hardware accelerators and CPUs

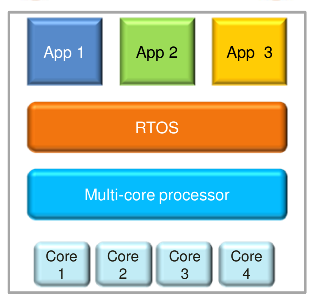

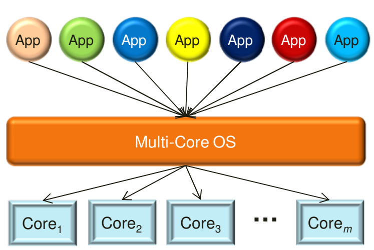

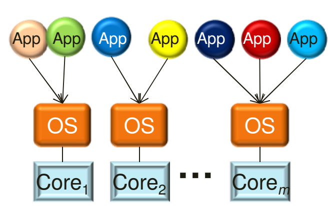

Carnegie Mellon

Cores, Cores, Everywhere

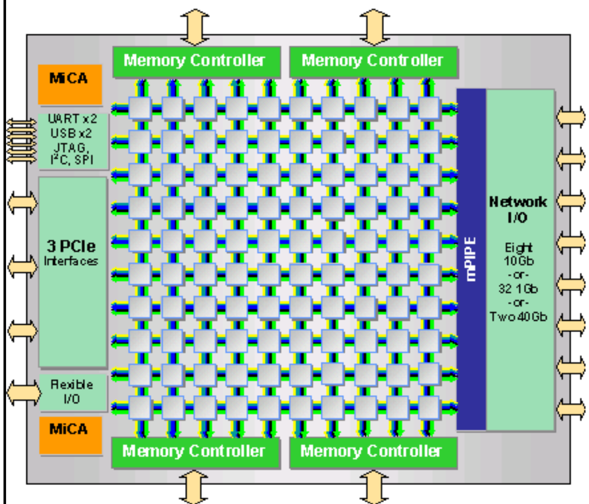

Issues

- Allocation
- Synchronization
- Communication
- Shared physical resources
- Timing predictability?
- Adaptation

CarnegieMellon

Impact of Inter-Core Cache Interference

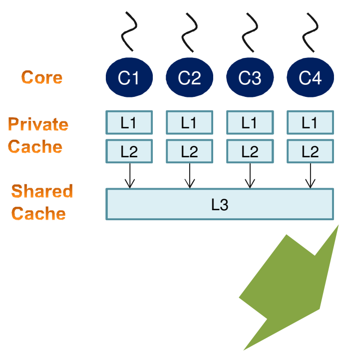

PARSEC benchmarks

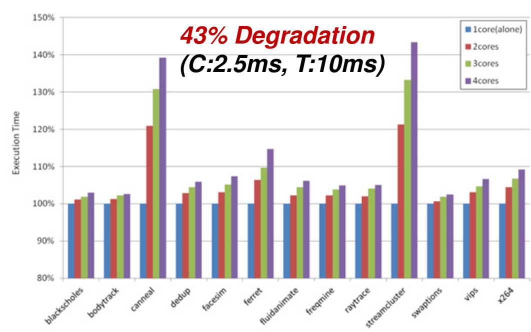

Impact of Intra-core Interference

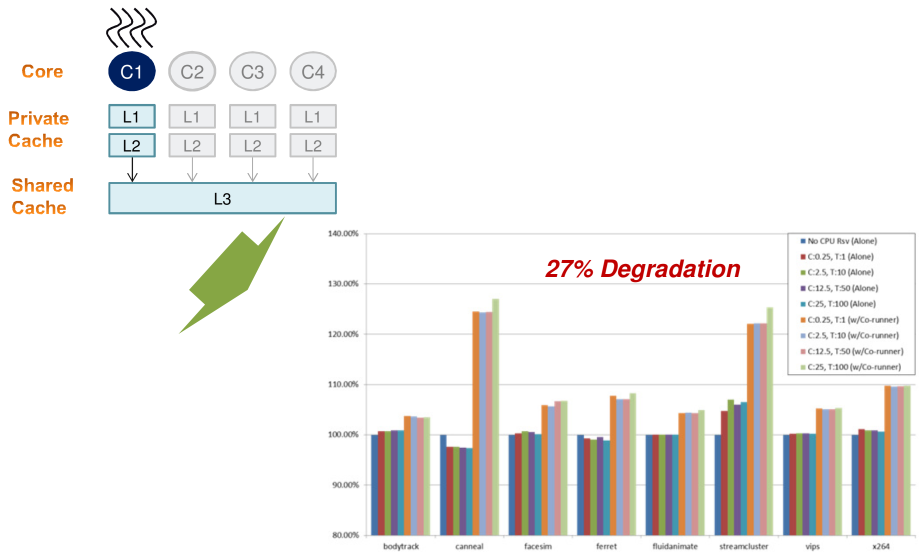

Real-Time *Parallel* Processing

- **OpenMP**:
  - Parallel processing on multicores and multiprocessors
- **CUDA**:
  - for nVidia GPUs

Carnegie Mellon

OpenMP

Performance comparison for path cost calculation as a function of the # of OpenMP threads

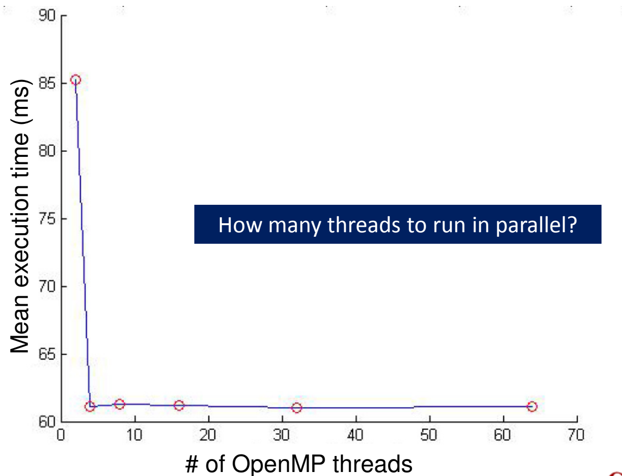

How many threads to run in parallel?

On a 4-core platform

Carnegie Mellon

CUDA vs OpenMP

Comparison of performance for path cost calculation among **only CPU**, **OpenMP** and **only GPU**

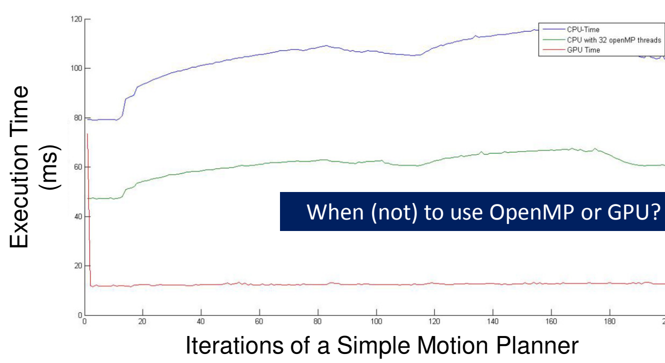

When (not) to use OpenMP or GPU?

- GPU gives a 10x improvement over the original CPU code
- GPU gives a 5x improvement over the optimal OpenMP code

GPU Parallelism does not always pay

Performance comparison for path velocity calculation with little parallelism

Execution time (10s of ms)

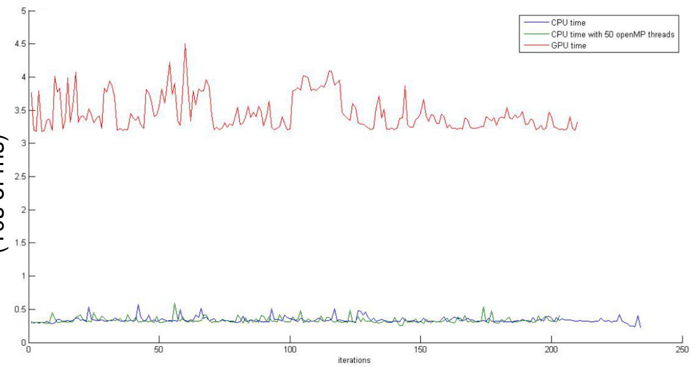

Iteration #

Carnegie Mellon

Multiple QoS Dimensions

- Timeliness
- Dependability
- Secure V2V and V2I communications
  - 10Hz Basic Safety Messages
  - Each message is likely to be encrypted and signed
  - About 750 messages per second must be authenticated and acted upon on a regular and timely basis
  - New hardware platforms required
  - New solutions required

Carnegie Mellon

Humans in the Loop

- Real-time monitoring of the driver's status
- When to take over control from the human driver?
  - Transition of control
  - Additional workloads

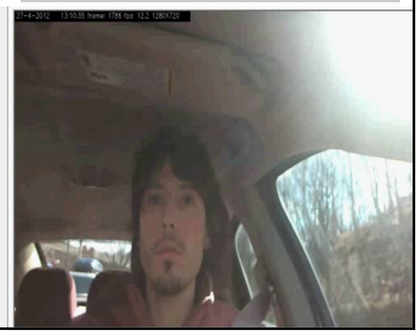

Mode Changes

- What happens **when GPS fails** (or is corrupted)?
  - What other mechanisms can kick in?
  - How to reallocate?
  - How to meet the new timing constraints during transition?
- Graceful degradation
  - Actuator adaptation
  - Can differential braking be used to steer in case of steering control failure?
  - Cut down fuel in case of brake failure?

Carnegie Mellon

Challenges Revisited

- Real-time processing, communication and coordination challenges
  - Resource allocation and multicore allocation on multiple cores
  - Parallel processing of multiple real-time tasks
  - Coordination of hardware accelerators and CPUs
  - Widely varying worst-case execution times (potentially even unbounded)
- Multiple QoS dimensions
  - Timeliness
  - Dependability
  - Security & Privacy
- Humans in the loop
  - Real-time monitoring of operator status
  - When to take over control from the human operator?
- Unexpected conditions
  - Graceful degradation and actuator adaptation
  - What happens when localization fails?

Modeling, Analysis and Synthesis of Highly Dynamic CPS in Uncertain Environments.

Carnegie Mellon

Summary

Timing requirements are dramatically different for dynamic CPS operating in unpredictable environements.

Dependability, privacy and security constraints must be satisfied simultaneously.

Carnegie Mellon