A Community Report of the 2008 High Confidence Transportation Cyber-Physical Systems (HCTCPS) Workshop

Community Draft

Updated on: July 22, 2009

Authors: Radha Poovendran, University of Washington (Technical Co-chair)
Raj Rajkumar, Carnegie Mellon University (Technical Co-chair)
David Corman, The Boeing Company (Area Co-chair: Aviation)
Jim Paunicka, The Boeing Company (Area Co-chair: Aviation)
William P. Milam, Ford Motor Company (Area Co-chair: Auto)
K Venkatesh Prasad, Ford Motor Company (Area Co-chair: Auto)
Shige Wang, General Motors Company (Area Co-chair: Auto)
Jim Barhorst, The Boeing Company
Christopher Gill, Washington University at St. Louis
Sandeep Gupta, Arizona State University
Krishna Sampigethaya, The Boeing Company
Jonathan Sprinkle, University of Arizona
Douglas Stuart, The Boeing Company
Wayne Wolf, Georgia Tech. University
Rahul Mangharam, University of Pennsylvania

Contact: All inquiries about this community report must be directed to grant PI Dr. Radha Poovendran (rp3@u.washington.edu).

Disclaimer: This workshop was funded in part by NSF Grant # 0850549 to the University of Washington, as well as by the sponsorship from Boeing Company and the University of Washington. All statements and conclusions of the report are that of the community and do not reflect any position of the NSF, United States Government, The Boeing Company, Ford Motor Company or General Motors Company.

Table of Contents

Community Draft

EXECUTIVE SUMMARY

Transportation cyber-physical systems (CPS) are intimately connected to the daily life and economic fabric of the United States. Everyday tasks or events, such as the drive or train ride home or an airplane landing, depend on the complex yet flawless interactions between functions within the vehicles' computer (cyber) systems and physical systems, all typically mediated by human operators or end-users. Today's transportation systems are being designed to be more competitive within their respective industries, with more complex features and capabilities to support increased energy efficiency and (in the case of military CPS) to maintain a capability edge over our adversaries. From an industrial competitiveness and economic perspective, these industries are extremely important to the U.S. For example, the transportation manufacturing and services sector contributed about 9% of the Gross Domestic Product (GDP) of the U.S. during 2007[1] and around 9 million jobs during 1999-2007.[2] Prior to the current recession the Detroit automotive manufacturers supported 6.1 million automotive-related jobs.[3] The aerospace sector provides over 2 million jobs and accounts for more than 20% of the value of total capital goods exported.[4,5] The ability to affordably design and field energy-efficient transportation CPS supports U.S. economic, national security, and environmental objectives. In the case of military transportation CPS, the ability to affordably design, build, and verify the safety and correctness of new and leading-edge on-board capabilities is critical to national security.

Important challenges threaten the ability of the transportation sectors to quickly and affordably meet future demands for growth in capacity and capability. The complexity of the software programmed into transportation platform computers may soon make it too costly to design, test, and verify the feature sets needed for competitiveness. Moreover, our current inability to deal with multi-system complexity will threaten our ability to network multiple transportation platforms together to create the transportation system of the future that would feature breakthrough safety, advanced features, affordability, and dramatic reductions in energy consumption.

This report summarizes the results of the High Confidence Transportation Cyber Physical Systems (HCTCPS) Workshop held November 18-20, 2008. The workshop brought together a diverse collection of stakeholders (researchers, certifiers, policy makers, end users) from a broad spectrum of science and technology (control engineering, certification, software engineering, aerospace, automotive, rail) to identify the community's shared vision of the technical and economic challenges we face in developing the next-generation, high-confidence, transportation cyber-physical systems. Our report establishes a 15-year roadmap setting forth an aggressive research agenda that will enable breakthroughs in affordable design, analysis and verification of more capable and efficient cyber-physical systems. The research agenda also incorporates adaptation and self-healing, breakthroughs in the interaction between cyber-physical systems and their passengers and operators, and the changes in engineering practice and education needed to institutionalize these advances.

Community Draft

[1] http://www.bea.gov/national/index.htm#gdp
[2] http://www.bts.gov/
[3] America's Auto Industry Economic Contributions & Competitive Challenges, Automotive Trade Council Report. Jan. 2008 8. Source: http://www.autotradecouncil.org/Upload/Domestic%20Auto%20%20Contributions.pdf
[4] http://www.aia-aerospace.org/industry_information/economics/year_end_review_and_forecast/
[5] http://www.ita.doc.gov/press/press_releases/2009/export-factsheet_021109.pdf

Background and Scope

Cyber-physical systems (CPS) are physical and engineered systems whose operations are monitored, coordinated, controlled, and integrated by a computing and communication core. CPS are intimately connected to the daily life and economic fabric of the U.S. Everyday tasks or events, such as the drive or train commute home or an airplane landing, all depend on the complex yet flawless interactions between functions within computer (cyber) systems and physical systems, such as an automobile or an aircraft, all typically mediated by human operators or end-users. Transportation systems are CPS on a grand scale. As they grow to meet the future demands of society, they will become increasingly complex systems-of-systems, often involving time-critical interactions between purely physical elements and highly intangible cyber elements. CPS technology is fundamental to the development, analysis, and verification of these systems, and enormous research challenges must be met to create a healthy, competitive future for our nation's transportation infrastructure. Many of these research challenges are clearly in the CPS domain. However, additional progress must also occur in related technologies including lightweight materials and "green" fuels to mention two, as well as in public policy, to assure this vibrant future.

In this report we focus on high-confidence transportation cyber-physical systems (HCTCPS), a timely and critical area that promises to accelerate the development of individual transportation systems while creating a framework to maximize the sharing of tools, design and manufacturing processes, technologies, architectural elements, and people skills across transportation systems. The HCTCPS effort draws upon the collective wealth of our nation's unmatched concentration of world-class research universities and the unquestionable combined power of our aerospace, automotive, and rail R&D and manufacturing skills.

The national workshop for research on High-Confidence Transportation Cyber-Physical Systems: Automotive, Aviation and Rail was held on November 18-20, 2008 in Washington, D.C. This workshop was a culmination of 12 weekly planning teleconference meetings that started in July 2008. Participants of the planning meetings included representatives from academia and the automotive, aviation, and rail industries. The workshop was sponsored by the Federal agencies that participate in the High Confidence Software and Systems (HCSS) Coordinating Group (CG) of the Networking and Information Technology Research and Development (NITRD) Program, including NSF, AFRL, NIST, NASA FAA, and NTSB. Planning and technical support for the workshop were provided by the National Coordination Office (NCO) for NITRD.

The best insurance for the future vitality of the transportation sector economy will be a rapidly retrained workforce for today and a future workforce skilled to meet the needs of designing, deploying, and maintaining future transportation cyber-physical systems.

Purposes and Format of Workshop

The purpose of the HCTCPS workshop was to provide an open forum for leaders and visionaries from industry, research laboratories, academia, and government to identify shared development and deployment needs and opportunities for CPS in the aviation, automotive, and rail sectors. The workshop's primary aim was to establish a compelling CPS science and technology research

Community Draft

agenda within and across these sectors to address societal demands such as increased mobility, comfort, convenience, and accident reduction amid increasing system complexity. The workshop facilitated discussions to identify common objectives, challenges, shareable best practices, and synergies for industrial and cross-domain involvement, by soliciting domain stakeholders' perspectives on the current state and future advances of the transportation community as well as other critical infrastructures such as energy and medical that face similar challenges. In addition, the workshop included a discussion by the Air Force Research Laboratory of the open literature components of its mixed-criticality architectural requirements program to identify topics that are vital to national security. The workshop focused on developing a persuasive and competitive agenda including presentations that elucidated global market and shareholder value-based drivers for transportation, and emphasized strong consideration of European Union (EU) views on global trends and advances.

The HCTCPS workshop included plenary and panel discussions and breakout sessions. The following five topics, essential to HCTCPS, were addressed:

1. **Verification and Validation of Transportation CPS**. In striving to respond to driving market forces such as fuel efficiency, traffic management, and pervasive connectivity – as well as regulatory needs such as reduced gas emissions – canonical CPS designs for automobiles, airplanes, and rail are now approaching the most complex CPS ever developed. This complexity presents a significant technical challenge in the verification and regulatory requirements that add effort, time, and costs unacceptable to customers and manufacturers. Without dramatic improvements in the capability and usage of automated, time-efficient verification and validation technologies, the costs of developing next-generation vehicles may be prohibitive.
2. **Mixed Criticality.** In order to optimize parameters such as weight and volume for manned and unmanned vehicles, it is tempting to mix non-safety and safety-critical functionality within the same computation platform. Similarly, future roadways and airspace will have different classes of vehicles exchanging content of mixed criticality. However, current approaches to certifiability and assurance of such systems warrant that *all* of the functionality be verified and tested at the highest criticality level, which can be an unnecessarily expensive proposition. Advances in time-space partitioning methods, hardware with predictable and testable behavior, and reference architectures are becoming important for transportation, due to their ability to isolate critical from non-critical content.
3. **Platform and Infrastructure Challenges.** The interactions of cyber-physical components in transportation domains consist of both mechanical and electrical engineering as well as computer science and computer engineering elements. Platform and infrastructure elements are usually in a hierarchical structure and, depending on the application, can range from sensors, actuators, and processors, to cars and airplanes, to highways, airports, and train stations. Additionally, information processing is used both for gathering and disseminating data to the distributable computation nodes that make up each vehicle and associated ground infrastructure. In current vehicles, computing and communications modules are used for electronic control of the vehicle's physical components such as power-train, wheels and chassis, and traffic monitoring as well as the vehicle's infotainment and comfort settings. The seamless integration of the communications and computation modules on-board as well as off-board requires a

Community Draft

foundation for reliable and timely communication, with the ability to integrate new components into the system.

4. **Autonomy and Control.** The evolution of transportation systems toward increased autonomy promises to increase traffic throughput, reduce fuel costs, and even save lives. But these goals cannot be realized by replicating technology designed to work on a single vehicle. The challenges of scalability affect both how to control and how to collect operational information about the system. At scale, influence is arguably more important than direct control, especially where human operators have tremendous authority, for example, in ground vehicles and next-generation air traffic management, e.g., Next Generation (NextGen). The challenges of observation involve sensing a vehicle's health, as well as the interaction of vehicles, not directly, but through sensors.
5. **Infotronics and Infotainment.** *Infotronics* and *infotainment* (also called *e-Enabling* in aviation) are the information-level abstractions of the interactions between the built-in systems and across the brought-in (on-board passenger electronics) and beamed-in systems (via V2X or "vehicle-to-X" communication where X can be an arbitrary airborne, space, or ground-based entity). Advances in this area are vital for the ability to get the right information at the right place and time within the vehicle and its connected elements to deliver the expected set of services to the vehicle itself and to the driver/operator, passengers, and to society at large, so that everyone has all the information they need to react at the very best moment. Vehicular communication and coordination, involving several microprocessors as well as sensors and actuators, have primary problems with the partitioning of resources and coordination of control loops on a single communication bus. As V2X extends the boundaries of a single vehicle to a network of vehicles and ground infrastructure, major challenges arise concerning system interoperability and the need to better understand real-time distributed computing for spatio-temporal networked control systems.

Each breakout group in the HCTCPS workshop was asked to define a coherent and compelling vision for its topic area to support both the shared and individual needs of the different sectors; to summarize the state-of-the-art in practice, development, and research; and to identify and create a roadmap for the R&D needs and challenges. Most of the presentations of the breakout groups, keynote speakers, and panelists, along with all of the submitted position statements of participants, are available at the workshop web site: http://www.ee.washington.edu/research/nsl/aar-cps

Grand Challenges

As this report is being written, the U.S. is in the midst of a global economic crisis of a magnitude not seen since the 1930's. The impact of this crisis will be felt in the coming years, and yet many grand challenges transcend current economic conditions and present research and development opportunities for CPS that are of potentially significant technological and economic benefit. These challenges and opportunities are likely to have a large-scale impact in re-establishing American preeminence in the world economy, ensuring the safety of our citizens, and the preservation of our strategic interests.

Transportation systems are grand-scale cyber-physical systems with many cross-cutting capabilities and interactions that are pervasive across the transportation sector. Demands for higher system performance and lower costs are requiring next-generation transportation systems

Community Draft

be highly networked and dynamic in nature, and their complexity is growing at an exponential rate. Some examples are noteworthy. The 747-400 that first flew in the late 1980's contained software that required 10 megabytes (MB) of memory. The software for the Boeing 777 that first flew in the early 1990's was an order of magnitude larger at 100 MB, on the order of 10 million source lines of code (SLOC). As systems such as the 787 evolve, software size and system complexity will increase by two or more orders of magnitude. Wireless networking both inside and outside the skin of the aircraft will increase rapidly. The trend for automotive systems is similar. Today's automobile may have on the order of 10 million SLOC with 1-10 networks. To accommodate emerging needs, the automobile of the future will have 100's of millions of SLOC and 10's of networks internal and external to the vehicle.

In the following paragraphs we describe a list of grand challenges and the important role of CPS in achieving the vision. Our primary focus in this report is on the challenges in developing a future affordable, safe, and secure automotive and aerospace transportation system. We also briefly address two national security challenges in an area where cyber-physical systems can make significant contributions: providing persistent surveillance and enabling zero-collateral-damage destruction of time-critical targets. For each grand challenge, we identify subsidiary elements of the challenge and describe how CPS technology contributes to meeting the challenge. While each of the transportation domains (air, automotive, and rail) has unique features, our vision for transportation reflects many common threads – the need for safety, cost-effectiveness to operate, profitable producibility, significantly reduced energy footprint, and the need for a feature-rich and comfortable user experience.

Grand Challenge – Developing and delivering a transportation system that includes uncompromised all-weather safety and security, comfort and convenience anywhere and anytime, unmatched performance, coupled with a dramatic, sustainable reduction in the environmental footprint, and at an affordable rate for the user or operator.

Substantial safety and affordability advances have been made in the automotive transportation industry through the integration of driver warning and assist systems, a breakthrough in both affordability and driver safety. Congestion reduction, which requires both on-demand and autonomous control algorithms inside and outside the vehicle, can also be achieved with the integration of advanced networking and new sensing devices. Progress can include highway-only autonomous driving, cross-traffic collision avoidance, and fault-tolerant single and multiple vehicular control. This will require advanced CPS research in engineering large complex systems, advanced sensing and control algorithms, and affordable fail-safe operational and fault-tolerant software architectures. Additionally, CPS research in safety-critical mobile information services and integrated diagnostics and prognostics are essential to achieving safety and affordability goals.

Safety – Safety and security are preeminent requirements for any transportation system. These translate into behavioral challenges for individual vehicles in the presence of increasingly crowded airways and roads. Aviation safety requires ultra-high-reliability behavior with certification by the FAA. While the occasional "Blue Screen" may be painful in the office environment, it can have extreme consequences in the air. Automotive safety is similarly regulated by National Highway Transportation Safety Administration (NHTSA). Implicit in safe behavior is the need for mixed criticality requiring that the safety-critical CPS vehicular components be certified (aviation) or certifiable (automotive), and that the lower-criticality software elements, e.g., passenger entertainment, do not negatively affect higher-criticality

Community Draft

software. Today's systems exhibit unprecedented levels of safety to the public. The challenge is to match and increase the level of safety and security as the complexity of transportation elements and interactions between elements increase exponentially.

Certification Certification of today's systems is a long, difficult, and extremely costly process. As today's vehicles morph into the next generation and beyond with the exponential increase in system complexity, the certification challenge becomes immense, Without advances in research that can support validation and verification of the interactions between system modules, the cost and time frame for developing next-generation CPS of this scale may be prohibitive, significantly impacting the competitiveness of the American transportation industry.

Measurable Efficiency The airways of today are highly congested. The Vehicle-to-Vehicle (V2V) separation requirements for today's air traffic management infrastructure and existing avionics suites result in substantial delays, especially during periods of high demand or in reaction to system shocks such as isolated pockets of adverse weather. The cost in energy consumption alone resulting from these delays is substantial. Advanced CPS technology can make a large contribution to providing safe travel in future air traffic control sectors, with significantly improved on-time performance, by integrating new affordable sensing and intelligent control algorithms in a widely distributed system.

Dual use for Defense and National Emergencies – Increasingly, we will have to deal with the presence of manned and unmanned systems in the same space (either air or on the ground). When the Global Hawk UAV first flew in controlled airspace, air traffic was re-routed for miles around the flight path, and ground traffic was halted in the vicinity while the UAV approached the airport. Recently, Predator UAVs performed reconnaissance missions in the vicinity of Fargo, N.D., to provide data on risk of flooding. While not yet a common occurrence, the incidence of such events in the air domain is becoming more frequent. The DARPA Grand Challenge of 2007 similarly showed the potential for autonomous ground vehicles to traverse through a crowded city landscape. The challenge of the future is to ensure that encounters of manned and unmanned vehicles, regardless of the domain, are safe and secure. This will require unprecedented advances in CPS technology for vehicle interactions in dynamic environments.

Economic Impact – Ensuring cost-effectiveness and profitability of production in the automotive and aviation industries is critical to ensure American competitiveness in the world marketplace. While the production scales may differ, the value proposition is to, for example, design vehicles that provide system behaviors demanded by the customer and regulating authorities, yet require minimal development, software, and certification costs, enable push-button Verification and Validation (V&V), and demand zero prototypes, with zero defects and no recalls. Achieving this on a small scale is difficult; achieving this on the scale of future systems poses enormous challenges and requires urgent attention. The retention of high-quality jobs in automotive, aviation, and rail engineering is directly tied to the ability to increase innovation, quality, and productivity in the design, development, production, certification, operations, and maintenance processes. CPS research in model-based techniques, automation, and virtual analysis are critical progressions on the path to achieve "correct by design" automotive systems. Moreover, product-focused technologies will be required – including software reuse, architectures, real-time theory, languages, and product-line architectures – to achieve system affordability by recouping investments across multiple system developments. Integrated vehicle health management (IVHM) that enhances system reliability and reduces logistics costs, will lower recurring as well as maintenance costs. A recent study of fighter

Community Draft

aircraft has shown that incorporation of IVHM technology may, in fact, pay for the cost of incorporating advanced networking technology on the platform.

Growing using Green Systems – Breakthroughs in fuel economy and use of renewable fuels are critical to increase cost-efficiency and reduce environmental impact. Major steps include the development of plug-in electric and hybrid-electric power trains and the development of hydrogen fuel cell power plants. While CPS research cannot contribute directly to the development of renewable fuels, it can play a major role in making the integrated system that utilizes these fuels more efficient and economical to operate. CPS research is important in energy harvesting, and closed-loop control of automotive power plants including energy storage, regeneration, distribution, predictive management, and monitoring of emissions. A significant step towards achieving dramatic improvements in fuel economy is through the reduction of vehicle weight. CPS research can contribute to weight reduction and cost-efficiency through (1) use of predictable and secure wireless networking protocols to replace heavy wire bundles used for control and information transfer, and (2) reduction in the number of on-board computers and their interconnects through leveraging of mixed-criticality architectures. Integration of energy-harvesting wireless sensors for aircraft in particular has the potential to greatly reduce weight due to long wire runs, resulting in substantial fuel savings.

Improving Quality of Passenger Experience – Personalization of automotive electronics and achieving the "connected" vehicle experience featuring connectivity to external information sources, streaming audio and video, enhanced navigation services, connectivity for traffic congestion management, and dynamic routing are all pathways towards creating the "optimal" vehicle experience. Integration of location-based concierge services and advanced vehicle health management are further contributors. CPS technologies for wireless networking of embedded processors and sensors are at the heart of the "connected vehicle."

For the everyday traveler, the ability to personalize vehicles and provide state-of-the-art comforts and convenience for both business and pleasure would be a major CPS achievement. Automobile as well as airline travelers expect to be able to integrate and/or use their personal electronic devices and obtain entertainment services, e.g., Internet access and live events, from on-board infrastructure. The static road and air vehicle capability configurations of the past are no longer sufficient. The future will demand the ability to customize and reconfigure functionality, performance and personalized features that require post-sale upgradability and change of usage profiles for cyber and physical components. A determining success factor is how performance and vehicle customization and post-sale configurations can be validated and certified. The objective is to continually meet evolving market forces regardless of the rate and reason of evolution.

The design of fuel-efficient, network-enabled operating vehicles will improve the financial bottom line of operators, i.e., auto drivers and airlines, since fuel prices and vehicle maintenance are large costs for operators. Business-owned vehicles will benefit from advances in automation that will enable on-board readable/writable/executable operations and allow different levels of control by local/remote authorized entities. Such vehicles will enable collaborative decision making and new business models and strategies for cost-effective vehicle operation and maintenance, making travel cheaper for commuters.

Situation Awareness in Industrial Aviation Applications – Gate-to-gate situation awareness is a process developed to improve and produce major economies in air travel. Key elements include precision aircraft tracking, seamless situation awareness spanning the entire travel path, and fully

Community Draft

integrated logistics. Major benefits will include travel delay reduction through common airspace awareness, seamless wireless connectivity, integrated health management and prognostics, reduction of maintenance costs and time lost due to repairs, and reliable luggage delivery. CPS research will contribute to safe and secure mobile communication protocols enabling exchange of critical vehicle information, and distributed sensing and reporting of maintenance data.

Grand Challenge for National Security – Persistent Surveillance: Missions requiring 24-hours per day, 7 days per week persistent surveillance drive new generations of intelligent, autonomous systems with increased capabilities to simultaneously track and identify vast numbers of entities.

Existing and future irregular warfare scenarios represent enormous challenges to the safety of military forces as well as peace-keeping forces deployed in regions of national interest. Providing actionable information to military forces requires persistent surveillance on a 24/7 basis. Similar persistent surveillance is required for a number of civil applications ranging from border and harbor patrol, to disaster response (e.g., hurricanes, floods, wildfires), to traffic control. There are also potential law-enforcement applications that add privacy and civil rights constraints. Such capabilities can only be achieved through the development of highly intelligent, agile, unmanned systems that can provide long-duration coverage of wide areas and, consequently, generate enormous quantities of data that must be rapidly reviewed and integrated to generate coherent and actionable information. CPS technology advances are essential for the high level of intelligent autonomy demanded by these systems and for the complex fusion and integration of mega data streams into knowledge. CPS technology is also critical to create required light weight and distributed sensing and is fundamental to ensure the safe and secure vehicle control of highly energy-efficient unmanned systems. From the civilian perspective, an aviation-related grand challenge can be stated as "Design the air traffic control system so that passengers always got to their destination on time, with a plane that was always 90+% full, and with no delays due to weather anyplace in the country."[6]

Grand Challenge for National Security – Responsive Strike: Global strike with near zero timelines and zero collateral damage requires highly precise, very-long-range, and very fast weapon systems.

Strategic interests of America are distributed across a global landscape. Terrorism and armed conflict place these interests at risk on a daily basis. The capability to respond to an incident with near-zero time latency and zero collateral damage represents a highly desirable end-state protecting national interests. Achieving this goal is not possible without highly precise, rapidly deployed and extremely fast weapon systems. Technology challenges for CPS are enormous including the need for high-resolution sensing, intelligent discrimination, agile and secure vehicle control, and networking. A variation of this technology can perhaps be used on the civilian side for critical search and rescue missions in remote locations.

Community Draft

[6] http://www.cds.caltech.edu/~murray/topten/

Current State-of-the-Art in Transportation CPS

The transition from traditional embedded systems to CPS represents a radical change of perspective necessary to address the dramatic and rapidly evolving nature of transportation systems. The notion of what is a vehicle and what is a system is rapidly changing. The vehicle has shifted dramatically to become a highly collaborative computational system that is reliant on sensors and actuators to sense and effect change. Even more dramatic, the notion of a system is changing to include infrastructure as well as vehicles in a system-of-systems, creating a uniquely large scope and context in which to build systems with predictable and provable behaviors. These changes are reflected in four drivers that are increasingly impacting the ability of current techniques and approaches to deliver new vehicles and new transportation systems that cut across all aspects of CPS development throughout the transportation domain. We have reached a tipping point in our ability to deal with system complexity. Current approaches to system certifiability based on process and exhaustive testing do not scale and are not producing sufficient evidence of system dependability[7]. Increasingly, interoperability with other vehicles and transportation system infrastructure is key to achieving system goals. The role of the human should be incorporated into the design dimension as opposed to being a peripheral element.

1. **Complexity** – Ever-increasing system complexity is directly related to the increasing demand for new functionalities. We are reaching a tipping point where the current tools and techniques are unable to deal with the unintended and undesirable emergent behaviors arising from complex runtime interactions between various sub-components, between components and infrastructure, and between components and humans and the environment. With the emergence of greater regulatory requirements to address fuel consumption, emissions, and diagnostics, the average automobile soon will exceed 100 million lines of code.[8] This increase in complexity increases the cost for developing dependable systems in a timely manner. The Hansen Report states that electronics and software account for 30% of the total cost of the automobile, and in order to reduce the impact of additional complexity and to minimize cost increases, the automotive industry is using platform consolidation.[9] CPS must additionally address physical and human issues in their design and development. The inadequacy of current tools and methodologies to deal with CPS complexity, combined with the need to contain cost, has led to a focus on mixed-criticality research and development initiatives.
2. **Certifiable Systems (Software, hardware, and their interaction) –** Although each transportation sector has different regulations with varying certification requirements, the multi-sector community collectively perceives the need for tools and techniques for developing safe and reliable certifiable products. Current techniques and tools have a significant impact on system cost, and are not scalable to tomorrow's system sizes. New certifiability approaches that are both persuasive and affordable, and that can be used to verify the dependability of large-scale dynamics and the adaptability of the systems-of-

[7] Jackson, D., Thomas, M., and Millett, L., Eds. Software For Dependable Systems: Sufficient Evidence? National Research Council. National Academies Press, 2007; books.nap.edu/openbook.php?isbn=0309103940

[8] USCAR CPS Summary presented to this workshop, slide 8.

[9] Hansen Report Volume 19 Issue 9 – 2006 Convergence Panel on Automotive Electronics

Community Draft

systems emerging in the transportation sector, are essential to realizing the anticipated benefits of these systems.

3. **Interoperability –** Transportation interoperability includes being able to build systems with components from different vendors, interacting systems with different generations of vehicles, or upgrading a vehicle to meet new requirements (e.g., lower emission). For the most part, interoperability has been a secondary issue in system development. As we move forward, it is crucial to ensure the interoperability of cyber-physical tools and technologies – not only amongst themselves but also with legacy systems, since new and legacy systems will have to co-exist for the foreseeable future. This would, for example, enable the retrofitting of legacy systems with new cyber-physical technologies.
4. **Human-in-the-Loop –** Active or passive human participants are an important component of every transportation sector. However, the degree to which human behavior has been modeled and incorporated into system design has varied. Many of the envisioned goals, i.e., zero fatalities, can only be met by designing systems with a comprehensive understanding of human behavior under varying situations including emergency or stressful scenarios. With the increasing complexity of interactions of humans with a cyber-physical system, there is a need to develop advanced models for human-machine interactions as well as, in a broader sense, *human-cyber interaction* ,i.e., interactions among people and computers, mediated by the system.

R&D Needs and Challenges

Many existing transportation systems are already equipped, or have been designed, with a limited awareness of and attention to CPS requirements. Advanced control systems bridging the physical world and the cyber systems have been implemented and applied in all transportation sectors. Examples include flight control systems in aircraft, full-speed-range smart adaptive cruise control in vehicle systems, cell phone or GPS based location services, and advanced signal systems for railroad management. Research that directly focuses on CPS is needed to enable the development shift from solutions focusing on either cyber or physical aspects independently, to those focusing on the integrated system. To make this shift towards a holistic cyber-physical approach a reality, fundamental theories, design methods, deployed systems, and education are needed that advance the state-of-the-art. Despite a common shared interest across different transportation sectors, current research thrusts are typically dedicated to individual sectors, or even individual domains within each sector.

While all CPS domains (e.g., transportation, health care, manufacturing) share similarities in the challenges they face, there are several unique aspects of transportation CPS that trigger a specific and distinct research agenda, though the benefits in many cases are not limited to transportation CPS. The bandwidth of functionality in vehicles demands a diverse set of complexities for controls including mixed hybrid control systems. The customization and variability of system architecture weaves a level of variant complexity rarely observed outside of the industry. The ultra-competitive global landscape mandates ever-evolving requirements for enhanced capabilities, resulting in the need to rapidly adapt systems. The changing landscape requiring the hardening of systems to external exposure is leading to a more pervasive set of cross-cutting requirements for system level qualities such as safety and security. In addition to the competition in the pursuit of vehicle consumers, the global competition at the component level is equally fierce, resulting in a diverse and constantly evolving set of component suppliers that must provide products to be integrated into the whole.

Community Draft

The areas of CPS research that primarily impact the transportation sector fall into the following five broad areas:

- **Theoretical Foundations** – To capture knowledge regarding new system development and certification techniques, and to ease the design burden faced by system engineers, new abstractions (and the foundational science that underpins them) must be identified for cyber-physical transportation systems. While some aspects of CPS design presently have such foundations (e.g., controls engineering), these concepts are not pervasive. Furthermore, there is no unifying foundation to enable reasoning across cyber-physical dimensions that are required to compose and analyze these systems.
- **Model-Based Analysis** – The evolution of our analytical techniques must accelerate considerably. While current tools are designed to help us reason about single dimensions, vehicles operate in a multi-dimensional cyber-physical state space. A dramatic increase in the ability to perform true cyber-physical co-design – where the physics of surface friction, moments of inertia, and computer hardware and software behavior can be simultaneously observed – is critical to advance both how systems are engineered and the kinds and the quality of systems that are deployed.
- **Adaptation and Self-Healing** – The future's complex vehicles and interoperating infrastructures must be able to adapt rapidly to anomalies in the environment and to embrace the evolution of technologies while still providing critical assertions of performance and other constraints. Whether it be degraded performance of sensors, the failure of another vehicle, or an infrastructure malfunction (e.g., a traffic light failure), these systems must adapt to whatever situations they encounter and virtually heal to yield the best possible system performance under those conditions. Advances need to be made in how these systems are designed and implemented, and in the supporting infrastructures, to make this capability viable.
- **The Human Role –** Understanding and developing systems that are intuitive and that integrate with human behaviors in high-stress environments are essential in communicating critical information and supporting decision processes of drivers, pilots, and other vehicle and infrastructure operators. While we have significant data and experience with the systems of today, there is a need for a more formal understanding of these issues, as the growth of assistive and autonomous behaviors in vehicles continues to expand.
- **Engineering Practice** – Significant growth is needed in the integration of science into engineering practice that can be executed by the generally trained engineer. Even essential technologies such as model checking (for proof of correctness ranging from software components to entire systems), have seen only limited adoption due to lack of accessibility to the generally trained engineer. Promising technologies such as this must be integrated with existing practices to expand the professional toolkit of those facing the challenges of designing CPS.

Just as there are a number of common themes that impact our ability to develop CPS systems using current approaches, there are also a number of common challenges that cut across all of the research thrusts needed to spawn new approaches. All of the challenges below must be taken into account in all of the proposed research directions if the proposed research is to have an impact on

Community Draft

the real vehicles and transportation systems that will respond to the grand challenges in transportation CPS.

- **Integration –** Industrial demand is focused on system composability. The objective is to create the "plug and play" concept at multiple levels. This notion of multi-level composability includes software components, electronic components, networks, physical composability, and V2X technologies (e.g., vehicle-to-vehicle or vehicle-to-infrastructure) technologies. The integration challenge includes the achievement and maintenance of cyber-physical properties of the system in light of both obvious and subtle disturbances in both the cyber and the physical environments**.**
- **Virtual Development Enterprise –** We are in an environment where systems are increasingly composed of components provided by multiple vendors as part of a virtual enterprise. While open architectures allow the integration of independently developed and selected software (and hardware) components based on their interfaces, issues arise in integration, verification and validation, and certification. The interface boundary is there to support both abstraction and protection of intellectual property. This issue is already a factor in relationships between suppliers and integrators today when system faults are being diagnosed and questions arise as to the contribution of the internal workings of components, or the interactions between components from different suppliers. Similar issues can also arise when interacting with or upgrading legacy systems. Any transitionable solution to the CPS R&D grand challenges must support a virtual development enterprise where system-level properties must be achieved with imperfect knowledge of system components.
- **Systems Engineering** – With the increasing complexity and cost of system development, there is resurgence in system engineering. It is important to ensure that the CPS perspective is taken into account as we move forward. Managing and exploiting the interactions between the cyber and physical components distinguishes CPS from traditional systems. It may be extremely difficult to predict all the interaction issues early in the development process; these may become evident only later during integration of various cyber-physical components. System engineering should evolve to include the CPS perspective by acknowledging a greater interdependence with architectural definition that arises because of CPS. Today, systems engineering includes decomposition of requirements to hardware, software, and operator actions, and identification of interfaces and architecture. With the emergence of V2X leading to ever-larger systems of systems, a more scalable method for partitioning system design and analysis is needed. Design decomposition in CPS becomes ever-more intricate. For example, physical properties (e.g., distance). can both couple and isolate components. In CPS, innovative decomposition strategies are required.
- **Verification, Validation and Certification** - With increases in the complexity and the criticality of electronics in the control of the physical components and vehicle behavior, the importance of effective methods in validation and certification is accentuated. The current practices of "testing in quality" or "post-facto certification" prove to be methods that sometimes are ineffective. New holistic methods initiated from the beginning of the design process must be created to ease (1) the burden of proof, and (2) the execution effort required to meet the scale and quality demanded in the future, and to focus V&V

Community Draft

activities on the most relevant areas of the, necessarily vast, cyber-physical system state space to minimize costs.

- **Design Representation, Tool Support, and Tool Chains** – With the scale and dimensions of designs increasing rapidly, the need to integrate knowledge utilizing multiple semantics and multiple syntaxes for analysis is critical. The cross cyber-physical analysis is critical to assess the performance and viability of system designs accurately and safely, independent of the physical properties of vehicles. This challenge is exacerbated by the number of cross-cutting concerns, such as safety, which have implications in multiple representations. These new approaches and representations must manifest themselves in tools and tool chains so that practitioners can apply the technologies to the design and development of real vehicles and transportation systems. Tools embody advances in CPS knowledge in ways are usable by practitioners. As the complexity of systems increases so also is the complexity of tools and tool chains. There is an urgent need for tools that are usable by practicing engineers.
- **Human Factors** – The most uncertain, and also often the most important, factor in transportation vehicle operation is that of "human interaction." The ways in which vehicles can interact can be either beneficial or detrimental to focusing the operator on the most important tasks. More confounding is that these behaviors adapt and evolve over time such that interactions must change and the initial learned behaviors may reduce effectiveness over time. Focused research on human factors in vehicle design is imperative to provide the strategic and useful aids required in the future.

Commonalities/Synergies Across Transportation Sectors

The transportation sector, as the fabric for movement of people, animals, grains, raw materials and finished goods, is a cornerstone of the U.S. economy. It is also a key to U.S. national security as (a) part of the logistics infrastructure of our armed forces, (b) the provider of strategic and tactical mobility, and (c) the supplier of many of the advanced weapon systems used by our war fighters. Unmanned air vehicles and other advanced transportation technologies developed for military use also have important civilian applications, such as the recent use of Predator drones for surveillance of rivers and levees[10] in times of danger caused by flooding.

The transportation domain includes a number of sectors: automotive, aerospace, and rail. Each sector is a manifestation of a cyber-physical system, and includes both vehicles (cars, trucks, airplanes, and trains) and infrastructural components (highways, airports and railroad tracks). There are some obvious differences among these transportation sectors. Vehicle speeds vary significantly from sector to sector. There are also substantial differences in the number, the size and the cost of both vehicles and infrastructure components among these sectors. Meta-characteristics can also be quite different. For example, the relationship between highways and automobiles is rather passive, while the next generation air traffic control (e.g., NextGen) is very tightly interwoven with passenger aircraft takeoffs, flight paths and landings. Operator and vehicle qualification varies widely, with the automobile sector being the least restrictive, and

[10] "Red River crests below forecast in Fargo", The Associated Press, 03/29/2009. http://www.stltoday.com/stltoday/news/stories.nsf/nation/story/DBAB495AEE626C8D86257588000DC99B?OpenDocument

Community Draft

aerospace being the most restrictive. The ownership of the infrastructure differs as well, which has an impact on the rate of change. Directly related to vehicle qualification is the rate of technology change and refresh, and the persistence of legacy systems. The rail sector in particular is severely constrained by legacy infrastructure and technologies.

The transportation industry includes a number of sectors: automotive, aerospace, and rail. Each sector is a manifestation of a cyber-physical system, and includes both vehicles (cars, trucks, airplanes, and trains) and infrastructural components (highways, airports and railroad tracks). There are some obvious differences among these transportation sectors. Vehicle speeds vary significantly from sector to sector. There are also substantial differences in the number, size, and cost of both vehicles and infrastructure components among these sectors. Meta-characteristics can also be quite different. For example, the relationship between highways and automobiles is rather passive, while the next-generation air traffic control (e.g., NextGen) is very tightly interwoven with passenger aircraft departures, flight paths, and landings. Operator and vehicle qualification varies widely, with the automobile sector being the least restrictive, and aerospace being the most restrictive. The ownership of the infrastructure differs as well, which has an impact on the rate of change. Directly related to vehicle qualification is the rate of technology change and refresh, and the persistence of legacy systems. The rail sector in particular is severely constrained by legacy infrastructure and technologies.

While differences have a serious impact on the nature of the cyber-physical systems that arise in each transportation sector, the commonalities among these sectors are even more substantial and significant. This is particularly true when looking at the research challenges facing the sectors in implementing transportation cyber-physical systems. All these sectors share a common vision of future systems as well as core technology areas including mixed criticality, autonomy and control, platforms and infrastructure, infotronics and infotainment, verification and validation, model-based development, and tool chains.

In terms of a common core need, all transportation sectors are increasingly technology-driven but must demonstrably satisfy safety, reliability and security requirements. All sectors are moving to an infrastructure characterized and enabled by distributed collaboration. For example, cooperative cruise control will soon be a reality. The next generation air traffic management system is based in part on distributing the control currently exercised by air traffic control throughout the airspace. There is also a growing trend towards autonomy in all the transportation sectors. The automotive sector is already witnessing the introduction of automated parking systems, and the recent DARPA Grand Challenges and Urban Challenge have demonstrated that the reality of autonomous automobiles is drawing ever closer. The role of UAVs in military aerospace is growing exponentially, as seen in Iraq and Afghanistan, while their use in civil aerospace has already begun, as is evidenced by the use of a NASA UAV to provide data on California forest fires, and the use of DHS UAVs to patrol our northern and southern borders.

All transportation sectors are also steadily progressing towards a future of long-lived platforms and infrastructure. It is very likely that today's aircraft may be flying for an additional 40 years. The airports, highways and railroad tracks that we have today will be with us at least as long. This means that all the transportation sectors will have to (1) deal with issues of legacy integration and migration, (2) develop CPS that deal with widely varying capabilities, and (3) support maintenance and upgrade needs while maintaining availability. Another shared objective across all of the transportation sectors is the exploitation of advances in cyber-physical systems to enhance safety. For example, the automotive sector is working towards a grand challenge goal

Community Draft

of *zero fatalities*. Given that there are more than 40,000 automobile-related fatalities annually in the U.S. (and more than 1 million across the globe), even an asymptotic trend towards this goal will have major benefits for society and the economy.

These common goals of transportation sectors can only be realized if there are significant advances in the science of cyber-physical systems, as well as in the practice of developing such systems. Accordingly, each of the sectors also envisions a future with development environments that provide significant support for engineering such systems. Realizing this shared vision requires advances in a number of technology areas in ways that will benefit all transportation sectors. While the motivation and respective details may differ, all transportation sectors require improvements in the techniques for dealing with systems containing elements of different criticality levels, i.e., elements that have different levels of impact on system safety and consequence of failure. Whether this is a brake-by-wire subsystem sharing a data bus with navigation data in an automobile, or flight-control software sharing a computer with waypoint navigation in an aircraft, techniques for providing the required isolation in the face of the shared resources that will arise from the use of emerging multi-core processors will need to go beyond the traditional federated systems approaches with "logical" separation replacing physical separation.

New approaches to deal with autonomy, partial autonomy, and mixed human-autonomous system interactions will be required for both the vehicular and broader transportation infrastructures. First, a better framework for defining and reasoning about autonomy is required. Another key challenge is to formulate approaches that accommodate multiple sources of initiative without being too conservative. Humans are good at this type of behavior, but current autonomy approaches based on reachable states may not be practical. For example, cars passing in adjacent lanes are likely to have already compromised reachable state-based safety margins.

Another area of commonality between the sectors is verification and validation (V&V). Current techniques are challenged by the scale of emerging systems, by the increasing demand for advanced capabilities such as autonomy and adaptive control in safety-critical systems, and by the hybrid nature of CPS that combine both discrete and continuous aspects. New V&V approaches are needed to support compositional system development, and they must enable V&V of systems in the context of their environment. This echoes findings of the *Software for Dependable Systems: Sufficient Evidence?* National Academies study[11]. All transportation sectors also need V&V techniques that provide these needed capabilities in a form that scales to industrial-sized CPS with tens of millions of lines of code and tens to hundreds of nodes per platform, and that are usable by the engineers developing these systems. Additional scalability challenges arise when attempting to verify and validate sophisticated cyber-physical systems-of-systems such as the interactions of vehicles and infrastructure in NextGen air traffic management, cooperative cruise control, and smart highways. An additional consideration, though one that varies across sectors, is the role of new certification V&V techniques. As systems become increasingly complex with an increasing number of components and interactions, the probability that undesired emergent behaviors could arise during runtime increases. Techniques for incorporating V&V runtimes could result in more dependable cyber-

[11] Jackson, D., Thomas, M., and Millett, L., Eds. *Software For Dependable Systems: Sufficient Evidence?* National Research Council. National Academies Press, 2007; books.nap.edu/openbook.php?isbn=0309103940.

Community Draft

physical systems, which would benefit all the transportation sectors. Other considerations include *Infotronics and infotainment;* these are mixed-criticality systems because they not only enable control, coordination, communication, and navigation of transportation, which impacts safety and raises new security issues, but they also provide entertainment which increases the systems robustness.

All transportation sectors are migrating towards the use of model-based development that relies on sophisticated tool chains to automate the development process. Most existing model-based development approaches focus on specific aspects, such as control models or component connection models. New approaches to deal with multiple system views are needed. A very real issue for CPS development environments is how to debug a CPS? How to set breakpoints in a cyber-physical system? This is challenging even in a purely simulated environment. These challenges are multiplied as system development progresses to "hardware in the loop" and target platform testing (e.g., flight test and vehicle test tracks). Related challenges include feedback and feed-forward techniques between the various environments to improve the fidelity of models, and collecting, managing, and mining data to support such CPS environment modeling. The data-sets involved can be quite massive. Additionally, managing and tracking various artifacts involved in the design, development, and testing of CPS itself is a major common challenge.

Addressing the need to keep *humans in the loop* is also common and is *often critical* across all transportation sectors despite the trend towards increasing autonomy. Humans of varying capabilities are involved in the design, development, and operation of a CPS. In order to verify and validate the entire system, humans must be included as part of the system to strive to attain goals such as zero accident fatalities. Thus, how to abstract and represent human behavior demands focused attention for all the sectors. Bridging the gap between natural language representation and formal language framework for CPS design and development would help decrease the cost for all the sectors.

These common challenges faced by the transportation sectors are not necessarily unique to transportation CPS. Similar challenges can also be found in the domains of medical devices, smart structures, and energy generation/distribution. The medical device domain is faced with many of the same system-of-systems V&V and certification challenges. For example, the CPS formed by the network of medical devices collaborating to provide care for a patient is in many ways similar to a platoon of autonomous cars navigating the same section of a smart highway, or the aircraft collaborating via the NextGen air traffic control (currently under development) to ensure separation transiting an airspace sector.

In summary, there are significant commonalities across the various transportation sectors, and other safety-critical CPS sectors. However, a research agenda must be developed to take advantage of these commonalities with the goal of making a significant impact on U.S. competitiveness and quality of life. The creation of integrated challenges that provide focus and context for emerging CPS research will greatly facilitate the realization of common goals. One such cross-sector challenge that could focus attention on common issues is global multi-modal transportation optimization[12]. Such a multi-sector challenge would include issues such as door-

[12] Multi-modal transportation occurs when a person must utilize different modes of transportation (e.g. some combination of car, taxi, train, plane, and bus) to go from point A to point B.

Community Draft

to-door optimal routing, environmental modeling including weather effects, and issues of autonomy and traffic, offering abundant scope for concrete problems and experiments supporting all of the common technology areas.

Education

A strong, vibrant, and knowledgeable workforce is crucial for the short-term energizing and long-term health of the socially and economically crucial area of transportation CPS. These high-quality technical personnel must be both comfortable and productive in working in a multi-disciplinary realm with complex, safety-critical, and life-critical requirements. The current educational framework is not sufficient to meet these needs. For instance, many experts in the cyber-domain today focus only on information management and security in the virtual world. They may obtain little or no exposure to the principles of engineering and physical dynamics. In fact, a surprising number of computer science (CS) degrees do not even require a course on freshman physics. Conversely, most non-computer science students often learn programming as a craft for occasional use. They do not learn the general principles of expressing and satisfying different cyber-physical attributes (such as timeliness, safety, and reliability) that must be applied to CPS. Fortunately, embedded system designers are at least aware of resource constraints in the computing domain (like memory, energy, and processing power limitations), and serve as a starting point for providing CPS development skills.

The following dimensions represent critical facets that must be supported by educational curricula in order to train tomorrow's CPS experts.

- **Training in Problem Solving** – The problems in the cyber-physical domain are complex and often difficult to predict. These characteristics demand significant intellectual capability to decipher and solve problems across multiple different domains of expertise. The capability to resolve systemic problems, rather than mask their symptoms, is a difficult one that must be fostered to enable the critical and creative skills of problem solving. Our academic environment must foster these skills to engage and prepare the future workforce.
- **Mixed-Discipline Development** – The challenges of cyber-physical systems are no longer the purview of single traditional "stove-piped" disciplines such as electrical engineering, mechanical engineering, or computer science and computer engineering alone, but are rather an integrated application of the collective knowledge and techniques of multiple disciplines. This dramatically changes the background and specialization in the fields. A new focus is needed in all of these disciplines to introduce the cross-disciplinary nature of cyber-physical problems and the challenges these interactions evoke. Additionally, there needs to be a growth in the capability and compatibility of these trained engineers working in a cross-discipline team where collaboration and respect are both developed and well founded in project execution.
- **Cyber-Physical Science and Engineering as an Academic Path** – Much as computer engineering started as part of the root discipline of electrical engineering, transportation CPS must distinguish itself as a unique engineering discipline requiring specific skills and techniques. To this end, our current embedded systems programs must be extended to meet the larger challenge of cyber-physical systems. Additionally. the community should continue to be educated through workshops and conferences dedicated to the topic. The

Community Draft

CPS information should be incorporated into the K-12 curriculum as well as in bachelor's, master's, and doctoral programs.

- **Re-training** – Until the recent economic crisis, the automotive sector represented over 5% of the U.S. private-sector GDP and employed one out of every seven Americans in the work force, while the American aviation sector is the world's largest exporter of commercial airplanes, again with the largest aviation workforce in the U.S. The community feels that more active approaches, such as supporting the creation of professional master's programs, are needed.
- **Motivation** – To implement the proposed changes successfully, we must excite the student community. More needs to be done to provide exciting and accessible challenge problems especially at the undergraduate level. The growth of the talent in the desktop and Web community is driven directly from that accessibility and intrigue, and this must be emulated in the CPS innovations. The development of "grand challenges" in addition to well-defined projects and representative challenges at the undergraduate and master's level will be critical motivators.

Educational Curriculum Requirements

To address the needs of transportation CPS, engineering and computer science programs must strongly interact and develop integrated curricula that combine the relevant core aspects of each discipline – for example, breaking down barriers between traditional control theory and embedded system computing courses so that elements of control theory are taught side-by-side with topics such as software performance analysis and real-time resource scheduling. Computer science students must learn substantial amounts of engineering and computational physics, e.g., energy consumption, as well as experience the realistic and imperfect nature of physical systems (sensor noise, variations in vehicle-to-vehicle performance characteristics due to normal manufacturing variability, etc.). Software must be pedagogically considered as something that's engineered rather than produced by artisans. Similarly, the engineering curriculum should cover important computer science topics, such as algorithms and search and operating systems, as well as skills such as software development environments. System-level concepts such as composability and scalability must be incorporated into engineering curricula to develop students' understanding of how transportation systems, such as in aviation, differ from other software-intensive application disciplines. Control engineers should be able to view control software as a necessary part of the control system, and not as a toolbox for simulation. Furthermore, the engineering curriculum across the board must expand in order to include topics on formal specification, verification, composition, interfaces, and hybrid systems to suitably prepare the future workforce.

Academic projects that adopt a lifecycle view of systems, i.e., system architecture, specifying requirements, identifying testable requirements, and using model-based design and formal methods to verify software and protocols, must be taught. By training students in building production-quality code to be within schedule and budget, the software engineering curriculum can be infused with appropriate systems engineering content, and similarly, software engineering can be incorporated into curricula of disciplines such as electrical and aerospace engineering that will be writing code for future systems.

It also would be beneficial to introduce senior design projects in the undergraduate curriculum that emphasize cyber-physical themes. There is an increasing need for engineers who are trained in project-based "capstone" design courses and who have the experience with model-based

Community Draft

design techniques. Education programs should utilize state-of-the-art tools and techniques, such as model-based software engineering for embedded software.

It would likewise be desirable to incorporate cyber-physical concepts into *high school education*, to ensure that students consider the CPS discipline when choosing higher education. Ultimately, a few key ideas in cyber-physical systems should also be included in *K-12 education*. Many university faculty report that incoming students are ill-prepared in basic computer science and engineering principles. Updating the curriculum to include CPS concepts could be an important part of a broad refresh in the K-12 computing curriculum. Introducing some basic control concepts regarding computing could also bridge the gap between traditional mathematics and the mathematics of computer science. Competitions such as the For Inspiration and Recognition of Science and Technology (FIRST) robotics challenge[13] have proved to be successful venues for building interest in engineering and teaching basic principles; such competitions are also ideal places to experiment with CPS concepts for K-12 education.

Currently, practicing engineers and other workforce members in the transportation sectors must also receive adequate training to transition to a CPS view of next-generation transportation systems. Such training is necessary to maintain the human-in-the-loop CPS skills, while retaining and expanding the industry workforce. For example, software developers and testers need to move from the traditional testing-based methodologies to new formal-analysis methodologies. Policy makers must be educated with an understanding of the new safety, reliability, and security challenges and requirements involved with the emerging cyber-physical interactions.

University-Industry Interactions

CPS education curriculum should also aspire to train students through realistic design experiences, ideally through partnership with industry. It is important that industry be involved in the transportation CPS education process, since there is not enough domain expertise currently in academia. Much can be gained by providing challenging problems from industry to the academic research community. Past examples have included the DARPA and AFRL MoBIES, SEC, PCES, and CerTA FCS programs, as well as the NSF-ITR on Embedded Systems experience. Such hands-on projects in the nation's educational programs would enable budding engineers and scientists to experience the joy of creating solutions. University-industry collaborations can offer internship and mentoring opportunities, recognition of personal and team accomplishments, and a focus on current challenges and environments. These activities must balance an attractive lifestyle and compensation that properly recognizes the practical value of contributions to transportation CPS. Industrial involvement in education would also offer incentives for the broader community, such as creation of an open-source model-based design tool chain, that could be used both in the education of CPS engineers and in industry. Meanwhile, getting involved in the CPS community could provide the transportation industry with opportunities to regenerate and rebuild its workforce. The industry is currently faced with a leading challenge of regenerating and maintaining its professional workforce.[14] The current workforce is aging, and there are too few young people opting for careers in engineering and

[13] Founded by Dean Kamen to make science and technology fun through teamwork and friendly competition. Sponsored by Boston Scientific, Baxter International Inc., and Johnson & Johnson

[14] Inside Aerospace Report and Recommendations, May 13-14, 2008, Arlington, VA. http://www.aiaa.org/pdf/public/Inside_Aero08_Report_and_Recommendations.pdf

Community Draft

science, creating a potential shortfall of engineers and scientists in transportation. Due to factors such as lack of inspirational activities and role models, aerospace has not been able to capitalize on two-thirds of the nation's future workforce, specifically women and minorities.

Government-University Interactions

In order to incorporate and fund transportation CPS, investments should be made in educational research projects via the agencies and the scope of national fellowship programs such as the NSF Graduate Fellowship and NDSEG, should be broadened.

Roadmap and Milestones

The participants in the HCTCPS Workshop characterized the state of the art and future directions in transportation CPS. There were two significant workshop outcomes: (1) The recognition that there is a *demand* for new techniques, approaches, and systems to meet the need for safe, efficient transportation; and (2) Though there are differences in detail between the needs and challenges of the individual transportation domains, there is a large degree of commonality that can be exploited to craft a research agenda whose fruits will benefit the entire transportation domain. To help define such an agenda, each of the five breakout sessions (the reports from the breakout sessions follow) was charged with developing a roadmap for its topic area. This section presents a higher-level roadmap for the transportation domain as a whole, with 5-, 10-, and 15-year milestones, both in terms of the identified research areas and the transition of research into the U.S. transportation system, aimed at meeting the grand challenges posed at the beginning of this workshop report. Following is the workshop's high-level roadmap for transportation CPS:

5-Year Plan

- Foundations - Developing a theory of probabilistic hybrid models
- Analysis - Building runtime architectural supports to guarantee components continue to meet specifications
- Adaptation - Developing a vehicle-to-infrastructure (V2I) to support controls and communications, for example to allow V2I exchange of embedded software updates and infotainment data, certifiable for non-safety critical functions
- Human Role - Conservative management of manned and unmanned aerial systems in airspace
- Engineering Practice - Infuse BS and MS curricula in Engineering and Computer Science with CPS principles
- Transition and Deployment
    - Creating a reference architecture for networked automotive CPS with basic sensing, control, and actuation based on single and single or multi-hop neighbor interactions
    - Demonstrate improvements in traffic management, fuel-efficiency, and noise reduction of vehicles
    - Open Virtual Networked Vehicle Test-bed for transportation CPS with a deployed network of 50 real vehicles and 1,000,000 virtual vehicles supporting V2V communication

10-Year Plan

- Foundations
    - Real-time theory for spatio-temporal models in CPS

Community Draft

    - Theories of composable CPS
- Analysis
    - Verification of equivalence of components to enable “plug-n-play”
    - Safety-critical certified V2I software and systems
- Adaptation = V2X software and systems for exchange of information such as traffic, weather, infotainment data, certifiable for non-safety critical functions
- Human Role - Efficient and safe management of manned and unmanned aerial systems in shared airspace
- Engineering Practice
    - Standards and specification for infrastructures to support CPS
    - Model-driven system generation (rather than hand coding)
    - Establishment of CPS professional Masters Degree programs and CPS professional certification
- Transition and Deployment
    - Transition to more efficient and available fuels
    - Reduction of point-to-point travel delays
    - Minimization of point-to-point travel delays
    - Enhanced test-bed for transportation CPS with additional real vehicles and live and virtual infrastructure with supporting V2V and V2I communication, and capabilities for testing multi-modal interactions.

- **15-Year Plan**
- Foundations - New paradigms for affordably and scalably ensuring safety and security of ultra large scale dynamic CPS.
- Analysis - Verification of adaptive, probabilistic, and networked systems
- Adaptation - Dynamic end-to-end performance optimization for platform adaptation across multiple transportation vehicles, certifiable for safety-critical functions.
- Human Role - Self-aware systems that can disable aspects of their operation based on self evaluation of inability to achieve mission objectives
- Engineering Practice
    - Zero prototype production releases with full virtual validation
    - Tools for checking non-functional properties of models
    - High-assurance certified V2X software and systems
- Transition and Deployment - Minimal environmental footprint of transportation

Community Draft

Verification and Validation

Participants: Ashish Tiwari, Bill Milam, Beth Latronico, Ron Garcia, Andre Platzer, Matt Behr, Frank Vehid, Pam Binns, Natasha Neogi, Steve Miller, Todd Belote, Basil Krikeles, John Baras, Ben Watson, David Garlan, Bruce Krogh

Introduction

Verification and validation are core interests for all three of the industry groups. Vehicles in all three domains are canonical cyber-physical systems. They are now approaching the sophistication of the most complex systems ever designed by humans, but added time and cost constraints make their development a significant challenge. In addition, the challenge of including the physical behaviors in the automated V&V methods is unparalleled. Automated technologies for V&V are now indispensable for developing safe and reliable transportation.

Before continuing, we note that there are three major research themes under V&V:

1. **Specification** – How do we describe desirable behaviors for CPS? Is temporal logic enough to do this? What more do we need? How do we go about reasoning an approach to find what constitutes a suitable specification method?
2. **Design** – Given a specification for a component or subsystem, what is the tool set for implementing a design against that specification?
3. **Verification** – How do we verify that the entire system of subsystems satisfies the overall specification?

Where Are We Now?

Several aspects of verification and validation are related to these questions. One is *functional correctness*: Does the system perform the intended task correctly? Another is *extra-functional:* Will the tasks be executed as scheduled even in a worst-case execution time scenario? Will the processor utilization exceed a defined threshold? Specific limitations are:

- Physics is a big unknown. A key area limiting our ability to deal with cyber-physical systems is our ability to model the physics of the system. A great deal of work has been done in modeling discrete systems that lends itself to formal V&V tools such as model checkers.
- Incomplete and evolving requirements and specifications. One of the challenges of CPS is that often the requirements are incomplete at the start of the project. For example, we might state that a requirement is to keep idle speed within 2% of the set-point, no matter what the disturbance. However, to accomplish this requirement, other requirements emerge such as sensing impending transient loads in order to use a feed forward mechanism in anticipation of the change in load..
- What is good enough? An interesting area of conflict between computer scientists and control engineers concerns the concept of fully proving correctness, versus accepting what is good enough according to the engineers' judgment.

What Are The Challenges?

- CPS models – Formal models of the physical components of a cyber-physical system are difficult to obtain. Even when formalized, the resulting dynamics are nonlinear, non-deterministic, and stochastic. Moreover, CPS work in a context/environment that is difficult to predict and model. CPS are high-dimensional systems that span multiple time

Community Draft

scales. These systems are also often adaptive and intelligent; for example, they can dynamically reconfigure. Cyber-physical systems are multi-entity systems that also include human interaction. These aspects of CPS contribute to the complexity of the models that represent them.

- System and component properties – Formalizing system and component properties is also challenging because CPS exhibit desired and undesired emergent behaviors. Properties of interest are often based on performance metrics and resource usage. Security is an important requirement; CPS are mixed-criticality systems. Defining full ”correctness” is difficult, and assessing coverage is challenging.
- Scalable V&V processes – The main challenge is coming up with scalable verification approaches that can handle these new complexities in the system model and the property/specifications.
- Education – Engineering and computer science curricula should train students in the following: calculus-based mathematics, discrete mathematics, thinking about abstractions, and more broadly mastering complexity.

Specific V&V Research Challenges

Research challenges in verification and validation of transportation CPS include the following:

- Make verification technology scalable to large CPS - There is a wide gap between what has been demonstrated for limited sample problems and the ability of current technology to verify and validate large, real-world problems.
- Design formal performance models – What is a suitable modeling language? A key aspect of verifying the deployment of control software is ensuring that it will be able to run within the specified scheduling regime on the platform. Capturing key performance metrics from the control design is one aspect of this challenge.
- Establish quantitative metrics to prove properties – How do we capture specific requirements metrics such as response time and define failure management behaviors.
- Provide probabilistic guarantees from stochastic models of CPS – Uncertainty in the physical world can be described using stochastic models. Verification techniques are needed to generate probabilistic conclusions regarding such models. In general, the level of uncertainty associated with CPS requires probabilistic models to capture key properties. Unfortunately, tools for reasoning quantitatively about probabilistic systems are only now emerging, and have many limitations.
- Develop appropriate new modeling languages and tools – For example, what are the appropriate enhancements to Simulink/Stateflow that will enable CPS design, testing, and development? Currently, such languages represent the state of the art for a number of application domains, such as automotive and aviation. Simulink/Stateflow has become a *de facto* standard for control system design and implementation through code generation. However, this leaves a gap with the formal languages used for verification, which are more rigorously defined.
- Develop verification systems for human-computer interaction modeling – Modeling of human operators for V&V is an area that is still largely underdeveloped. The focus has been on modeling the physical behavior and control system, but little has been done to

Community Draft

address the human interaction with the system. A car responds to the driver's commands via the steering wheel, throttle, and brakes. On the other hand, to avoid a rollover the vehicle itself attempts to assess the actual vehicle status, infer the driver's intentions, and determine if action is needed to prevent loss of control.

- Identify the dynamic or runtime verification technologies for CPS – When a complex system cannot be statically verified, it may still be possible to improve reliability by monitoring the system as it is running. Can such techniques (probabilistically) guarantee lifetime safety?
- Improve testing-based analyses with V&V technology – The use of counter examples from formal analysis tools can be used as test vectors in functional testing. There are tools on the market that generate test vectors for code or model coverage. However, test vectors that demonstrate correct behavior or the lack of undesired behavior are a challenge.
- Write a compositional algebra to formalize the different levels of abstraction of CPS – One of the potential avenues to be pursued in making industrial-scale V&V problems tractable is the use of multiple variants of a model that represent different abstraction representations. How do we ensure that they are consistent with the original model and can we recompose the abstractions to capture newly added attributes?
- Identify architectures that can aid the design of verifiably safe and secure CPS – Experience has shown that by limiting the architectures so that we can constrain system designs to specific well-understood patterns, we can produce systems that can be verified more easily than systems with arbitrary architectures. However, it is unclear what kinds of architectural styles and patterns are appropriate to CPS. Architectural modeling approaches do not currently support the integration of both cyber and physical elements and properties. What kinds of architectural models are needed to do this?
- Design verification to be compositional – Can domain-specific modeling languages (DMSL), relied upon to promote the use of model-based development, also be composed to check system properties and behaviors? For example, we want to compose a model of the platform behavior with a model of the control algorithm and a model of the physical plant.
- Generate interfaces and check interface compatibility – A system consisting of several interacting components can be analyzed compositionally if individual components specify the assumptions they make concerning their inputs and the guarantees they generate regarding their outputs. How do you specify and generate such interfaces for CPS components and check interface compatibility? This becomes a particularly difficult problem when issues such as timing and other extra-functional properties are taken into account.

Strategies and Roadmap

In order to promote research activity that addresses our concerns, a common open experimental platform or several platforms must be created. Some examples exist of industry donating vehicles and components to universities to develop and validate their work. However, the ad hoc nature of this approach makes collaboration difficult. Creating experimental platforms also allows for comparison and evaluation of various methods against a common example.

Community Draft

5-Year Milestones

- Connected vehicles –V2V and V2I
- Smart roadways sense vehicles, provide information to vehicles
- Accident mitigation
  - hybrid safety verification
- Maneuver assist, e.g., parking assist
- Continuous verification of system integrity (V&V inside)

10-Year Milestones

- Active accident avoidance
  - Cross, oncoming, grade crossing
- Optimize navigation for time, traffic, weather, and energy
  - Instrumented/smart roads provides necessary information

15-20-Year Milestones

- More efficient utilization of infrastructure
- Higher traffic loads, dynamic response to conditions
- Fully/partially autonomous vehicles and supporting infrastructure
- Cooperative vehicle networks
- Near-zero emissions
- Near-zero accidents

Community Draft

Mixed Criticality

Participants: Jim Ritcey, Peter Dibble, Mo-Yuen Chow, Martine Fritzsche, Chris Walter, Kurt Doppelbauer, Peter Stanfill, Craig Treece, Jim Barhorst, Wayne Wolf, Sandeep Gupta, Patrick Goertzen, Alex Doboli, Douglas Stuart, Mark Swick, Jim Paunicka, Patrick Stokes, Ashish Agarwal

Introduction

A mixed-criticality system is an integrated suite of hardware, operating system and middleware services, and application software that supports the execution of safety-critical, mission-critical, and non-critical software within a single, secure compute platform. Implicit in this definition is the concept that the lower-criticality software does not, in any manner, negatively affect the execution of higher-criticality software. Safety-critical systems have traditionally been separated from less critical systems. This is designed and implemented through various methods including physical hardware separation. This is also desirable from a certification standpoint as it robustly separates the higher critical processes from lower criticality ones, thus preventing a lower-criticality function from adversely affecting a higher-criticality function, leading to unpredictable system behavior.

The next generation of vehicles across the transportation domain requires greater software complexity and higher level functions. This increased complexity challenges the long-held definitions of safety critical and mission critical. For example, functions and data that historically were considered mission critical are now safety critical. In manned vehicles, the human operator/pilot is the ultimate decision-maker determining what less-than-safety critical data to accept. In autonomous systems, the embedded software often has to decide which data to trust often co-located on the same processor. Additionally, vehicles no longer operate in isolation, rather there is collaborative on-and-off vehicle to accomplish goals such as retasking, distributed sensing, collaborative cruise control, etc. All of this has to be accomplished in the context of long lived vehicles and even longer lived infrastructure, e.g., rail infrastructure lasts 50+ years. Infrastructure and vehicles must support flexible upgrades of varying criticality. Consequently, the status-quo of process isolation is no longer adequate or feasible.

There are common advantages to employ mixed-critical systems throughout the transportation domain, in areas ranging from UAVs to passenger vehicles. We have been pushing the limits of increases in capability and vehicular performance to out-do adversaries or competitors; employing ever faster time constants; and face potential loss of vehicle control by missing a few processing deadlines. In the realm of UAVs, high-rate safety-critical challenges arise from a number of aspects, including on-board reasoning and the need for continuous high-rate processing to maintain stable vehicle flight. Failure to maintain this high-rate processing, or failure of on-board reasoning, can have disastrous consequences, including damage or loss of the vehicle itself, as well as injuries or loss of human life and property in the vicinity of a vehicular incident. These issues point to UAVs as one of the most technically challenging CPS. The increasing need for increased levels of on-board autonomy highlights the necessity for increased system functionality and software complexity. Current semi-autonomous systems are only fore-runners of envisioned truly autonomous vehicles that are already on the drawing boards of vehicle developers, e.g., descendants of ScanEagle and the Tartan Racing DARPA Urban Challenge vehicle. A key requirement for the next generation vehicles is the enabling - through onboard software - of higher level functionality normally performed by a pilot (whether in the air

Community Draft

or at a ground station) or driver/passenger. These higher-level functions are hallmarks of true autonomous systems capable of self-monitoring, self-evaluation, and autonomous (though potentially collaborative) decision-making, giving rise to real problem-solving automata.

Another driving factor in the transportation domain is the constant struggle to reduce weight and volume associated with on-board processing, which translates directly to increased range (more fuel capacity, or fuel efficiency), greater performance, or capability to carry larger and more diverse payloads, e.g., sensors, passengers, etc. To take full advantage of the ever increasing power of processors, co-locating more of the vehicle's software onto a single processor is very appealing. New technologies supporting this co-location of software could help enable future vehicles to reduce their weight, power, volume or cost footprint. For example, in the automotive industry, the air conditioning compressor is electric in hybrid vehicles, and mechanical in traditional vehicles. The cyber-physical characteristics of these components differ significantly between the hybrid and traditional vehicle applications, but in each case the roles the components play and the analysis and design criteria that must be met in integrating them, are similar.

Mixed criticality has both safety and security dimensions. For example, execution guarantees require ensuring sufficient resources for computation, communication, and physical processes. Physical (e.g., fuel, power) as well as cyber (e.g., CPU cycles, CDMA transmission slots – and their bandwidth and latency interpretations) resource dimensions are interconnected with a cyber-physical "process" abstraction. For example, fuel availability and consumption rate affects vehicle trajectory, e.g., speed, acceleration, etc., as well as electricity generation for sensors and actuators involved in vehicle monitoring and control. Finding the optimum cost path (e.g., to return a vehicle to its origin) is situational and application-dependent in all of the above dimensions.

There is an increase in software size and complexity of transportation systems resulting from the addition of these advanced capabilities (e.g., autonomy, collaboration, personalization), further resulting in increased need for on-board computing resources – computers, processors, interconnecting networks, wiring, power, etc. Both miniaturization and mixed-critical co-hosting of software are important trends for coping with this need for more on-board computing resources. But, leverage of mixed-critical co-hosting renders the task of certifying these systems, significantly more challenging and costly. Accordingly, these systems that must satisfy safety, security, and other dependability requirements demand new methods and new tools to support affordable development.

In the automotive domain, computing and communications modules are used for the traditional controls applications of the *powertrain* and chassis, as well as infotainment, comfort, and convenience applications. The integration of the communication and computation modules requires a foundation of reliable and timely performance, with the ability to plug-in new components (e.g., hardware and software) into the integrated system. This demands the creation of new concepts of systems architecture and infrastructure components to support the high confidence applications of the future.

Where Are We Now?

Safety-critical and mission-critical systems have traditionally been physically separated in multiple ways. Each system had its own physical hardware and software components and

Community Draft

interfaced with the other systems via external buses, and I/O (analog and discrete). Over time, the inherent inefficiency of this approach led to the concept of consolidating like functions into a federated Vehicle Management Systems (VMS) computer separated from federated Mission Systems Computers. These exposed new internal issues and drove the need for process isolation in these federated systems. One way this process isolation has been implemented is ARINC 653, a standard for time and space partitioning for avionics systems. In this approach all applications' safety-critical and mission-critical segments are allocated to separate processor containers and the data-flow between these containers is tightly controlled through a combination of middleware and OS / kernel level mechanisms.

The criticality of an element, however, depends not just on the inherent criticality of the function it provides, but may also be inherited from the elements that depend on it. For software, this includes indirect dependencies, such as side effects due to sharing computational resources. This is where certification requirements impact solutions for mixed-critical systems. In the absence of separation mechanisms such as those provided by ARINC 653, software elements of lower criticality running on the same processor as a high-criticality element would have to be treated as higher-criticality. This can significantly increase the V&V effort required over a federated solution. Accordingly, one of the objectives for a mixed criticality solution is to provide separation between elements of differing criticality levels sufficient to guarantee that the lower criticality elements cannot interfere with the functioning of the higher criticality elements.

Current approaches to mixed-criticality systems involve a federated model that segregates functionality of different criticality levels either to different hardware components, or on newer systems, to virtual resource partitions such as those defined by ARINC 653. Multiple criticality levels are particularly prevalent in intra-vehicle networks in the automotive sector. Current automotive solutions include using separate buses based on throughput speed: Control Area Network (CAN) bus for high- speed, high-critical subsystems such as anti-lock breaking information to detect wheel slip; communication from ABS module to allow the vehicle to sense and react to a skid; and Local Interconnect Network (LIN) bus for low-speed, less-critical subsystems such as lights and door locks.

Current practice of systems design involves dividing subsystems or functions into levels such as safety-critical, mission-critical, and non-critical. Newer capabilities may also add levels such as maintenance-critical, to describe advanced functions such as vehicle health management. For example, the FAA DO-178B software certification standard defines criticality levels A through E, with level A representing the potential for catastrophic failure including loss of life and/or the aircraft, and level E having no effect on safety. The required safety of a subsystem is typically a function of its criticality and is often stated in terms of the likelihood of mishaps of the corresponding severity. Requirements for criticality level A address a catastrophic mishap with the probability that it is unlikely such a mishap will occur over the planned service life of the aircraft fleet. Producing evidence to demonstrate the required level of safety imposes an increasingly steep V&V burden as criticality increases. Therefore, affordability demands that system elements not be assigned inflated criticality levels.

Mixed-Criticality Grand Challenges

A challenge to realizing the benefits of mixed-criticality systems is that often we do not need to, nor can we afford to, test all of the elements of such a system to the highest level of assurance. For instance, the software that provides emergency stopping functionality in a light rail system is

Community Draft

of a much higher criticality than the software that provides displays functionality on the operator panel. Lives depend on the emergency stopping code; more testing is needed. A mixed-critical system would enable a single processor to perform both functions. New system development approaches are needed for mixed criticality CPS in which less-tested, lower-criticality code can safely exist on the same processor as the more safety-critical, better-tested code. This involves ensuring proper time-space partitioning – a cross-cutting technology that needs further development.

Central to CPS are physical properties including relationships between timing jitter, resource limits, and other cyber issues, and the physical portions of mixed-criticality systems must be explored. Current real-time analysis and design techniques need to be generalized to incorporate physical resources and limits (fuel, battery life, thermal thresholds). Mode-dependent changes in criticality also require a generalization of criticality away from a declarative property to an evaluated property, governed by rigorously defined adaptation properties over which verification and other forms of reasoning can be performed.

Social embedding of mixed-criticality systems (liability, ethical practice, and intellectual property) – need an ethically and legally sound and rigorous foundation (which will require contributions from technical, legal, and ethics experts) for designing safe harbors for system design, implementation, certification, and maintenance that both (1) promote and enforce principled engineering practices, and (2) encourage innovation, particularly where current impediments to progress are largely due to an absence of such a foundation. Dual problems spanning such diverse areas as intellectual property protection and verification (for example, whether properties can be guaranteed without requiring either exhaustive exploration of all states, or visibility of states that would give away trade secrets) raise a compelling set of technical research questions in this domain.

Safety for CPS needs models for expressing and evaluating safety constraints, defined over fundamental and formal (though in some cases possibly stochastic) representations of cyber-physical models. Security could be based on appropriate access lattices that generalize away from simple permission label-based approaches to approaches in which access lattices are in fact semantically rich interaction channel lattices.

Mixed-Criticality Research Challenges

Essential mixed-criticality research challenges include: communication abstractions supporting composability and adaptation; MILS concepts applied to mixed criticality; multidimensional modeling; model-based development; dynamic resource allocation incorporating fault tolerance, resource optimization, and mode change; systems mixing real-time and non-real-time, and secure and non-secure elements; and systems engineering and component integration. The notion of communication changes in a CPS view: for example, using a physical medium as a communication channel, exploiting physical locality to perform shared sensing rather than transmitting sensed values over a network, etc. CPS dependences that span communication links also involve (especially latency and other temporal) properties that cross the link.

Current practice suggests that future CPSs will be constructed using composable and configurable components. Composition of high-dependability CPS where *certifiability* is required would be enhanced by the development of CPS composition approaches that can compose not only the system, but also the evidence of its dependability required for certifiability.

Community Draft

Components and their associated metadata would then constitute the basic initial set of building blocks for future CPS development and certification. The overarching challenge is to identify, develop, and implement both an approach to certifiability and a composability framework that will support composable and incremental evidence of dependability. A systematic means to identify and assess component certifiability levels in designated contexts, even before they are implemented, is necessary to support composable and incremental certification. Formal and analytical methods for embedding attributes of certifiability within components (or their specifications) are envisioned to play a key role in the process. These certifiability attributes will have to include the system context (cyber and physical) that characterizes the applicability of the attributes.

A significant challenge for transportation CPS is the unification of safety and security. In future configurable CPS, the distinction among criticality levels for components will be blurred and variable depending on which function a component provides relative to the particular set of mission or operational requirements. In UAV systems with mixed-criticality levels, safety is still likely to be the predominant concern, yet security impacts safety and some UAV mission operations may mandate security and survival over safety. Techniques and tools have been developed and studied that characterize both system safety and security levels, but little work appears to have been done that leverages disjoint security and safety certification credit by identifying the common approaches used to address both characteristics.

A unified approach to safety and security that combines the views of threats and hazards might leverage separation mechanisms to satisfy both safety and security requirements. Information-flow analysis tools have been developed for verification of software in the security domain that can determine all possible interactions between functions and shared memory, whether directly communicating or not. Similar analysis techniques might be useful in determining non-interference runtime guarantees required for constructing Worst-Case Execution Time (WCET) estimates. The highest levels of safety often have stringent timing requirements that might not be applicable to security requirements. A standard set of core approaches that simultaneously satisfy system safety and security isolation guarantees at varying levels of criticality could reduce system development and certification costs.

Multi-core processors are rapidly becoming the industry standard. What mechanisms are required so multi-core processors can provide support for the partitioning that enables mixed-criticality CPS? In multi-core processors, many cores concurrently compete for access to different hardware resources ranging from on-chip caches to off-chip memories and I/O devices, resulting in a need for partitioning that goes beyond current approaches that focus on the CPU and memory. The problem is exacerbated by the diversity of emerging multi-core processor architectures. Novel approaches to partitioning in multi-core processors are needed – e.g., no commercial multi-core Real-Time Operating Systems (RTOS), certifiable to DO-178B level A, is currently available.

Mature well-integrated toolsets and methodologies are needed to aid in the cost-effective development of certifiable CPS. Tool chains that embody the advancements in understanding of mixed-criticality CPS are needed in order to cope with the complexity and scale of transportation sector CPS. Compositional development approaches are required, and advanced analysis and composability tools must capture the essential compositional metadata for components, systems, and architectures. Ideally, design tools will use this information to automatically integrate and

Community Draft

generate the software elements of the CPS. An additional challenge is developing V&V technologies and certification approaches that are also compositional, so that individual components in a composed CPS retain their modular certification, and tools automate the process of generating the V&V artifacts and certification evidence of other composed systems leveraging these same components. A key aspect is that tools should be able to (1) predict system behavior and performance, as well as other quality attributes, and to (2) include the environmental context of the system in these predictions.

Research Strategies and Roadmap

Mixed-criticality cyber-physical systems require new approaches both to meet increasingly acute demands for more and more sophisticated technology on Size, Weight and Power (SWaP)-constrained platforms and to deal with the assurance requirements of systems that are increasingly seeing the networking of safety-critical onboard systems with other vehicles and the transportation infrastructure. A new holistic approach to CPS offers the opportunity to change the existing trajectory of static solutions and exponentially growing V&V obligations that inhibit the deployment of advanced technologies such as adaptive control and autonomous systems that are even now emerging. This section posits a number of milestones on the road to realizing the full potential of CPS for the transportation sector.

5-Year Milestones

- Compositional development of partitioning mechanisms for multi-core computer platforms providing required isolation mechanisms for safety and security comparable to those available for current single processor systems. This would provide runtime support to guarantee that components continue to meet specifications in the presence of failure of other components on the same multi-core processor.
- V&V and certification approaches that provide the necessary evidence to permit deployment of mixed-criticality multi-core systems in safety- and security-critical applications.
- Multidimensional resource management approaches for mixed-criticality systems.

10-Year Milestones

- Mixed-criticality compositional approaches that enable very dynamic and transient workloads. Existing theory and practice are not adaptive enough to support quickly changing workloads of the mobile cyber-physical systems of the transportation domain. Workloads can change in terms of functions and resource requirements, but criticalities can also change based on operation mode. This includes approaches for dealing with workload changes due to fault detection, isolation, and recovery.
- Compositional V&V and certification approaches for demonstrating safety and/or security of *dynamic* mixed criticality CPS.
- Mixed criticality approaches adequate for system-of-system level interactions in CPS. This includes connectivity for advanced vehicle health management (prognostics, self-healing vehicles), distributed control (next-generation air traffic management, smart highways, etc.).
- Mixed-criticality approaches supporting mixed initiative and autonomous CPS.

15-Year Milestones

New paradigms for affordably and scalably ensuring the safety and security of ultra-large-scale dynamic CPS. This includes demonstrating appropriate levels of assurance for all functions of

Community Draft

the system in all modes, and coping with both short-term and long-term changes in cyber and physical aspects. Architectures, architectural mechanisms, and infrastructure at the software, cyber, physical, and cyber-physical levels will collaborate to provide the support for creation of CPS with appropriate assurance guarantees for all system capabilities.

Community Draft

Platform and Infrastructure

Participants: Ashish Agrawal, Rance DeLong, Yasser Fallah, David Garlan, Daniel Mosse, Calton Pu, Raj Rajkumar, Harini Ramaprasad, Frank Vahid, Shige Wang, Hongwei Zhang

Introduction

Transportation CPS platform and infrastructure provide essential services at various levels for application development and deployment at a broader scope across all transportation sectors (avionics, automotive, and railroad). Examples of these services include system state monitoring, location identification, instrumentation, and time. The services can be provided by physical components such as motors and sensors, or computing devices such as processors and controller boards, or integrated systems of both such as cell phones in the traffic monitoring systems and the wayside signal control in the Automatic Train Control Systems (ATC).

The fundamental requirements and challenges for CPS platforms and infrastructure are similar across all transportation sectors, regardless of the design and implementation differences of the platform and infrastructure for different sectors to meet different business objectives. Specifically, the need for platform and infrastructure service abstractions is essential for design, analysis, and deployment of components and subsystems, and for their compositions and evolutions with extended, desirable transportation capability. The need for verifiable and certifiable platform and infrastructure services is also essential for individual components and integrated systems at both design time and runtime. The differences in platform and infrastructure for different transportation sectors are then limited to different selections and configurations of their components and services, considering also the costs of ownership, impacts of failures, skills and ages of operators, and system maintenance procedures.

CPS for transportation systems requires some unique platform and infrastructure services that either do not exist or are weakly supported in today's systems. Examples of such services include instantaneous location and velocity detection, situation-aware end-to-end travel management, and large-scale system debugging. Instantaneous location and velocity detection require fault-tolerant sensing and communication of real-time, safety-related information, which cannot be achieved with today's devices like GPS and roadside sensors. Situation-aware, end-to-end travel management, which may involve vehicles from different transportation sectors, requires the services to provide timely and reliable integration, analysis, and delivery of travel path information regarding weather, traffic, and route changes. Large-scale debugging service is critical for fast and correct platform and infrastructure implementation, especially when a platform and infrastructure service is provided collaboratively and coordinately by components and subsystems from multiple transportation sectors.

Platform and infrastructure may benefit from the solutions and technologies developed to address the challenges in other areas of transportation CPS, including mixed criticality, autonomy and control, infotronics and infotainment, and V&V. As an example, the techniques for compositional and incremental V&V will help platform construction and evolution, given the fact that a component providing the platform service could be a CPS itself, whose implementation must be verified and validated. Depending on the deployment and runtime management strategy of a platform service, the criticality of the service may change, which requires the techniques and solutions from the research in mixed criticality.

Community Draft

Where Are We Now?

Many of today's transportation systems and devices are capable of providing basic infrastructure services for design, implementation, and deployment of CPS. In aircraft, devices and controllers that provide services enable flight control in the Traffic Collision Avoidance System (TCAS), by-wire control, and an enhanced ground proximity warning system. In automobiles, platform services include cell-phone and GPS-based services for location identification, traffic detection, and path rerouting; advanced control services for smart adaptive cruise control, lane departure warning, traction control, and automated parallel parking assistance; and offense prevention services such as drowsiness and drunk-driving detection, and phone call/text messaging blocking for safe driving. Similarly, the railroad systems are devised with platform services provided by the automatic signal system and advanced civil-speed-enforcement system. Although these services have been deployed in existing systems, their capability and accuracy are very limited. For example, the platform service using commercial GPS can become unavailable or inaccurate in the city area, causing the location-based applications to fail.

Further, advanced technologies for computing platforms allow more and more smart sensors and actuators deployed in the automobiles, aircraft, train systems, and roadside infrastructure to be used for monitoring and control. Various networks and protocols for wireless and sensor networks, dedicated short-range communication (DSRC) and radio frequency identification (RFID), and wired networks such as high-speed and time-triggered Ethernet, CAN and FlexRay also are widely adopted in transportation CPS. New processor technologies such as multi-core, FPGA, and system-on-a-chip have been introduced in transportation CPS. Adoption of these new technologies greatly improves the capability of the platform. The services for hardware detection and fault separation also exist in today's hardware, operating system, and middleware.

Existing development platforms and infrastructure support creation of individual components/subsystems of transportation CPS. Theories and techniques for modeling, architecture design, analysis of a given single property, code generation, code verification, and simulation are mature and have been applied to many stand-alone systems/subsystems. Development tools for these activities are commercially available. Services such as data logging and benchmarking also exist. However, all of these techniques, tools, and services are applicable to only a small set of applications, and often require significant efforts to integrate them into a system development.

Despite the existence of CPS-enabling services and CPS-capable devices in current transportation platforms and infrastructure, they provide only a partial solution to the inherent challenges. Adequate abstractions to represent existing services and new services, as well as corresponding methods and techniques, are necessary to help system design, composition, verification and validation, and the reconfiguration of transportation CPS systems.

Platform and Infrastructure Grand Challenges

The grand challenges of transportation CPS platforms and infrastructure are to improve the use of existing infrastructure resources and services and provide fast, safe, and reliable transportation. Specifically, the challenges include:

- Minimizing traffic-related negative impacts – Across the avionics, automotive, and railroad sectors, the ultimate goal is to eliminate fatalities and accidents. Travelers desire minimal end-to-end travel delays with seamless, multi-modal transportation between the original location and the destination, which consequently requires dynamic travel planning and traffic management to avoid traffic jams. Energy consumption is another

Community Draft

challenge for future transportation. Not only should travel delays be reduced, the total energy consumed during travel should also be minimized in order to maintain sustainable modes of transportation under increasing energy demands and costs. Future transportation CPS is expected to help protect the environment through reducing greenhouse emissions. Given that transportation CPS may be large in scale and evolve over time, these systems should be able to maintain overall safety in the face of emergent behaviors from system components and/or operators.

- Maximizing utilization of the existing transportation infrastructure for affordable travel – Transportation CPS enables real-time monitoring and management of the traffic and vehicles, e. g., airplanes, automobiles, and trains. This makes it possible to optimize the usage of existing transportation infrastructures, including airport gates and runways, highways and city streets, and railroads and train stations. With such optimization, more and better transportation services can be delivered to the passengers and customers without increasing the investments for additional new infrastructure, which should eventually lead to stable, predictable, and affordable modes of travel.
- Educating and training the public and society to exhibit safe behavior in the new transportation CPS environment – As the CPS introduces new capabilities and services, such as autonomous driving and seamless multi-modal travel, it is critical that the general public understand the rules and limitations of these novel transportation services, and regulate their behaviors. The transportation CPS, on the other hand, should include mechanisms to minimize false alarms and avoid negative social impacts.

Platform and Infrastructure Research Challenges

Addressing the grand challenges of transportation CPS demands new research in transportation CPS platforms and infrastructure. The research challenges include theoretical foundations and effective methodologies for design and implementation of dependable, affordable, and trustworthy platforms and services. Specific research challenges include:

- Support for autonomous transportation – An autonomous transportation system is key to achieving near-zero fatality, eliminating accidents and traffic jams, and reducing emissions. To create an autonomous transportation system, it is essential that applications execute on intelligent, adaptive platforms. The services and components of the platforms and infrastructure must be situation-aware and self-configurable so that they can adapt to the new operational environment to maintain safe and dependable services when unexpected events happen. Consequently, the components and subsystems in the CPS platforms and infrastructure should be interoperable. The technologies used in current platform and infrastructure implementation provide very weak support, at best, for adaptability and interoperability.
- Runtime assurance of system properties such as performance, safety, security, and reliability – Different from traditional information processing systems and embedded control systems, components of both cyber and physical systems must provide designated levels of assurance. Existing technologies are designed mainly for cyber system assurance. New research is required to provide assurance of physical elements and their interactions with cyber systems in a CPS platform. Further, transportation CPS typically requires fail-operational rather than fail-silent modes. The platform, therefore, is required to provide diagnostic and prognostic services for runtime health monitoring and management of both physical components and cyber systems. Dynamic management of

Community Draft

resources and capacity, including power, processors, and I/O devices, as well as physical elements like brakes and engines become critical services in the transportation CPS platform. To assure safe emergent behaviors, sensors and actuators should have safety management and fault-tolerant guards built in them. How to implement such services in sensors and actuators connecting to physical components in a dependable and affordable way has not been well understood.

- Methodology for transportation CPS platform design – The transportation CPS requires a comprehensive design environment that supports a dependable-by-construction design approach. Existing modeling techniques and tools do not meet these needs because the design requires that the environmental and physical aspects, with their uncertainties, be captured. Transportation systems are usually safety-critical, so the safety specifications, which may involve multiple vehicles in different types, must also be captured. This requires new model formalism to represent the platform and new methods to reason it at the system level, with a proper abstraction.
- System evaluation and certification for reduced and focused testing -– It has been well-known that system testing is costly and cannot ensure detection of all errors. Recent model-based system development and integration have shown to be promising as they allow formal system analysis and reasoning. Extensive research is necessary to expend the model-based solutions to larger and complex systems with platform and physical elements included. It is desirable to feed the runtime behaviors back into the design, especially when abnormal events happen, in order to identify and resolve the platform-related issues that may not be detected during development. New theories and methods are required for component and system certification, as the platform may be composed of cyber and physical elements from different sources.
- Theories and models for fault modeling and composition. – As the scope of transportation CPS may involve multiple vehicles of different types, traditional fault models designed to capture the faults of individual vehicle or transportation mechanisms are not adequate. The new fault model should capture not only the system-level faults beyond the scope of individual vehicles, but also the implicit assumptions of the conditions causing the faults and their impacts. Theories for composition of fault models should also be developed.
- System deployment without brittleness – The existing transportation infrastructure cannot be totally replaced in a short period of time. Therefore, the new advanced technologies must be deployed incrementally. The platforms and infrastructure for transportation CPS must support system operations with co-existence of the components constructed using both old and new technologies. The deployment of newly designed systems should evolve from the current practices, and should not disrupt the operation of the existing system.

Research Strategies and Roadmap

Addressing the platform and infrastructure challenge requires close collaborations and joint efforts from government, academia, and industry. New research should focus on new platform and infrastructure services adequate for transportation CPS, new development environments and methodologies for safe, affordable, and assured transportation, and education that helps improve understanding and deployment of CPS platforms and infrastructure.

Community Draft

5-Year Milestones

The short-term research should focus on providing platform services for each sector.

- Understanding the dependencies among the systems, vehicles, and human and operational environments – The research should define services and interactions to support construction of a better and safer platform for each transportation sector, and allow the construction of communicating systems, such as vehicle-to-vehicle and vehicle-to-infrastructure networks, that operate with given communication modes. Techniques to improve human-machine interfaces, including information overload avoidance and symbolic display of system status and operation, should be investigated to provide better operation assistance. Services for safe behaviors in human-machine interactions, such as disabling personal communications during flight departure and preventing text messaging during driving, need to be developed.
- Developing services for reliable and assured transportation – These include fault-tolerant support with unified quality-of-service metrics (time, reliability, security, etc.), particularly in the areas of wireless communications, tamper-proof hardware and software, fault-tolerant controls for X-by-wire applications, basic data logging for an individual device/component/vehicle, and basic situation monitoring and environment sensing. Services for instantaneous location and velocity detection should also be investigated extensively, which may consequently require fault-tolerant and accurate GPS services in any environment. Service supporting dynamic performance optimization and power management to allow platform adaptation at the individual vehicle level should be provided.
- Creating platform development methods and tools for component-level predictability and system assurance in each sector – This research focuses on the development methods and tools for analysis, simulation, and evaluation. The research should also include the methods and solutions for platform abstraction to support separation of time and space when building applications with hard, soft, and non-real-time constraints as well as with mixed criticality. Model-based methods and solutions for modeling various cross-cutting properties such as fault-tolerance, safety, performance, and power, for reasoning about and analyzing these platform properties, and for allocating the partitioned requirements to platform and infrastructure components, are also essential. Domain-specific models, as well as models capturing the interactions and interfaces of different domains, need to be developed. Advanced tools and methods are needed to analyze interactions and interfaces among networks, real-time computation, and control.

10-Year Milestones

The research should focus on providing integral solutions to cross-sector CPS services, which involve more than one sector.

- Developing platform services and techniques for system-environment interactions across multiple sectors – The platform should define a contract between operators and vehicles in different sector and allow use of different modes of communications in the dynamic, uncertain environment to achieve fault-tolerance and reliability.
- Providing additional runtime services uniquely for transportation CPS with new quality-of-service measures – One such new runtime service is cyber-physical markers, which utilize the road-side, railroad-side, and infrastructure devices to offer real-time safety-critical and convenient information. Services for software updates, runtime data log, and environment and context sensing with security and privacy protection should also be

Community Draft

included in transportation CPS platforms and infrastructures. Platform services for integrated, timely, and robust control with physical constraints in different transportation sectors should be studied and improved to enable effective controls corresponding to different physical characteristics. For example, the by-wire control can be different for automobiles, avionics, and trains. Further, the quality-of-service measures, including time, fault-tolerance, safety, security, and adaptation, are different from the measures used in traditional stand-alone transportation systems. New quality-of-service measures across multiple transportation sectors and vehicles need to be defined and evaluated.

- Enhancing the development environment to improve end-to-end predictability and composability of platform methods for capturing decision-making points for adaptation and degradation – Data-mining methods and tools are also desired to analyze the runtime data logs to identify potential design flaws and undesired behaviors. Debugging services are needed for transportation CPS platforms across multiple sectors. Multi-dimensional feasibility analysis and management methods should be developed to support the design tradeoffs to better achieve end-to-end requirements.

15-Year Milestones

The research should focus on providing system-level, end-to-end platform solutions involving all transportation modes.

- Developing platforms with services that support system-environment interactions at the end-to-end system level that could be applied in all transportation sectors – The key issues to be addressed include integrated networks and end-to-end communications among all kinds of vehicles and transportation infrastructures, common symbols for warnings, alerts, and emergencies across different transportation means, convergence of different communication modes, and unified quality-of-service across all involved networks and communication modes.
- Developing complete platform solutions to provide situation-aware capability, support fully autonomous control with better sensing capability to capture the dynamically changing operational environment and context, and adjustable system configuration to maintain the desired system behaviors and properties – The solutions should include the composition of needed services that manage environment uncertainties, such as road conditions, weather conditions, and status of sensors and actuators at different conditions. Additionally, dynamic end-to-end performance optimization for platform adaptation across multiple transportation vehicles should be developed.
- Building a comprehensive development environment for transportation CPS platforms with methodologies that support dependable-by-construction; system-level assurance and certification with reduced and focused testing; analysis and evaluations of integral functions with multiple levels of criticality, performance, safety, and security; and methods and tools to detect and automatically fix assumption mismatches.

Community Draft

Autonomy and Control

Participants: Jonathan Sprinkle, Emilio Frazzoli, Sriprakash Saratim, Cliff Wang, Simon Cobb, Timothy Chang, Greg Sullivan, Xiuzhen Cheng, Linda Bushnell, Patrick Benito, Jeff Maddalon, Bill Schoening, Panos Antsaklis, Dan Work and Alex Bayen.

Introduction

The "vehicles of tomorrow" seem always to lie just beyond next year's models. Science fiction promises self-driving cars, personal air transport, intelligent highways, and other applications that always seem to be just beyond the reach of today's technology.

The cruel irony is that the shrinking cost and expanding computational power predicted by Moore's Law extend the reach of that technology to single handedly tackle these problems. This is in contrast to traditional information technology, which improves with increases in speed. In the application of CPS, the ability of our transportation systems and infrastructure to continue to support higher layers of autonomy does not simply require more computational power – if it did, we would see new and dramatic capabilities for autonomy rolled out each year in consumer automobiles. In fact, we are nearing the bounds of our existing approaches to design, analysis, implementation, test, validation, and verification of autonomous cyber-physical systems. New methods, tools, strategies, and abstractions are necessary in order to allow these systems to interact with one another at societal scales, and with an acceptable societal risk.

The evolution of autonomous transportation systems promises to save lives, increase throughput, and even reduce fuel costs. But an inductive approach, where technology shown to work on a single vehicle can simply be replicated, will not solve the grand challenges that this domain faces. The challenges of scale affect both how to control or influence the system and gathering knowledge (or estimates) while it operates. At scale, influence is arguably more important than direct control, especially when human operators have tremendous authority (e.g., ground vehicles). The challenges of observation involve the interaction of vehicles through sensors (not direct communication), as well as sensing vehicle health. The challenges of interaction involve direct exchange of (state) data, as well as scalability to some extent.

Thus, the grand challenges can be summarized as:

- **Challenges of Scale** – (1) *Control/Influence*: Multi-level supervisory systems that make decisions and perform roles outside their original design parameters,. Because our state- of-the- art requires explicit design boundaries, even small extensions to the application of an unmanned vehicle may violate certain assumptions. Yet transportation systems today, when under human control, frequently operate while exceeding the designed thresholds of the infrastructure (e.g., a traffic jam). In a flexible supervisory system, the interaction between groups of vehicles should be indistinguishable from interaction between a central controller and vehicles. This would permit vehicles to come on/off the grid easily, without sacrificing vehicle safety, or the goals of the mission. Moreover, influencing some portion of a group of vehicles to change behavior may be sufficient, but individual vehicles are not controlled by a central authority (e.g., traffic calming). (2) *Knowledge/Estimation*: Sensing the overall state of the system at scale is difficult, as algorithms for where to deploy sensors must take into account the changing physical state of the system. Mobile sensors may improve this ability, but knowledge of how to move sensors to gather the best information is nontrivial. This is further complicated if sensors must navigate in the medium they are sensing (e.g., water-based mobile sensors). As such, estimates of parameters (*a priori*) may need to be updated as the system evolves.

Community Draft

- **Challenges of Observation** – In order for autonomy to increase in transportation, the fact that a vehicle is under human, or computer, control should not be distinguishable to a human observer. This is the so-called Autonomy Turing Test. There are many different aspects of this challenge. Is the behavior of the system experiencing a loss of one of the critical axes of CPS (communication, control, computation) distinguishable from that of a human in the same scenario? Would a human operator be able to pass the Autonomy Turing Test? In the event of loss of vehicle health, how will the controller know to operate in degraded mode? These issues *must* be solved before unmanned and human-operated vehicles can safely coexist.
- **Challenges of Interaction** – Interaction between autonomous vehicles and human-controlled vehicles is safe, intuitive, and understood, including cross-platform interaction (e.g., train crossings and ground vehicles). Important tests would include "mob scale" interaction (one autonomous vehicle, lots of human-controlled vehicles, and vice-versa). These challenges dramatically extend the state of the art in swarming, but also rely on expected behaviors for the rest of the vehicles. In the event that the global state of the system can be estimated, and vehicles in a neighborhood can be observed, in what way should an autonomous vehicle interact with its neighbors, or control itself, to globally optimize the situation?

What Can We Do Well?

We can define autonomous systems for highly structured and controlled domains. This is widely evidenced by the recent successes of teams participating in the DARPA Urban Challenge (where software was a major factor in the system's behavior). In this case, the participating teams were heavily constrained by competition rules, in addition to the rules of the road. There are many other examples. Driverless rail trams at airports shuttle passengers back and forth between terminals. Unmanned military vehicles can perform loitering and steady-flight in surveillance operations. Commercial aircraft regularly perform auto landings, and steady flight is performed by autopilots.

In the design phase of autonomous systems, we can establish bounds of "good behavior" and design for these boundaries. Further, given anticipated "good" behavior by other vehicles, we can interact safely with small numbers of vehicles. This may include vehicles under the control or jurisdiction of the designer, or vehicles under the control or jurisdiction of a common authority (e.g., the air Traffic Collision Avoidance System—TCAS). Even though many of these systems do eventually rely on a human to make a decision (TCAS is especially noted here), the decision-makers rely heavily on information they get from the computing and communication components of their physical vehicles.

Technology that *can* enable increased autonomy is already finding its way into consumer automobiles, which are remarkably well instrumented. Data and information exchange between vehicle components is now a state of the practice, and standard system components can receive state information from other components (based on design time decisions), providing velocity, acceleration, traction information, etc., for the purposes of infotainment, navigation, and other conveniences. More advanced sensor systems, such as radar and lidar, and cameras are used to notify drivers of potentially unsafe situations, such as lane departure, or impending collision. However, such systems do not *take control* of the vehicle to prevent collision, or maintain the current lane, as the intent of the driver still trumps the vehicle's sensor fusion algorithms. Rather, safety behaviors such as tensioning the seat belt prior to an impending collision, or shaking the

Community Draft

steering wheel to get the attention of the driver, do not immediately alter the situation but rather confirm the driver's intent, to ensure prompt driver reaction.

Why Can't We Declare Victory?

As the number of interacting vehicles increases, the weaknesses of our current state of the art are brought to light. Fundamentally, it is the uncontrolled physical environment that prevents us from simply improving existing designs and attacking these grand challenges. For example, consider automotive autonomy, where many of the "inexpensive" technologies that have permitted driving autonomous ground vehicles depend on a static world, with a vehicle that has precise global measurements. Vehicle-to-vehicle and vehicle-to-infrastructure communication technologies *could* enable mitigation of these weaknesses, but there are key weaknesses in the size and stability of large-scale (greater than 500 nodes) ad hoc networks. Control designs for autonomy must be robust enough to withstand outages in communication, be able to make decisions based on noisy sensors in a local environment, and scale to varying speeds in all kinds of driving conditions.

As traffic increases (in the air and on the ground), the dynamics of the flow of vehicles change dramatically. If a field of vehicles must operate together, meaning individual vehicles will converge on their own behaviors based on the behavior of others, changes in the flow of vehicles must be sensed. In aerial vehicles this is particularly important, due to the velocities required to maintain lift. Clearly segments of the entire field of traffic can be partitioned and solved, rather than treating the entire field of traffic as one system that must communicate and coordinate in real time. However, understanding the boundaries of these relevant segments, and the timing requirements for communication among necessary vehicles, is not well understood. Moreover, vehicles must be able to operate alongside human-operated vehicles, so notions of sharing future predicted state, etc., are not feasible.

These issues in control are fundamental, yet breakthroughs in these areas will not immediately enable solutions to our grand challenges. Once solutions are in place, the implementation and designs *must* be validated and verified, as replacing a human with a computer in a societal scale requires rigorous proof of safety. This places strenuous demands on the existing methodologies for design and implementation, which are not verifiable at scale. Given that the complexity of these systems will continue to increase, the scalability of V&V methodologies cannot keep pace.

Significant R&D Challenges

Our research challenges support the grand challenges.

1. **Dynamics –** In today's systems, we constrain the dynamics to an operable range, and if the system departs from these dynamics, we change strategies (e.g., a hybrid systems model). In a truly cyber-physical design, however, the design can be functionally dependent on the dynamics, rather than use the dynamics as a design criterion. In order to permit such an approach, however, we must address the following kinds of questions:
    - How do we intuitively represent the scope for autonomy, with respect to operating conditions?
    - If operating conditions change, is the entire design subject to modification, or are the operating conditions tersely represented?
    - If the dynamics of the platform (or of the system) change, but the operating conditions remain the same, is the design subject to major modifications?

Community Draft

If a large number of autonomous vehicles are to be controlled by a single operator, we must have some understanding and expression for the interaction between vehicles, and the protocols for this communication. However, the latency of this communication, as well as various issues such as leader/follower and collision avoidance are all necessary considerations. The following kinds of questions must be addressed from this perspective:

- How do we establish the bounds and protocols for interaction between vehicles (and humans, and groups of vehicles) in expression of high-level decisions? This should go beyond design-time protocols (or even pre-mission defined protocols). The emergent behavior (as such) should be trustworthy.
- How do we define the dynamic allocation of responsibilities and collection of various data?

In essence, we address in this research challenge how we can safely, and with confidence, stretch the system past its original design, as new capabilities or instabilities are added. Again, if such capabilities, dynamics, and other constraints are part of the *execution* parameters, then changes could be made based on ranges of values, or perhaps it could be determined in real time that the new dynamics are still possible to execute safely.

2. **Methods and Methodologies –** Although tremendous progress has been made on control strategies for the grand challenges and research challenges, the implementation of these strategies as well as the V&V of the software are still of concern.

   This concern motivates research in methods – such as domain-specific modeling, verifiable software synthesis, parameter identification, system verification, etc. – and the methodologies that permit such methods to be carried out. What scalable analytical methods and methodologies exist for CPS? In order to make progress in other areas, we would benefit from scalable methods and methodologies (including tools) that permit expressiveness of boundaries and dynamics.

   Methods that provide some risk of failure of control/computation/communication and sensing would assist in giving measures of confidence. These methods could give new dimensions of confidence, in addition to the well-known PLOC (probability of loss of control), e.g., from avionics. Of particular interest are failures of communication that are amplified in a way that results in failures of control (likewise for computation). Understanding these subtle interactions is a new field of research in CPS, because the three areas of concern can no longer be logically separated for all systems, and some fundamental interconnections between physical, computational, and communication aspects are implicit in the system's application, construction, or execution.

3. **Composition –** It is difficult for an observer to understand instability of interacting systems. With respect to the stability (control) or intuitive behavior (autonomy), this understanding is crucial for vehicles and humans that will interact (or perhaps intervene) in the event that instabilities are observed. This aspect is crucial for transportation systems where passenger comfort is important. For example, erratic steering and hesitant acceleration by a student driver are expected, but passengers in the car sense these erratic behaviors more acutely than observers outside the vehicle. Nonetheless, a human driver with jerky control of the vehicle will most likely be able to safely drive the vehicle to its destination. This is not necessarily the case for an autonomous vehicle, as the jerkiness of

Community Draft

the controller could be an indication of saturated input values, or an optimizer that is failing to converge.

The main point is that sensing state values may not be enough to determine that a subsystem is behaving erratically. For systems where state is also noisy to measure (e.g., camera sensors detecting pose of another vehicle), it could mean that the composition of two vehicles (or two components of a subsystem) may no longer be valid, and transition to a new control mode should take place.

Additionally, there are outstanding research questions regarding the preservation of various properties under composition. Stability, timing, and latency may all be affected by the composition of various behaviors or models. Understanding how composed models behave with the addition of other vehicles, as well as the changes in the dynamics of the global system, or the change of dynamics of subsystems, is a significant research challenge.

4. **Cost Constraints –** Affordability is a major constraint for CPS transportation systems. The current state of the art is approaching price limits due to the complexity of today's systems: As more communication and computational components are introduced, the cost continues to climb. Given constraints on affordability, what amount of robustness is acceptable? In aviation transport, robustness is a design constraint, and cannot be sacrificed; thus, the cost of the system climbs until the robustness demanded by the marketplace is met. However, this system cost cannot continue to rise uncontrollably, and the cost for certification increases more rapidly as more software is introduced into the system. Sacrificing safety for affordability, however, is a troubling alternative. A better solution is continued research into verification of system properties under operational constraints.

   For ground transport systems such as automobiles, there are no well-understood robustness constraints for autonomy, and for control, the state of the art continues to expand these boundaries. Clearly, consumers will not accept arbitrary increases in end-user costs, but will decreased safety be accepted if the autonomous system is perceived as "less intelligent" than a human's behavior in the same situations. This issue must be addressed by policy, as well as technically, to determine what overarching design constraints must be satisfied (e.g., an autonomy band of airspace, or special lanes/times for autonomous vehicles) in order to support increased autonomy in transportation.

5. **Context and Intelligence –** The design and validation of predictable and safe behaviors in new contexts, or changing contexts, is also a significant research challenge. Generally, responses to change must be well reasoned and safe. Implicit in responses to changes in context is reliable sensing and estimation of the state of the global system, as well as local system knowledge. As context and operator changes affect the system, the system should behave as expected. Consider an aerial vehicle that identifies an on-board fault and must now operate in degraded mode. Changes in context include: What portions of my mission are necessary and what portions are expendable? What is my plan for landing? What airspace or fuel restrictions constrain my future actions? Have my dynamics, or controllability, changed due to my change in context? Holistic methods to address these shifts in context may extend existing work in control, particularly hybrid models and fault-tolerant systems.

Community Draft

Closing the loop on vehicle and system intelligence is a significant research challenge. Not only must approaches be scalable, but they must also be transparent (to permit certification and open-box testing). Adaptability and learning will be important areas, as well as reasoning and meta-reasoning.

Research Strategies and Roadmap

There are many organizational strategies that can move this research agenda forward. Joint industry/academic/government collaborations benefit the entire community. Benefits include participation in realistic scenarios, which in turn implicitly push the boundaries of academic pursuits. Industry can then examine the output of academic pursuits, and validate it, or suggest areas of interest to improve it.

Discipline integration is a further strategy to move the research agenda forward. Joint courses in CPS, and "touchable" examples, illustrate interaction with system elements outside the regular abstractions.

There are also technical approaches to advance this research agenda. Hybrid and embedded software systems can be considered a forerunner to this research area. The continued examination of this abstraction, and its enhancement to capture the relevant issues of CPS, is a promising approach. Distributed parameter systems and the optimization (or discovery) of those parameters constitute another promising area of investigation. Given the ubiquitous monitoring now available due to the success of the wireless sensing agenda, we have more sensor sources (and data) than many algorithms may have been designed to consider.

Model-based design abstractions have traditionally depended on the ability of humans to correctly specify the design, or structure, of a system. For CPS, some aspects of the model (whether its dynamics, sensing profile, etc.) may need to be discovered and transformed from data into a model. This can be considered an extension of system identification (where the dynamics of the model are discovered), but also applied to communication paths, available computational models, etc. Further, generic decisions may be designed that do not have (during the design phase) the exact data needed to validate them.

The consideration of time (as well as timing) as a critical design component for the system, or for its components, is a crucial area for investigation. The traditional dependence of control, communication, and computation on "fast enough" technology begs the question of whether "fast enough" or "exactly fast" is the most appropriate abstraction. In some scenarios, "too fast" may cause instabilities or undefined behaviors. Timing also relates to issues of embedded humans. Specifically, how does latency (input, or observation, or both) affect the human's ability to safely or stably provide inputs to a remote system, or even a local system. Operation with imperfect (or uncertainly imperfect) information is another crucial area for investigation. The distributed sensing ability of many systems necessarily introduces latency, or uncertain accuracy, of timestamps. Operation in this environment is a necessary advancement for CPS design and implementation.

The final point of discussion in this area focused on the well-established needs for an agenda for testing, verification, and validation. Specifically, V&V in today's systems are centered on the adherence of system designers to the design process, and not necessarily based on evidence that the design meets requirements. Many requirements checks emerge from the testing phase, but it is impossible to cover all of the application space in testing, due to the combinatorial explosion of states as new computational components are added to the system. Methods and tools

Community Draft

supporting methods to verify design parameters from the implementation space to the application space are of critical importance to permit the continued scaling of complexity.

5-Year Milestones

- New models to account for latency and timing in human interactions with autonomous vehicles
- Foundations of new theory for large-scale interoperation of vehicles
- Availability of a family of open testbeds, or interface to simulators, to permit wide-ranging interactions among similar platforms and between researchers
- Proofs and acceptance of timing-enabled programming models to support large-scale interoperation, especially for networking/communication between vehicles/infrastructure
- New and appropriate sensing models that apply to global state, and their effects on control or "persuasion" decisions
- Composable design and implementation models, which permit appropriate abstractions to reduce complexity, while preserving design- and implementation-space coverage
- Evidence-based software verification becomes a certification criterion for safety-critical portions of autonomous systems
- Foundations for smart infrastructure, including how much infrastructure is needed to support legacy vehicles
- Demonstrated, scalable interactions between large groups of vehicles, using information exchange
- Large-scale interactions between vehicles in simulation or hardware, show convergence of algorithms in "mob-style" situations

10-Year Milestones

- Scalable methods for validation and verification of algorithms, based on requirements satisfaction
- Results in CPS identification – going beyond just identification of dynamics to also include discovery of networking and computation models – begin to scale up
- Robust behaviors in the event of damaged sensors (or computers), or poor vehicle health
- Demonstrated safety of autonomy features of vehicles in the field
- Smart infrastructure is tested in certain locations

15-Year Milestones

- Pervasive use of model-based technologies, with certification based on evidence, in the design, analysis, implementation, validation, and verification phases
- Public and practitioner confidence in operation alongside autonomous vehicles, including handoffs between autonomous controller and operator. This includes situations where the vehicle recognizes that its health no longer permits autonomous operation
- Demonstrated improvements in passenger comfort, fuel efficiency, long-term reliability, safety, or other metrics
- Demonstrated resilience of autonomous vehicles when the network is under attack or faces other pathological issues of communication, including bounds for safe operation

Concluding Remarks

Many of the technological foundations in networked systems and x-by-wire behaviors have provided a substrate for extending these systems with autonomous behavior. However, just because communication is possible, and actuation is made simpler, it does not follow that

Community Draft

systems can immediately use this infrastructure in a safe and verifiable manner. New theory is needed in understanding how the physical limitations of the transportation infrastructure, as well as those of the vehicles, affect the global state of transport, and how vehicles should be encouraged to converge on global (and not local) optima. A new demonstration of more robust autonomic vehicle behavior that is indistinguishable from human control – especially in degraded modes of operation – is necessary in order to deploy autonomous systems on a societal scale. Finally, "legacy" transportation systems will not disappear overnight, and any new autonomous control systems will need to operate alongside human systems. We must have some measures of how safe this interaction will be, while not sacrificing the benefits of autonomy.

Community Draft

Infotronics Including Infotainment

Participants: Rahul Mangharam, K. Venkatesh Prasad, Radha Poovendran, Rance Cleaveland, Krishna Sampigethaya, Nirupama Bulusu, Mingyan Li, Dio de Niz, John Scoredos, Weisong Shi.

Introduction

Today's commuter daily spends an average of: 89 minutes in ground vehicles[15], or 127 minutes in a domestic flight[16], or 50 minutes in a train journey[17]. During the past two decades, the interior comfort and conveniences of automobiles have steadily grown to offer the average driver or passenger one of the most precious entities in their lives – a private block of time with a personalized ambience (e.g., climate and audio controls). Similarly, the modern aircraft as a means of mass transportation aims to continually improve the quality of air travel through on-board personalization controls (e.g., ambient light and window transparency) and in-flight entertainment (e.g., choice of first-run movies).

With the ubiquitous presence of personal cellular phones, travelers quite naturally are increasingly using their transport vehicles as mobile and digitally connected offices and living rooms. Automobiles have undergone a significant metamorphosis – from being a system composed of just "built-in" modules to a "system of systems" whose information architecture involves interactions between combinations of built-in modules, passenger brought-in devices, and "beamed-in" services from the external infrastructure (road-side beacons, other automobiles, satellites, or the more generic "smart-cloud"). With this disruptive transformation, the automobile is in the process of transcending traditional mechanical, electronic, and software boundaries and has taken on a great new technological challenge of orchestrating the complex interactions within the built-in system and across the brought-in and beamed-in systems to seamlessly deliver a new set of vehicle- based services. Driven by market forces, such as making airplanes fuel-efficient, cost-effective, and passenger-friendly, air transportation is also moving towards system- of- systems architectures – recently with the transition to an integrated modular avionics architecture, now with digital links to connect aircraft with geographically spread off-board systems of air traffic control/airlines, and in future with connections between aircraft as well as potentially third parties offering services to brought-in passenger devices.

Meanwhile, the demand from consumers and society at large is incessant. Automotive consumers seek continuous connectivity with the external infrastructures (e.g., real-time traffic information, instant travel-related notifications) and expect to be productive and entertained with their digital media devices and services. Airlines demand seamless connectivity with ground infrastructures (e.g., updating flight schedules of airline's fleet), while their passengers seek convenience and entertainment. Society is expecting to harness the power of inter-networked systems to reduce roadway/airspace congestion by improving traffic-flow management, such as during rush-hour conditions, accidents, special events, or emergency circumstances. Automobile manufacturers are responding by investing in the research and engineering of semi-autonomous systems, e.g., adaptive cruise control, x-by-wire systems. Over all, with the "commoditization" of key computing, communications, and control system technologies, automobiles are to become more

Community Draft

[15] United States Department of Transportation, Federal Highway Administration, Highway Statistics, 2006.
[16] Bureau of Transportation Statistics, Airline On-Time Performance Data, http://www.transtats.bts.gov
[17] Center For Urban Transportation Research, Public Transit in America: Findings from the 1995 Nationwide Personal Transportation Survey, (www.cutr.eng.usf.edu), Table 4-13.

programmable so that aspects of the major vehicle systems including power train, chassis, body, and infotainment systems will be amenable to remote diagnostics and life-cycle maintenance. Similarly, aircraft manufacturers are investing in research and technologies to make the modern aircraft information-rich, automated, and more autonomous, e.g., condition-based IVHM, decentralization of traffic control.

In anticipation of the challenging convergence of information systems with road vehicle electronics, the automotive industry coined the term *infotronics* in 1996.[18] Infotronics, broadly speaking, is the information-level abstraction of the interactions both within the built-in systems of the automobile and across the brought-in and beamed-in systems. The primary focus of infotronics is to deliver the right information at the right place and time.within the vehicle and its connected elements to provide the expected set of services to the automobile itself and to the driver, passengers and, more generally speaking, to society at large. Infotronics include the computation components and communication protocols for sensing, actuation, and control at both short time-scales for vehicle trajectory control and longer time-scales for route planning and navigation. One subset of the broad area of infotronics is *infotainment*, which focuses on the human consumption and generation of travel-related information and general in-vehicle audio or video (for the rear-seat passengers) entertainment. In aviation, *"e-Enabling"* can be considered as the counterpart of infotronics, allowing the aircraft to function as a node in an enterprise-wide information network, ensuring that all air travelers receive the optimal information required to react to a wide variety of situations.[19] Such e-Enabling infrastructures and services are expected to benefit airline operation centers, flight, cabin, and maintenance crews seeking enhanced situational awareness, information availability, and management features, as well as passengers who want to use personal electronics to access off-board applications.

As infotronics in both automotive and aviation evolve beyond the boundaries of a single vehicle to a network of vehicles that communicate in a peer-to-peer, or more appropriately, "V2X" (or vehicle to something outside the vehicle) manner, there are two major departures from traditional real-time systems: (1) Coordination and control must span multiple network and control domains in heterogeneous vehicles where, in addition, the vehicle operators may not be fully compliant with optimal control decisions and (2) the broad dynamics in environment and topology result in changing resource requirements and task priorities in time and space. With these fundamental changes in the system model, it is essential to extend our current understanding of real-time and distributed computing theory to the design of spatio-temporal networked control systems.

As vehicle-centric and consumer-centric information and control become tightly coupled across multiple levels of criticality, new technologies are needed to realize long-term automotive goals– vehicles don't crash, experience minimal congestion delay and are fun to drive – and aviation goals – airplanes are safe and experience optimal gate-to-gate flight times, supporting consumer business interests and public well-being. Key goals for infotronics include, but are not limited to:

- ***Vehicles as Sensors, Actuators, and Controllers*:** Automobiles will be equipped with sensors to allow for global efficiency in fuel consumption, travel times, and traffic network design evolution. Meanwhile, NextGen will use data-link-based air traffic management and onboard

[18] The 1996 Convergence Council comprising the research and engineering leaders of the Detroit 3 automobile manufacturers and the supply base.

[19] Boeing Frontiers, 02:04, August 2003. http://www.boeing.com/news/frontiers/archive/2003/august/i_ca1.html

Community Draft

GPS for highly accurate surveillance and close navigation of aircraft. As vehicles increasingly participate in wide-area sensing, they become agents of actuation for both short and longer-time-scale tasks. Currently, automobile safety focuses on features within the vehicle, while aircraft safety relies on well-defined protocols and procedures with heavy human intervention. However, with networked vehicles exchanging safety-critical messages, the safety parameters broaden. Vehicles approaching an unsafe or congested scenario will be informed in advance and operators can make safe and timely decisions. Vehicles will communicate with the traffic infrastructure to exchange equally time-critical as well as less time-critical, albeit locally relevant, traffic and route condition updates.

- ***Enabling the "Personalized" and "Programmable" Vehicle*:** Today, vehicles in both automotive and aviation are built in long design cycles and the models are static in both form and function. Future automobiles will be programmable, with services for the health and performance of both humans and machine. Currently, electronics and software for engine and cabin controls account for over 30% of the cost of an automobile. This figure is expected to increase as vehicles evolve from mechanical to electronic to software-controller, to service-based mobile CPS platforms. By networking vehicles, portals will enable remote diagnosis and re-programmability during the vehicle's life. Future airplanes will increasingly use smart sensor systems to continually monitor the condition of on-board parts, and will use networking to proactively send real-time health diagnostics to ground personnel and enable airlines to remotely change software configurations as well as retrieve flight and cabin updates. Ensuring safe and correct programming is paramount for such road/air vehicles.
- ***On-line Traffic Probing and Real-Time Prediction*:** Delays due to heavy traffic are now costing Americans $78 billion in the form of 4.2 billion lost hours and 2.9 billion gallons of wasted fuel[20]. Air transportation also incurs such costs from flight delays and system inefficiencies. V2X traffic information dissemination can help ameliorate the problem. As vehicles integrate mobile and infrastructure-based networks, the aggregated vehicle speed, position, and direction information will be a significant benefit. New end-to-end network frameworks will support online traffic probing so the detailed status of traffic networks can be analyzed online and timely alerts can be issued. The aggregated data from all communicating vehicles as well as historic data will be used to predict local vehicle trajectory and minimize delay on origin-destination routes.

What Can We Do Well?

Over the past decade, both automotive and aviation have seen a variety of electronic and software-driven technologies incorporated into the engine control, cabin comfort and overall safety and performance of the vehicle. Vehicles use information for three principal purposes: (1) augmentation of mechanical control (e.g., traction control), (2) replacement of mechanical function (e.g., digital X-by-wire), and (3) information-enabled functionality (e.g., traffic-light timing, approach-and-landing trajectory). Vehicle operators use information for vehicle diagnostics, navigation, and traffic updates. Passengers use information for entertainment and comfort. These advances are incorporated in model-based design as the functionality is expected to scale across several dozen microprocessors.

[20] Transportation Research Board, 2007 Urban Mobility Report, http://www.trb.org/news/blurb_detail.asp?id=8172

Community Draft

To date, a primary goal for automotive has been to develop a vehicle unit capable of alerting and assisting the driver by increasing the safety horizon, with sensing, control, and actuation tasks communicated over a common wired CANbus or FlexRay bus. None of these technologies are reliable enough for automatic actuation to assume control of the vehicle. While a technology such as lane departure warning and adaptive cruise control helps keep the driver alert, it does not prevent collisions with nearby vehicles. Single-vehicle-focused automotive technologies have thus stretched the limits of sensing and responding to activity outside the vehicle. In order to enhance safety and comfort significantly, it is necessary for vehicles to communicate their current status and intended actions to their neighbors, thus extending the safety boundaries well beyond the physical dimensions of a vehicle and intentions of one driver.

For aviation, a primary objective is to enable the airplane to share timely, useful, and actionable information about its condition and status to the ground systems, to enhance safety, airspace security, and efficiency. However, the achievable enhancements can be more significant if the information is shared between aircraft, which will improve situational awareness.

Why Can't We Declare Victory?

As vehicle safety, comfort, and navigation are assisted by interaction with infrastructure and other vehicles, the current single-vehicle and vehicle-to-ground operator-focused technologies, modeling, and protocols must be redefined. The inclusion of multiple domains of control, several of which, such as vehicle operator behavior, are not deterministic, increases the complexity of the network-based vehicle control in roadways and airspace. Transportation CPS require dynamics of the vehicle, the environment, operator reaction, and interaction with other vehicles to be considered in every communication and control decision. To provide such guarantees, timeliness and node-independent coordination must be expressed across the network.

The complexity of transportation CPS is beyond current real-time theory and practice in reasoning about large-scale mobile network protocols with multiple dynamic inputs. Our current rate-based and aperiodic event models must be extended to events whose priorities change based on environment. Finally, the operating systems in each vehicle today operate at a fixed operating point. In the transportation CPS context, the dynamic range of environmental variables and vehicle group dynamics is large and requires real-time operating systems (RTOS) to scale beyond multi-modal operation to adaptive, yet predictable and stable, responses.

Significant R&D Challenges

As we extend the automobile's and airplane's services across domains of control and communication, we realize a multi-dimensional problem across time, space, degree of autonomy, number of intra-vehicle sub-systems, number of communicating vehicles, and our current technology capabilities. We categorize the challenges in three classes, namely: real-time theory for spatio-temporal systems, scheduling for distributed and dynamic control, and real-time vehicular network protocols.

1. **Real-time theory for Spatio-Temporal Systems** – Future mobile cyber-physical systems, unlike the current real-time systems, will be spatio-temporal systems of systems that create computational environments, where the computations and their timeliness and hence priorities will be dependent on:
    a. the location of the platform in its environment
    b. the velocity with which the platform is moving
    c. the number of objects in the environment
    d. the velocity vectors of the objects in the environment

Community Draft

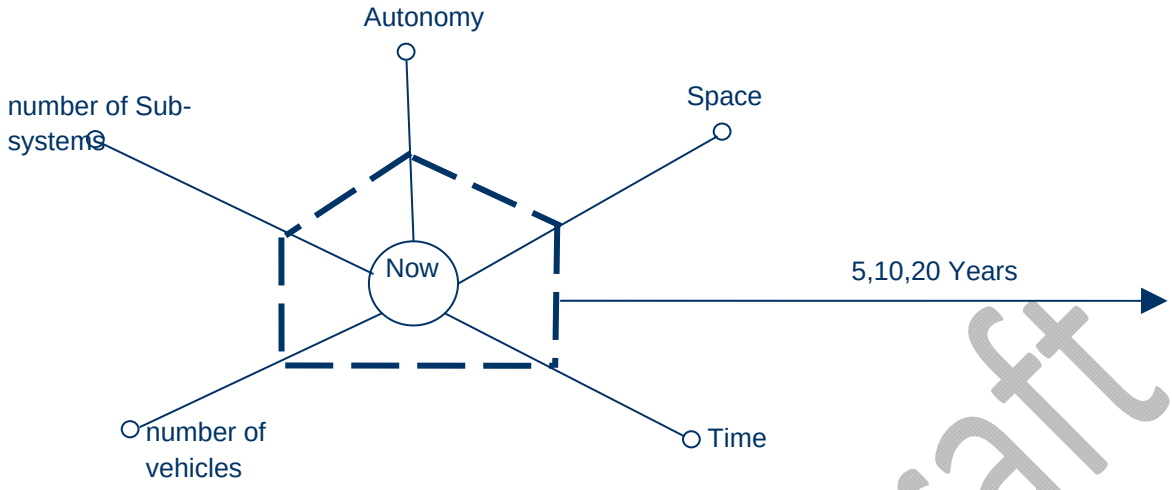

A computational environment consists of multiple levels of heterogeneous systems interacting with each other. These systems are heterogeneous in terms of processing capability, communication interfaces, and even time scales in which events are measured.

- **Managing Mixed-criticality Data and Information Processing –** Electronic Controller Units (ECUs) in vehicles will soon migrate from a large number of low-end microcontrollers to fewer many-core processors. The same processor will need to execute both highly critical and non-critical tasks. Such mixed-criticality systems will require new RTOS and real-time virtualization techniques. A significant effort is required to design and evaluate algorithms for dynamic mapping and prioritization of tasks to processor and network resources given both data-dependent input and time constraints. For groups of vehicles to efficiently adapt to environment conditions and share degrees of state, adaptive RTOS platforms with parametric and programmable control of runtime services are needed. This enables *fluid networks* which, based on network and environment, adapt the coupling between nodes.
- **Scalable SPUR: Security, Privacy, Usability and Reliability –** The SPUR challenge is providing anonymity of user geographic data while maintaining integrity of data for traffic surveillance and safety enhancement. As the in-vehicle informatics evolves to deliver content, there are significant issues with the travelers' privacy and security. New anonymization techniques are required for data mining of transactions and traffic information. The relationship between security and safety must be understood in a more dynamic context and at a larger scale in transport networks. While it is useful to know the location and identifier of a vehicle, how do we ensure that the appropriate levels of privacy are enforced across the different stakeholders with ephemeral relationships? The methodologies for SPUR-related developments must be native to the design of networked and autonomous vehicles.

2. **Scheduling for Distributed and Dynamic Control** – Real-time models and scheduling theory must be advanced in many ways. As vehicles trigger safety alerts based on the local state, the receiver must consider the vehicle state and the world state both before and after the response. To coordinate chassis control with network-based active safety, navigation data, and in the context of vehicles, it is necessary to prioritize tasks

Community Draft

dynamically based on environment and ensure that post-actuation state is safe for the target and its neighbors.

3. **Real-Time Vehicular Network Protocols** – Transportation CPS networks introduce new constraints on network protocols that distinguish them from conventional point-to-point connection and connectionless protocols. The high rate of change in network topology makes path-based and hierarchal address-based route discovery infeasible. As all relevant communication is due to local incidences and actions, we may rely on bounded delay broadcast-based protocols, which are hard to come by. In cases where traffic incidences result in disrupting traffic patterns for extended periods of time, the protocols must ensure the event alert is persistent in the region of interest. Finally, for bulk data exchange and stream-based communication, the protocols must support frequent connectivity changes in the network. The major challenge is to develop a suite of high-confidence protocols, even though the underlying physical network is unreliable and with experiences a high rate of topology change.

Research Strategies and Roadmap

A major effort is required not only to extend existing real-time theories and distributed systems but also to define fundamentally new ideas for communication and control across time and space with changing operation environments. Means for scalable global time synchronization and resource assignment for inter-vehicle collaboration are required for timely message delivery in the broadcast regime. Highlighted below are some of the major research themes that will need to be pursued to realize the networked transportation CPS goals.

- **Open Vehicular Test-beds and Experimental Virtual Systems** – CPS network protocols must be tested across vehicle populations and spatial constraints. We require an architecture for network virtualization for seamless model-based design, testing, validation, and code migration between virtual and real vehicles. As both vehicle types are based on the same assumptions of spatial constraints, navigating conditions, and network and control algorithms, we can balance the complexity of wireless communication, group coordination, and adaptive RTOS with the reality of a slow and gradual roll-out of experimental vehicles. Such an approach integrates systematic design and allows for logical abstractions and network management across large populations of vehicles. For automotive, large-scale, always-on, open-source vehicular network testbeds in association with fractional-rental services such as ZipCar will allow collaborative and competitive transportation CPS design.
- **Adaptive Real-Time Network Protocols –** Because the underlying physical network is unreliable, network protocols must not associate routes with nodes but with primitives such as time, position, driving direction, and average speed. By tightly synchronizing vehicles, we can derive reliable logical abstractions of the network through wireless interference control and spatio-temporal schedules across a range of vehicle densities and spatial constraints.
- **Distributed Virtual Machine (VM) architectures** – A level of abstraction is necessary to extend network-wide run-time services such as scheduling analysis, software attestation, dynamic task migration between nodes, policy negotiation, online fault diagnosis, protocol adaptation, algorithm activation, and data migration. To facilitate this, a virtual machine distributed across tightly coupled nodes in a fleet network will allow for an efficient shared state and coordinated decisions across vehicles in the group.

Community Draft

- **Vehicle Operator Behavior Models and Mobility Models** – The vehicle operator must be modeled in terms of behavior, non-compliance, and reaction times, etc. Vehicular networks have well-defined navigation trajectories on roads and airspaces, and models are needed for better understanding of the impact of spatial/time constraints, time-varying traffic congestion, and environmental factors on macroscopic and V2X interactive microscopic mobility models.
- **Methods and Tools for Certifiable e-Enabling** – Commercial solutions can cost-effectively meet the consumer demands, and are not subject to the stringent regulatory environment of aviation. Recent FAA guidance for next-generation aircraft and on-board wireless technologies suggests that threat assessment, mitigation, risk analysis, and radio interference impacts on on-board systems operation are integral to aircraft certification and airworthiness. Furthermore, new operating and maintenance procedures, such as the airplane's digital certificate management and detection and response to failure of an on-board cyber system must be well-defined and established. Automotive also benefits from certifiable vehicles and practices.

5-Year Milestones

- Develop real-time spatio-temporal models for transportation CPS supporting localized mixtures of application criticalities
- Develop real-time theory for spatio-temporal properties in mid-scale systems, supporting expanded mixtures of application criticalities
- Create comprehensive interaction across vehicle operator/vehicle/environmental models
- Design modeling of intra- and inter-vehicle cross-boundary information exchange and control
- Develop medium-assurance framework and methods for certifiable vehicle-to-ground (V2I)
- Design and develop infotronics for V2I exchange of embedded software and infotainment data for transportation CPS platforms
- Create open Virtual Networked Vehicle Testbed with a deployed network of 50 automobiles with basic standardized Dedicated Short Range Communication (DSRC) for V2X communication, and 1,000,000 virtual road vehicles

10-Year Milestones

- Create high-assurance framework and methods for certifiable V2I e-Enabling
- Develop real-time theory for spatio-temporal properties in large-scale systems, supporting fully dispersed and varying (at runtime) mixtures of application criticalities
- Reference infotronics architecture for transportation CPS with basic sensing, control, and actuation capabilities based on single and one-hop neighbor interactions
- Develop inter-vehicle coordination and control with tightly coupled time synchronization running on a network virtual machine across a set of vehicles
- Create open Virtual Networked Vehicle Testbed for aviation with a deployed network of 50 air vehicles with ADS-B and 802.1x data links, and 10,000 virtual vehicles

15-Year Milestones

- Deploy highly assured infotronics for V2I exchange for transportation CPS
- Create high-assurance frameworks and methods for certifiable V2X e-Enabling

Community Draft

- Design verification and validation safety and reliability technologies for the proposed protocols in a variety of network topologies
- Develop dependable inter-vehicle coordination and control for safety-critical tasks of transportation CPS

Summary Remarks

We find that the applicability of the automotive infotronics concept to aviation's e-Enabling vision exemplifies the commonality among the transportation sectors in networking and information technology advances being made to meet the demands of consumers and society. We highlight the major challenges in realizing road/air vehicles that don't crash, experience minimal traffic congestion delay, are enjoyable and lucrative to operate or travel in, and protect the public well-being. The community proposes new technologies including Spatio-Temporal Real-Time Scheduling, Network Virtualization-enabled Open Testbed, and Embedded Virtual Machines for fluid control across network boundaries. The proposed roadmap is a set of specific aggressive milestones that will drive basic R&D to realize truly intelligent networked vehicular systems in the next 15 years. Limitations in today's technologies create a number of challenges in building future mobile CPS that only a new, bold research initiative will overcome.

Community Draft

Conclusions

Significant breakthroughs in safe, sustainable, affordable and enjoyable transportation systems are on the horizon. We can now plausibly speak of achieving goals such as zero automobile fatalities. Advances in cyber-physical systems technology are one of the key ingredients in doing so. Industry sources at the workshop confirm that next-generation automobiles and aircraft that are steppingstones to these goals will have more than an order-of-magnitude growth in complexity, giving rise to exponentially increasing costs of development and verification using current techniques, to cite one area where CPS technology advances are needed. Similarly, supporting future growth in traffic with greatly reduced time delays will require smart sensing and collaborative vehicle-to-vehicle and human-vehicle interactions that likewise depend on the research agenda proposed here.

The community also understands that new and vibrant methodologies for design and verification are not the only technology elements required to achieve the future transportation vision. Advances in new lightweight materials are also critical. Furthermore, significant public policy hurdles will need to be overcome. The impact of collaborating networks of CPS and vehicle- to-vehicle interactions, including the exchange of vehicle state information, raises new challenges in security and privacy that need further policy and technology research.

The community has noted that CPS represents an important interdisciplinary research focus that has significant impact on K-12, undergraduate, and graduate education. We also see this as a great opportunity to ensure that unconstrained national fellowships such as the NSF and NDSEG programs dedicate resources for this visionary, cross-domain area to attract the very best and the brightest men and women to the science of CPS. Broad technology issues for CPS require understanding of both the technology and the physics of system behaviors and interaction with the real world. This is an acute need for the area of transportation CPS. Future developers of the transportation platform should be educated with rich knowledge of the physics of sensors, properties of actuators, computing architectures, system software, and control theories. Without this interdisciplinary training, advanced research breakthroughs critical to future developments may be forestalled.

The heightened attention to national infrastructure, including transportation infrastructure, due to the global economic crisis is prompting significant investments in its renewal and improvement, including economic stimulus funding targeted at the national infrastructure. CPS research offers the promise of not just renewal but transformation. The coming together of the renewed emphasis on infrastructure and the emergence of CPS research offer a historic opportunity to produce a new transportation infrastructure moving toward “zero-defect” and safety-assured automobiles and airplanes with a dramatically reduced energy footprint, and a transportation infrastructure (air and ground) that can safely and sustainably support not just current but future traffic loads.

Community Draft

APPENDIX A: Related References

Infotronics and Infotainment

1. United States Department of Transportation, Federal Highway Administration, Highway Statistics, 2006.
2. Bureau of Transportation Statistics, Airline On-Time Performance Data, http://www.transtats.bts.gov
3. Center for Urban Transportation Research, Public Transit in America: Findings from the 1995 Nationwide Personal Transportation Survey, (www.cutr.eng.usf.edu), Table 4-13.
4. The 1996 Convergence Council comprising the research and engineering leaders of the Detroit 3 automobile manufacturers and the supply base.
5. Boeing Frontiers, vol. 02, no. 04, August 2003. http://www.boeing.com/news/frontiers/archive/2003/august/i_ca1.html

Mixed-Criticality

1. A Research Agenda for Mixed-Criticality Systems. Public Release report authored by Boeing, Lockheed Martin, Northrop Grumman, and the Air Force Research Laboratory, Approved on 07 April 2009 by Wright Patterson Air Force Base Public Affairs, PA case number: 88ABW-2009-1383. http://www.cse.wustl.edu/~cdgill/CPSWEEK09_MCAR/RBO-09130%20Joint%20MCAR%20White%20Paper%20PA%20approved.pdf
2. The ScanEagle low-cost, long-endurance autonomous unmanned vehicle. http://www.boeing.com/defense-space/military/scaneagle/
3. Tartan Racing: A Multi-Modal Approach to the DARPA Urban Challenge. April 2007. www.darpa.mil/GRANDCHALLENGE/TechPapers/Tartan_Racing.pdf
6. ARINC 653. Airlines Electronic Engineering Committee. Avionics application software standard interface part 1 - required services. Technical Report 653P1-2, Aeronautical Radio, Inc., March 2006.
7. CAN bus. http://www.canbus.us/
8. RTCS SC-167 and EUROCAEWG-12. Software Considerations in Airborne Systems and Equipment Certification. Number RTCS/DO-178B. RTCA, Inc., Washington, D.C., December 1992.
9. MILS (Multiple Independent Levels of Security). http://www.ois.com/Products/MILS-Technical-Primer.html

Platforms

1. Victoria Transport Policy Institute, "Introduction to multi-modal transportation planning: principles and practices", http://www.vtpi.org/multimodal_planning.pdf, November 18, 2008
2. National Institute of Standard and Technology, "4D/RCS: A reference model architecture for unmanned vehicle systems, version 2.0", http://www.isd.mel.nist.gov/projects/rcs/ref_model/coverPage3.htm, August 2002

Draft

3. Transportation Research Board of the National Academies, "Railroad operational safety, status and research needs", in Transportation Research Circular E-C085, January 2006
4. Carnegie Mellon Software Engineering Institute, "Ultra-large-scale systems: The software challenge of the Future", July 2006
5. Federal Aviation Administration, "Introduction to TCAS II version 7", November 2000

Autonomy and Control

1. FlexRay Consortia, "FlexRay communications system protocol specification, version 2.1", December 2005
2. G. Pappas, C. Tomlin, J. Lygeros, D. Godbole, and S. Sastry, "A next generation architecture for air traffic management systems," *Proc. of IEEE conference on Decision and Control*, pp. 2405–2410, 1997.
3. Edward A. Lee, "The Problem with Threads," in *IEEE Computer*, 39(5):33-42, May 2006.
4. S. Thrun, et al. "Stanley, the robot that won the DARPA Grand Challenge." *Journal of Robotic Systems*, Special Issue on the DARPA Grand Challenge, Part 2, 23(9):661-692, 2006.
5. Christopher Urmson, et al. "Autonomous driving in urban environments: Boss and the Urban Challenge" *Journal of Field Robotics*, Special Issue on the 2007 DARPA Urban Challenge, Part I, 25(8):425-466, 2008.
6. Varaiya, P., "Smart cars on smart roads: problems of control," *IEEE Trans. Automatic Control*, 38(2):195-207, Feb 1993.
7. Spooner, J.T.; Passino, K.M., "Fault-tolerant control for automated highway systems," *Vehicular Technology, IEEE Transactions on*, vol.46, no.3, pp.770-785, Aug 1997.
8. Giridhar, A.; Kumar, P.R., "Scheduling Automated Traffic on a Network of Roads," *Vehicular Technology, IEEE Transactions on*, vol.55, no.5, pp.1467-1474, Sept. 2006.
9. Waydo, S.; Hauser, J.; Bailey, R.; Klavins, E.; Murray, R.M., "UAV as a Reliable Wingman: A Flight Demonstration," *IEEE Transactions on Control Systems Technology*, 15(4):680-688, July 2007.
10. D. Jackson, M. Thomas, and L. I. Millet, Eds., "Software for Dependable Systems: Sufficient Evidence?", ser. *Committee on Certifiably Dependable Software Systems*. Washington, DC: National Academy Press, 2007.
11. F. Dana, V. Gupta, J. Hespanha, R. M. Murray, and B. Hassibi, "Estimation over communication networks: Performance bounds and achievability results," in *American Control Conference*, New York, Jul. 2007, pp. 3450–3455.
12. C. Tomlin, J. Hansman, J. Springle, "Report on the workshop for aviation software systems for the second century of flight: Design for certifiably Dependable Systems (HCSS-AS)," October 206.

Community Draft

APPENDIX B: The case of Roadway Infrastructure CPS[21]

Contributors: Raja Sengupta and Mr. James Misener, PATH Project, University of California, Berkeley

Introduction

Currently intelligent roadway and vehicular systems are built in isolation and don't typically share data. For example, when a city designs an intelligent signal control system it builds and installs sensors, actuators, computing, and communication sub-systems without regard to the concurrent design and installation of a city bus route that includes an LED messaging and electronic communication system at bus stops driven by GPS[22]. For example, currently the GPS data from buses usually cannot communicate with the signal control system. Computing technology deployment typically occurs in isolated system or sub-systems, not in an interacting system-of-systems infrastructure.

The increasing need for transportation safety, security, efficient use of energy and traveler comfort and convenience demands a futuristic roadway where the proliferation of intelligence is considerably more efficient. We envision an embedded CPS roadway infrastructure, that can be used by anyone to create value in any manner consistent with its lawfully constituted business model. This new infrastructure model approach would be economically feasible if it were designed to require new service builders to only pay for their system specific services and share the costs to design and build the intelligent services across the infrastructure. For example, if the roadway offers precise positioning, WiFi, and street corner computing, a small entrepreneur might easily make the blind or elderly safer through our "Watch Out For Me" concept[23], i.e., the person's smartphone could multicast "watch out for me" as soon as he steps of the curb to any oncoming car[24]. The example illustrates why the infrastructure should be cyber-physical. Many valuable services of the future, requiring low marginal cost connectivity, such as MySpace or Facebook, can be cost effectively developed employing interconnectivity, in real-time, between groups of embedded sensors and actuators.

Over the last five years, we have created the Vehicle Infrastructure Integration (VII) California test bed, as a resource for the scientific community engaged in research on the smart roadway[25]. VII California is a technology infrastructure, scientific staff, and a network of partnerships enabling the rapid prototyping and evaluation of new safety, mobility, or green transportation services in a real environment - 40 miles of roadway with 7 signalized intersections on a major freeway and major arterial. The test bed is equipped with multiple wireless communication

[21] This project is described as an example and by no means should be construed as endorsing the effort by Berkeley.

[22] http://www.nextbus.com

[23] http://path.berkeley.edu/PATH_Downloads/To-Send/WC2008/WOCO2008-NYC.mpg

[24] This could avoid tragedies like http://www.dailycal.org/article/104588/berkeley_kindergarten_student_killed_by_truck

[25] www.**viicalifornia**.org/

Community Draft

services, positioning services, accessible computing in street corner signal cabinets, and data feeds from local buses, cars run by silicon valley automotive research laboratory partners, and traffic measurements sensors embedded in the pavement by the California Department of Transportation. We seek to make this test bed a tool for the advancement of a roadway infrastructure CPS.

Community Draft

APPENDIX C: Acronyms

ABS – Antilock Brake System

ADS-B – Asynchronous Dependent Surveillance Broadcast

AFRL – Air Force Research Laboratory

ARINC – Aeronautical Radio Inc.

ATC – Automatic Train Control

ATM – Air Traffic Management

CAN – Controller Area Network

CDMA – Code Division Multiple Access

CerTA FCS – Certification Techniques for Advanced Flight Critical Systems

CPS – Cyber-Physical System

CPU – Central Processing Unit

CS – Computer Science

DARPA – Defense Advanced Research Projects Agency

DHS – Department of Homeland Security

DoD – Department of Defense

DoT – Department of Transportation

DSML – Domain Specific Modeling Language

DSRC – Dedicated Short Range Communication

ECU – Electronic Control Unit

EPA – Environmental Protection Agency

EU – European Union

FAA – Federal Aviation Administration

FDA – Food and Drug Administration

FIRST – For Inspiration and Recognition of Science and Technology

FPGA – Field Programmable Gate Array

GPRS – General Packet Radio Service

GPS – Global Positioning System

HCCPS – High Confidence Cyber-Physical Systems

HCSS CG – High Confidence Systems and Software Coordinating Group

IEEE - Institute of Electrical and Electronics Engineers

I/O – Input / Output

ITR – Information Technology Research

IVHM – Integrated Vehicle Health Management

K-12 – Kindergarten through 12th grade

LIN – Local Inter-connect Network

MILS – Multiple Independent Levels of Security

NextGen – Next-Generation Air Transportation System

MoBIES – Model-Based Integration of Embedded Systems

NASA – National Aeronautics and Space Administration

NDSEG – National Defense Science and Engineering Graduate Fellowship

NHTSA - National Highway Transportation Security Agency

NIST – National Institute of Standards and Technology

NITRD – Networking and Information Technology Research and Development

NRC – Nuclear Regulatory Commission

NSF – National Science Foundation

NSA - National Security Agency

NTSB – National Transportation Safety Board

Community Draft

PCAST – President's Council of Advisors on Science and Technology

PCES – Program Composition for Embedded Systems

PLOC – Probability of Loss of Control

QoS – Quality of Service

R&D – Research and Development

RFID – Radio Frequency Identification

RTOS – Real-Time Operating Systems

SEC – Software Enabled Control

SLOC – Source Lines of Code

SPUR – Security Privacy Usability and Reliability

SWaP – Size, Weight and Power

TCAS – Traffic Collision Avoidance System

UAS – Unmanned Aircraft Systems

UAV – Unmanned Aerial Vehicles

UID – Unique Identification

USB - universal serial bus

USCAR – United States Council for Automotive Research

V2I – Vehicle-to-Infrastructure

V2V – Vehicle-to-Vehicle

V2X – Vehicle to something outside the vehicle

VM – Virtual Machines

VMS – Vehicle Management System

V&V - Verification and Validation

WCET – Worse-Case Execution Time

Community Draft

APPENDIX D: 2008 HCTCPS Workshop Roster

Program Co-Chairs
Radha Poovendran, Univ. of Washington
Raj Rajkumar, Carnegie Mellon University

Program Committee Members
Robert Baillargeon, General Motors
Alex Bayen, UC Berkeley
David Corman, Boeing
Edward Griffor, Chrysler
Bruce Krogh, CMU
Srikanta Kumar, BAE Systems
Nandish Mattikalli, BAE Systems
William Milam, Ford
James Paunicka, Boeing
K. Venkatesh Prasad, Ford
Krishna Sampigethaya, Boeing
Jonathan Sprinkle, U of Arizona
John Stankovic, University of Virginia
Janos Sztipanovits, Vanderbilt

Government Points of Contact
Ray Bortner, AFRL
Michael Branicky, NSF
Magdy El-Sibaie, FRA
Helen Gill, NSF
David Homan, AFRL
David Harris, AFRL
Jonathan Hoffman, AFRL
Paul Jones, FDA
Frankie King, NCO/NITRD
Alan Kushner, NTSB
Brad Martin, NSA
Scott Midkiff, NSF
Nelson Miller, FAA
Paul Miner, NASA
Russell Urzi, AFRL
Albert Wavering, NIST

Workshop Sponsors
Air Force Research Laboratory (AFRL)
National Science Foundation (NSF)
National Security Agency (NSA)

Workshop Participants
Ashish Agarwal, Boston University
Panos Antsaklis, University of Notre Dame
Hakan Aydin, George Mason University
Hamsa Balakrishnan, MIT
Jim Barhorst, Boeing Phantom Works
Robert Baillargeon, General Motors
Joseph Bergmann, The Open Group
Pam Binn, Honeywell
Sushil Birla, Nuclear Regulatory Commission
Todd Belote, Lockheed Martin
Robert Benito, MITRE
Michael Branicky, NSF
Linda Bushnell, UW
Qing Cao, GE Global Research
Timothy Chang, NJIT
Xiuzhen Cheng, George Washington University
Mo-Yuen Chow, North Carolina State Univ.
Rance Cleaveland, University of Maryland
Eric G. Cooper, NASA LaRC
David Corman, Boeing
Werner Damm, EICOSE
Rance DeLong, LynuxWorks
Alex Doboli, SUNY at Stony Brook
Ian Downes, Stanford University
David H Du, Univ. of Minnesota
Yaser Fallah, UC Berkeley
Emilio Frazzoli, MIT
Martin Fritzsche, Daimler AG
David Garlan, CMU
Soheil Ghiasi, UC Davis
Patrick H. Goertzen Boeing Phantom Works
Chris Greer, NCO/NITRD
Edward Griffor, Chrysler LLC
Sandeep Gupta, Arizona State University
David Harris, AFRL
Gene Hayman, Boeing/NextGen
Wenbo He, Univ. of New Mexico
Jonathan Hoffman, AFRL

Community Draft

David Homan, AFRL
Chuck Howell, MITRE
Paul Jones, FDA
James Kirby, NRL
Bruce Kim, Univ. of Alabama
Frankie King, NCO/NITRD
Basil Krikeles, BAE Systems
Bruce H. Krogh, CMU
Alan Kushner, NTSB
Elizabeth Latronico, Robert BOSCH
Insup Lee, Univ. of Pennsylvania
Mingyan Li, Boeing Research & Tech.
Scott A Lintelman, Boeing
Steve Liu, Texas A&M University
Xue Liu, McGill University
Jeffrey Maddalon, NASA
Rahul Mangharam, Univ. of Pennsylvania
Nandish Mattikalli, BAE Systems
Scott Midkiff, NSF
William Milam, Ford Motor Company
Mary Ellen Miller, Raytheon Company
Nelson Miller, FAA
Steven Miller, Rockwell Collins
Paul Miner, NASA Langley Research Center
Aloysius Mok,UT Austin
Daniel Mosse, Univ. of Pittsburgh
Cesar Munoz, National Institute of Aerospace
Kamesh Namuduri, Univ. Of North Texas
Natasha A. Neogi, UIUC
Dionisio de Niz, CMU
John Osterholz, BAE Systems
James L. Paunicka, Boeing Phantom Works
Andre Platzer, CMU
Radha Poovendran, UW
K. Venkatesh Prasad, Ford Motor Company
Calton Pu, Georgia Tech
Raj Rajkumar, CMU
Harini Ramaprasad, Southern Illinois University
Jim Ritcey, UW
Glenn Roberts, MITRE
Joe Salvo, GE Global Research
Krishna Sampigethaya, Boeing R&T
Prakash Sarathy, Northrop Grumman
Shankar Sastry, UC Berkeley
William W. Schoening, Boeing
John Scoredos, Northrop Grumman
Lui Sha, UIUC
Weisong Shi, Wayne State University
Arun Somani, Iowa State University
Cheryl Souders, FAA
Jonathan Sprinkle, Univ. of Arizona
Peter Stanfill, Lockheed Martin Aero
Douglas Stuart,Boeing Phantom Works
Greg Sullivan, BAE Systems AIT
Janos Sztipanovits, Vanderbilt University
Yosef Gavriel Tirat-Gefen, NCO/NITRD
Ashish Tiwari, SRI International
Craig Treece, Lockheed Martin
Frank Vahid, UC Riverside
Chris Walter, WW Technology
Cliff Wang, ARO
Shige Wang, GM R&D
Bennett C Watson, Lockheed Martin
Wayne Wolf, Georgia Tech
Daniel B. Work, UC Berkley
Hongwei Zhang, Wayne State University
Dieter Zöbel, University of Koblenz-Landau

Community Draft

APPENDIX E: Community Acknowledgements

The authors of this community report extend their sincere appreciation to the many academic and industry experts in the high-confidence IT R&D community who planned and participated in the 2008 High Confidence Transportation CPS workshop. Through their intensive intellectual engagement over time, these contributors articulated and shaped the theoretical, scientific, and technical framework for the research agenda discussed in our report.

Dr.Venkatesh Prasad from Ford Motor Company, the University of Washington and The Boeing Company are responsible for the idea and the seeds of this effort. Dr. Prasad was invited by the Dean of College of Engineering, Matt O'Donnell, who had recently moved from the University of Michigan. A Boeing-UW-Ford workshop on aviation and automotive systems in Seattle was planned in the summer of 2008. When NSF was approached for possible sponsorship, it became evident that the topic was of broader interest to the nation which led to a national workshop under the sponsorship of the NITRD Program.

We are thankful to Drs. Helen Gill, Ty Znati, Jeannette Wing, Scott Midkiff, and Michael Branicky from the NSF and Dr. Chris Greer from the National Coordinating Office (NCO) for their valuable time and strong support for our efforts. Apart from the NSF, significant funding help and logistical help were provided by Dr. Russ Urzi under the sponsorship of the AFRL, without which the program could not have been hosted in Washington, D.C and for that we are deeply grateful.

The workshop in D.C. was well attended with more than 115 participants. It included talks from Dean Shankar Sastry of UC-Berkeley, Dr. Amy Prichett of NASA, Dr. Chris Greer of the NCO, Mr. Don Winter of Boeing Company, Dr. Werner Damm from European Mission, Prof. Janos Stipanovitz, and Prof. Michael Branicky which prompted detailed discussions and the formation of working groups. Dr. David Homan from AFRL presented his view of the research challenges that helped the academic audience integrate the research needs from dual perspective. We also want to thank Professors Richard Murray and John Stankovic for their comments. Very special thanks to Prof. Claire Tomlin who participated in our panel via phone while also attending to parental duties.

We would like to acknowledge the very generous support of the Boeing Company for this workshop. Apart from committing their lead researchers Drs. David Corman and James Paunicka, they also provided the Washington DC Boeing facility and the St. Louis facility in 2008 and 2009 respectively for writing workshops. Boeing also generously provided their resources for the web and phone services for group meetings and provided financial support for the meeting in D.C. We thank the leadership of Boeing for the unwavering commitment to the success of this effort.

Last but not the least, we thank the NCO/NITRD staff, Ms. Frankie King who played an instrumental role in the development of our report, including assisting in research, editing and rewriting text, and publishing the final document.

Community Draft