Breakout Session Report

Moderator: Tian He, University of Minnesota
Scribe: Marco Pavone, Stanford University

The audience of this breakout session comprised approximately 20-30 people from different fields of expertise. We covered two important aspects of transportation Cyber-Physical Systems: V2X networking and control.

V2X Networking

Motivation: We believe that an application-driven approach helps to identify the networking requirements, which further drive the underlying research. Potential driving applications discussed include (i) safety-related applications (e.g., intersection collision avoidance, lane merge assistant), (ii) network-wide control (e.g., congestion control, route planning, taxicab dispatch), and (iii) multimodal transit coordination (e.g., public transit co-scheduling). These applications impose different sets of critical challenges for V2X networking.

Challenge: The need for low latency: For safety related applications (V2V mostly), we need to reduce round-trip latencies down to ~10 milliseconds for real-time safety control. **Roadmap:** To maintain network connectivity continuously, we need to design fast mobility management (hand off, network route discovery, maintenance and recovery) in the presence of shrinking cell size. We also need to embed distributed coordination designs in close proximity of vehicles in exchange for short end-to-end latencies, instead of exclusively relying on remote clouds. In addition, packet prioritization mechanisms should be investigated to ensure timely delivery of delay-sensitive traffic.

Challenge: The need for high reliability: Wireless communication is inherently unreliable under various types of interferences, while safety is the paramount concern in transportation Cyber-Physical Systems. **Roadmap:** We need foundational theories and mechanisms to ensure provable and defendable reliability. Possible strategies include: (i) context-aware adaptation through closed loop feedback, (ii) reducing interference with small cell and mmWave beam forming, and (iii) better transport layer congestion control.

Challenge: The need for high scalability: We expect future transportation Cyber-Physical systems will easily scale up to tens of thousands to millions entities across different transit modes. The network performance (e.g., latency, reliability, throughput) should be highly scalable under high-density and large-scale settings even with limited available bandwidth. **Roadmap:** we need to design scalable "narrow waist" network architecture, naming and addressing, synchronization schemes to support high-dimensional multi-modal data format in an unprecedented scale and accommodate time-varying network topologies. Information scalability is another concern. We need to establish theoretical foundations to intelligently extract the value of multi-dimensional information with time constraints. In-network collaborative sensing, coordination and information fusion are needed to achieve scalability in throughput by delivering value of the information instead of raw data.

Challenge: The need for hybrid and heterogeneous support: 100% penetration rate of homogenous networking devices might not be feasible in the near future. And also due to the privacy concern, some vehicles might decline to participate in networking. **Roadmap:** We need novel distributed designs to allow both active and passive information sharing, coordination and synchronization.

Challenge: The need for open experimental platform: Test beds of transportation Cyber-Physical systems are needed to support different levels of experimentations. **Roadmap:** While small-scale test beds are valuable to evaluate safety-related designs, it is critical for the community to have open and highly programmable access to existing large-scale urban infrastructures with real-world traffic and outdoor spectrum use.

Control

Motivation: The discussion about high-priority research directions to advance the field of controls of networked automated vehicles was structured around four main pillars:

1. Device-level control: control at the hardware level, anti-lock braking systems, etc.
2. Vehicular-level control: trajectory generation with noisy sensors, localization, speed management, emission control, etc.
3. Local vehicular control: for example, interaction with pedestrians, interaction with other vehicles, and verification and validation techniques for the local interaction scenarios that can arise.
4. Global control: for example, throughput optimization and congestion control, intermodal transportation, etc.

The topic "device-level control" was considered mature, well addressed by industry, and did not trigger any further discussion.

The topic "vehicular-level control" led to two main points of discussion:

Challenge: Planning under uncertainty, that is, the generation of the trajectory that a vehicle should follow. Issues: the trajectory planning process is going to be deeply affected by several types of "noise," for example, sensor noise (where is the vehicle exactly located within a lane?) and "situational awareness" noise (is the leading vehicle autonomous or human-driven? Where is a pedestrian going? What is the intention of the surrounding vehicles? etc.). **Roadmap:** (i) Modeling: Characterization of the different types of noises/disturbances affecting the trajectory generation process. (ii) Improved sensing: robust mapping and localization techniques in different weather conditions, possibly with a rigorous characterization of the confidence level. (iii) Extraction of contextual information from data under temporal constraints: novel approaches for fast scene understanding under temporal constraints, for example, algorithms that guarantee that x% of the "meaning" of a data set is extracted within y milliseconds. (iv) Planning algorithms: planning algorithms that take into account uncertainty and come with either probabilistic or worst-case guarantees. What should a planning algorithm do in case all available trajectories might lead to negative outcomes?

Challenge: Diagnostics: autonomous vehicles will be equipped with several heterogeneous and networked sensors. There is need for diagnostics algorithms that constantly monitor the "health" of the sensors. **Roadmap**: (i) Modeling: characterize the typical failures for sensors used by automated vehicles (for example, stopping failures, byzantine failures, etc.) (ii) Diagnostics: develop diagnostics algorithms that can detect in less than a second sensor malfunctions. (iii) Failure recovery: study *safe* failure recovery techniques (for example, sudden stop of the vehicle versus online reconfiguration of the sensing architecture).

The topic "local inter-vehicular control" led to two main points of discussion:

Challenge: Verification, validation, and "correct by construction" algorithms: once a vehicle has been equipped with some degree of automation, its safety needs to be proven. This seems to be a topic where principled approaches (beyond extensive experimentation) are missing. One of the key challenges is that the local control of autonomous vehicles entails both a discrete component (e.g., rules of the road) and a continuous component (e.g., the physics of the vehicles). While well-established approaches exist for the discrete case and for the continuous case, current approaches that address the hybrid discrete/continuous setting (for example, hybrid systems control) appear limited. A related challenge is the one of devising high-level decision making algorithms that are *a priori* correct. **Roadmap**: (i) Modeling: study appropriate frameworks to rigorously pose the question of safety (is the Lyapunov framework suitable?). There is a need for new "stability" metrics. (ii) Modeling (part two): develop methods to rapidly generate fault trees for autonomous driving, and to abstract or compress the huge amount of potential scenarios into a manageable subset amenable to computation. (iii) Analysis/design tools: develop *formal* verification analysis *and* synthesis techniques that address both the discrete and continuous components of the problem. A promising approach would be to rely on formal methods techniques and develop synthesis approaches that yield *cooperative* high-level decision-making algorithms that are correct-by-construction (within a given operational envelope).

Challenge: Robust verification techniques: robust stability analysis has been a topic of research for the past three decades, however, the topic of *robust* verification of cyber-physical systems (needed, for example, to address the difference between a *model* of local interaction and the real setting) does not appear sufficiently developed. **Roadmap**: design *robust* analysis/synthesis verification techniques for networked autonomous vehicles.

Finally, the topic of "global control" led to one main point:

Challenge: System-wide coordination: system-level aspects require analysis tools at the interface between robotics and transportation science, which appear to be currently *unavailable*. Key questions to be addressed include: How to optimally coordinate thousands of automated vehicles? Would autonomous cars decrease the number of cars that are needed to serve the population? What is the achievable throughput of transportation systems with automated vehicles? What is the cost of the system and is there potential for economies of scale? **Roadmap:** (i) Analysis tools for optimal routing: develop analysis tools to address system-level coordination problems for automated vehicles, for example, optimized routing with charging constraints, global emission control, intermodal deployment, etc. Can current tools (e.g., queuing theory) be extended to this case? (ii) Routing algorithms: devise algorithmic approaches to dynamically route the vehicles by taking into account safety, demand uncertainty, and charging constraint. (iii) Study of sustainability benefits: study the sustainability benefits in terms, for example, of usages cases (point-to-point versus last-mile), number of shared automated vehicles to serve the population, cost of the system and potential economies of scale, potential for congestion reduction, etc.

Finally, as for the networking aspect, an overarching need is the development of **test beds** where automated vehicles provide *sustained* (that is, daily, non-episodic) transportation services within a small, gated community.