Cyber-Physical Challenges in Transportation System Design

Robert Cartwright[1], Albert Cheng[2], Paul Hudak[3], Marcia O'Malley[4], Walid Taha[1]
[1]Department of Computer Science, Rice University
[2]Department of Computer Science, University of Houston
[3]Department of Computer Science, Yale University
[4]Department of Mechanical Engineering, Rice University

Transportation systems have a direct impact on the nations productivity, environment, and energy consumption. To design novel, higher efficiency transportation systems we must overcome fundamental technical challenges relating to the cyber-physical nature of modern transportation systems. In a *cyber-physical system (CPS)*, discrete computing components control or monitor continuous physical components in real-time. The transportation domain provides an ample supply of examples of such systems. We, the authors, are working together to develop a novel CPS modeling environment called Acumen. Our goal is to accelerate the development of highly reliable software and hardware solutions for CPS domains. Our past research efforts have addressed numerous aspects of such systems, including robotics, haptics, and real-time and embedded systems. In the transportation domain it is easy to see how leadership in the development of cyber-physical systems not only helps maintain and improve US economic competitiveness, but contributes to protecting and extending human life.

The key lesson learned from our collective experience is that the structure of the physical model used to understand a particular problem has a dramatic effect on how solutions to this problem are developed.

We are particularly interested in the development of new and innovative transportation systems that radically depart from the traditional design space. Success in such tasks critically depends on rapid prototyping and testing of new control algorithms. Whether the application is space exploration, freight transportation, or steer-by-wire systems for personal automobiles, developing new task-specific control algorithms is costly, time-consuming, and can have a significant impact on overall operating cost.

The fundamental limitations of current technology for developing transportation systems include:

Fidelity: Although it is well known that models for cyber-physical systems must involve digital (discrete) and physical (continuous) elements, there is a conspicuous absence of software tools and methods for establishing the correctness and validity of such cyber-physical models. For example, while individual components of automotive systems are increasingly combining physical and digital subsystems, existing tools generally force these components to be treated separately. The weak coupling between the tools only exacerbates the fact that these tools generally assume that the artifacts that are modeling live either fully in the discrete domain or fully in the continuous domain. Regrettably, there is generally no guarantee that the output from a tool is correct except that "it looks right". Given the safety-critical nature of transportation systems, this situation is far from satisfactory.

Reuse: Due to the enormous diversity of the technologies involved, transportation systems are fertile grounds for unique, creative innovations. Vivid examples can be seen in new technologies for autonomous and for high efficiency vehicles, where individual successful designs can be radically different. The uniqueness of such innovations makes developing reusable canned solutions impossible. Successful models are often ad hoc and depend greatly on the details of the system being designed. As a result, the timely development of new systems hinges on the ability to quickly model and validate the physical behavior of both the new technology and its potential environments. Existing methods and tools for building hardware and software solutions provide neither the integration nor the abstractions needed to build reusable cyber-physical components.

Cost: A transportation system that is experimental or highly specialized cannot benefit from the economics of mass production. As a result, the design and validation cost for such systems can be extremely high. We need to develop tools and methods that allow new models to be prototyped quickly and efficiently, leveraging when possible similar features from other models.

The most important research challenges are:

Reliable methods for modeling cyber-physical systems: While the difficulty in modeling cyber-physical systems comes from the diversity of these systems, the most promising approach to mitigating this problem is to develop very expressive, general modeling languages. In principle, mathematics can provide this kind of common language. Unfortunately, there are currently two problems with this ideal. First, because of round-off errors, existing programming languages rarely provide satisfactory implementations of either arithmetic or other mathematical concepts to the point where they would be acceptable for modeling purposes. In fact, even most symbolic mathematical systems fail to provide the semantic correctness guarantees that would be needed for mathematical modeling. At the same time, for specialized domains such as mechanical design, mathematics does not always provide the most standard or most directly usable notation.

Effective methods for analysis, simulation and validation of models: Once a model has been built, systematic design requires tools and methods to transform into an efficient machine-executable form. State-the-art compilation techniques can be used to perform this translation, but new technologies such as multi-core architectures and reconfigurable computing architectures can only be exploited effectively if we take into account the nature of the class of models being executed. At the same time, tools are needed for both the manual and the automated analysis of various properties of these models. For example, given a model of an automobile, can we verify that an automatic transmission controller does not oscillate when the car is driving on a fixed-gradient road? Can we verify that a cruise control system can stay within a certain error of the desired speed? Can we verify that it expends energy only within reasonable bounds of what is expected to be needed under such circumstance?

Methods and techniques for guaranteeing real-time properties: Cyber-physical systems enable their users to monitor and manipulate the physical world. Many such systems involve real-time constraints. While recent advances in the theory of hybrid systems provide a foundation for studying cyber-physical systems, these techniques must be adapted and extended to handle the modeling concepts that naturally arise in physics, and integrated with a wide range of analysis techniques to verify real-time properties of cyber-physical models.

The most promising innovations and abstractions are:

Integrated Methods for Design and Verification of CPSs: We need to find clear, well-defined methods for designing systems that contain hybrid components. Such methods will generally always be dependent on the technology available for capturing and validating designs. Thus, current methodologies are closely tied to existing tools, which are either traditional programming languages which are highly mismatched for modeling physical systems or mainstream integrated workbenches such as MATLAB which are in essence customized interfaces to specialized numerical libraries and threfore provide no guarantees about the fidelity of the models used. We need methodologies that are built around suitable formal verification tools that provide the basic tools for reasoning about real numbers.

Integrated Tools for Design and Verification of CPSs: Developing correct physical models is very similar to developing correct programs. It is extremely hard. Correct equations are no easier to write than correct programs. Therefore, just as with software, testing is of paramount importance. Surprisingly, we find that testing physical models with the tools available today can be particularly challenging. For example, very few commercial tools provide us with the ability to simulate a physical system in a manner that *guaranteers* that the results are correct to any stated degree of accuracy. Similarly, whereas there are many tools available for visualizing data, very few tools provide powerful tools for visualizing *designs.* Better use and integration of both textual and visual languages can enable dramatic improvements in helping engineers and scientists understand their designs.

Important milestones for the next 5 to 10 years will be:

Cyber-physical modeling environments that provide fidelity guarantees: Simulation systems need to tell us how accurate their results are.

Synthesis techniques that provide safety and real-time guarantees: Mappings from models that have been tested in simulators need to come with guarantees about how the software and hardware components will operate when embedded in the real-world.

Order of magnitude improvement in productivity for developers of cyber-physical systems: The ultimate measure of success of new technologies for cyber-physical computing will be in improved productivity in the development of reliable cyber-physical innovations and products.

Biographies:

Robert Cartwright: Professor Cartwright's principal research interests are programming language design and implementation, program specification, program testing and analysis, and software engineering.

Albert Cheng: Professor Albert Cheng directs the Real-Time Systems Laboratory at the University of Houston. His interest is in the analysis, verification, and scheduling of real-time systems. He is the author of the textbook "Real-Time Systems" published by Wiley.

Paul Hudak: Professor Hudak leads the Yale Haskell Group, which studies the use of functional programming, formal methods, and domain-specific languages to solve problems in diverse areas such as robotics, control systems, animation, and computer music.

Marcia O'Malley: Professor O'Malley directs the Mechatronics and Haptic Interfaces Lab at Rice University. Her research interests focus on the study of physical human-robot interaction, and the design and control of haptic (force-feedback) devices.

Walid Taha: Professor Taha leads the Resource Aware Programming research group at Rice University, where he advises two postdoctoral associates, three PhD students, and several undergraduate students. His interest is in developing and applying programming language techniques that can improve the productivity of software and hardware developers.