Abstractions for Cyber-Physical Systems Development: An International Opportunity

Roger D. Chamberlain
Chandler Ahrens
Chris Gill

Roger D. Chamberlain, Chandler Ahrens, Chris Gill, “Abstractions for Cyber-Physical Systems Development: An International Opportunity,” presented at *Visioning Workshop for International Networks to Advance CPS Research, Development, and Education Worldwide*, April 2018, https://cps-vo.org/node/48624

Dept. of Computer Science and Engineering
School of Engineering and Applied Science

College of Architecture
Sam Fox School of Design & Visual Arts

Washington University in St. Louis

Abstractions for Cyber-Physical Systems Development: An International Opportunity

Roger D. Chamberlain*, Chandler Ahrens†, and Chris Gill*

*Dept. of Computer Science and Engineering, School of Engineering and Applied Science

†College of Architecture, Sam Fox School of Design & Visual Arts

Washington University in St. Louis, St. Louis, MO, USA

Email: {roger,cahrens,cdgill}@wustl.edu

Layered *system architectures*[1] have proven effective in providing standardized interfaces to support portability and reuse of hardware and software, while allowing new innovations and abstractions to enrich system capabilities. Even when systems don't adhere strictly to standards such as the OSI [6] networking model or the POSIX [4] operating systems interface model, system designers, developers, and users still benefit from the structure those models provide.

This occurs largely because standardized interfaces establish clear boundaries of responsibility on which other layers may rely, while allowing an essential "permission to tinker" with various implementations between those boundaries. For cyber-physical systems, whose semantics include timing and physical properties not considered in previous cyber-only system architectures, there is significantly less experience with what interfaces, abstractions, and even broad system layers are truly common (and so perhaps could be standardized) versus which other aspects are more likely to diverge between systems, and so should be free to do so. Here, we articulate the case for a broader investigation of appropriate abstractions to support both new and existing cyber-physical systems and applications.

We also examine the potential for new constellations of abstractions to emerge from this investigation, as has also happened numerous times previously in the history of system architectures. For example, operating systems emerged to allow different computations to be run on a single computer, and then evolved to support concurrent use by multiple applications at once. Atop standardized operating systems, *concurrency platforms* [5] emerged to allow programs with internal parallelism to exploit multicore capabilities of each host. Such constellations of abstractions (and the platform hardware and software that embodies them) serve to isolate system designers, developers, and users further from lower-level design decisions and details involving compilers, run-time environments, operating systems, and underlying hardware, allowing them to focus instead on the dominant abstractions relevant to the applications they are trying to realize.

We contend that similar concepts are beneficial to the field of cyber-physical systems, but that insufficient attention has been focused so far on *consolidating experience, abstractions, mechanisms, and policies across the wide range of cyber-physical systems that are being developed today, into cohesive, common, and reusable platforms*. Each such *CPS platform* should define a cohesive set of *customized* abstractions under a common system architecture that provides tangible benefits to CPS developers. As with the concurrency platform example noted above, the semantics underlying each platform's particular set of abstractions must be supported fully by compilers, run-time systems, operating systems, and execution hardware. These *customized* abstractions then would make CPS applications easier to build, and also more robust, secure, etc., by avoiding semantic mis-matches between abstractions that may serve well for one application but do not address fundamental semantics of another application.

We articulate several notions of what might comprise examples of such CPS platforms in specific CPS application domains. In doing so we hope to illustrate what is needed and motivate broad, collaborative research in this direction.

Cyber-physical systems often entail multiple levels of control, with a lower level that uses classical control theory abstractions (e.g., position an object at a particular location, or maintain a temperature at a specified set point) and a higher level that is presumed to involve different abstractions (e.g., to optimize an objective while still meeting system constraints).

While not always the case, it is frequently true that such a separation between lower-level control concerns and higher-level decisions is defined in terms of the quantitative vs. qualitative nature of the models employed and control algorithms that are used. At the lower level, it is often clear how to identify the measured inputs (from sensors), the control outputs (to actuators), and algorithms needed in the cyber portion of a cyber-physical system. At the higher levels of decision making, however, it may be necessary to express objective functions in qualitative terms even if decisions also involve whatever quantitative information may be available to inform them. We illustrate this further with a diverse pair of applications: redirecting sunlight to improve conditions inside a building in one, and integrating physical and simulated portions of an earthquake engineering experiment in the other.

A. Catoptric Systems

Figure 1 is an image of a prototype catoptric (mirror) surface (called AMP) that was designed, fabricated, and in-

[1] We use the term *system architecture* to denote computer hardware/software architecture, distinguishing it from architecture for the built world around us.

Fig. 1. *AMP* prototype, TRex building, St. Louis, MO

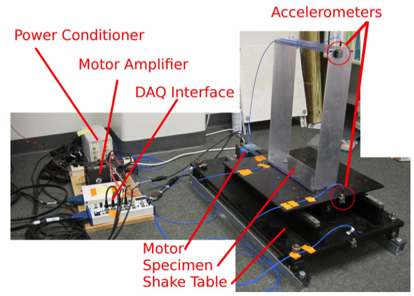

Fig. 2. RTHS experiment at Washington University, St. Louis, MO

stalled through an undergraduate architecture studio taught by C. Ahrens. The installation redirects light from gable ends of an existing building into the darker recesses of the atrium to create better natural lighting where it is desired. In the next generation of this system, which is currently under construction, the mirrors are under active, 2-axis, microprocessor-based control and therefore can be pointed in different directions dynamically as desired over time.

Using ray tracing algorithms and straightforward low-level control mechanisms, we can position a set of mirrors to direct light where it is desired. The high-level decisions revolve around the question of "where do we want the light to be?" This question becomes even more interesting if we expand the higher level options to include varying the intensity of daylight that is directed into an area according to the area's use (e.g., lower to improve contrast when computer screens are being used, or higher when people are reading paper materials), or the ability to direct incoming sunlight that is not needed for another purpose into a heat exchanger that can harvest energy for the building's HVAC system.

Given a bounded resource (currently available sunlight), we now are charged with balancing the relative benefits of natural sunlight illumination within a physical space vs. the harvesting of thermal energy which can potentially reduce operating costs for that same physical space. While this example is clearly an interesting cyber-physical system (and advances the state of the art in the specific area of smart and connected communities), it is less clear which of its mechanisms and services (1) could be supported by a generic CPS platform that would facilitate the efficient implementation of such a system, or (2) could be re-used by other CPS applications and if so to what extent. Before attempting to answer this question, however, we turn to our second application example.

B. Real-Time Hybrid Simulation of Civil Structures

Our second cyber-physical application example comes from the domain of *real-time hybrid structural* (RTHS) [1] experiments used in structural and earthquake engineering. In these experiments, a physical specimen is connected to a low-level control loop via sensors that can measure the position and/or acceleration of the specimen continuously, communicate those measurements frequently as digital readings through a DAQ board, and in turn receive and apply commands sent by the low-level controller through the DAQ board to actuators that are also connected to the specimen. Low-level control design includes both the dynamics of the specimen and compensation for the dynamics of the actuators to ensure that appropriate forces are applied to the specimen over time. The sensor measurements are used as inputs to a simulation (e.g., using a finite element model) that computes the behavior of the overall structure, including the next set of forces to apply to the specimen at the next time step. Within each iteration, both the low-level control loop and the simulation must complete within a millisecond or less to avoid losing high-frequency vibrations or other behaviors of the specimen.

Figure 2 shows a scale-model version of such an experiment, at Washington University in St. Louis. Full-scale experiments require hydraulic actuators and other capabilities found in facilities such as the Bowen Laboratory at Purdue University, and our experiences working with both versions has shown that scale-model experiments offer a useful setting for designing, developing, and debugging cyber-physical experiments before moving them to the full-scale setting.

Higher-level objectives include the ability to substitute simulated versus physical versions of different parts of a structure [3], and through advances in parallel real-time concurrency platforms to trade-off computational demand and time scales [2] allowing resources to be concentrated for portions of the structure that are most significant.

C. Discussion

Clearly, there are at least some common abstractions for both of these examples. For example, low-level control abstractions (e.g., PID controllers) are needed for both of them and for such well-understood abstractions providing them within a common software framework (with a suitable parameterization) should be relatively straightforward even if they entail semantic guarantees (e.g., timing constraints, or objective functions or fail-safe conditions imposed by the application).

However, achieving this in a reusable manner becomes more complicated at higher levels of abstraction, particularly when multiple low-level behaviors may interact. For example, pointing too many mirrors to one location may lead to overheating, and allocating too many cores to a particular RTHS component may cause others to miss deadlines. In addition to ensuring safety properties and optimizing application objectives, these platforms also may need to ensure resilience, privacy, and other properties that further complicate the design of appropriate abstractions.

D. *Vision for International Outreach and Collaboration*

Cyber-physical systems are an active area of research internationally, with many research groups throughout the Americas, Europe, and Asia rapidly advancing both the state of the art and the state of the practice. The vision we propose for consolidating and exploiting common abstractions for cyber-physical systems is necessarily international in nature, both to leverage the breadth and depth of CPS research world-wide, and to open up a larger conversation about how social, cultural, economic, environmental, and other aspects both inform and shape cyber-physical systems in different contexts.

With regard to the two CPS application domains we have discussed here, C. Ahrens is already collaborating with a colleague at Technion Institute in Israel, who runs an architecture/robotics lab that includes a 7-axis robot arm (a significantly more intricate device for manipulating light and other media). C. Gill's collaborators at Purdue University work with structural and mechanical engineers throughout the world, and real-time hybrid simulation is an active research topic at a number of universities across Europe and the Americas. Groups who are already looking at how light and other phenomena can be manipulated in the built environment, or how real-time hybrid simulation can be used to support new kinds of experiments that previously could not be attempted, would be natural targets of outreach, to build on existing directions and establish previously unforseen ones based on the unique perspectives of each participant.

One potential approach to facilitating international collaborations would be for NSF (or another agency) to fund explicit connections between research groups in the US and the already established broadly international research collaborations in the EU funded through the Horizon 2020 program or its successor.

Acknowledgments

Work described here has been supported in part by NSF through grants CNS-0931693, CNS-1527510, and CNS-1329861, and the Washington Univ. International Center for Energy, Environment and Sustainability (InCEES).

References

[1] S. Dyke, B. Stojadinovic, P. Arduino, M. Garlock, N. Luco, J. Ramirez, and S. Yim, “2020 vision for earthquake engineering,” 2010. [Online]. Available: https://nees.org/content/article/206

[2] D. Ferry, G. Bunting, A. Megareh, S. Dyke, A. Prakash, K. Agrawal, C. Gill, and C. Lu, “Real-time system support for hybrid structural simulation,” in *ACM International Conference on Embedded Software (EMSOFT)*, Oct. 2014.

[3] H.-M. Huang, T. Tidwell, C. Gill, C. Lu, X. Gao, and S. Dyke, “Cyber-physical systems for real-time hybrid structural testing: A case study,” in *ACM/IEEE Int'l Conference on Cyber-Physical Systems*, 2010, pp. 69–78.

[4] “IEEE Std 1003.1-2017 - IEEE Standard for Information Technology – Portable Operating System Interface (POSIX(R)) Base Specifications, Issue 7,” 2017. [Online]. Available: http://standards.ieee.org/findstds/standard/1003.1-2017.html

[5] C. E. Leiserson, “The Case for a Concurrency Platform,” *Dr. Dobbs Journal*, Nov. 2008. [Online]. Available: http://www.drdobbs.com/parallel/the-case-for-a-concurrency-platform/212001540

[6] H. Zimmermann, “OSI Reference Model – The ISO Model of Architecture for Open Systems Interconnection,” *IEEE Transactions on Communications*, vol. 28, no. 4, pp. 425–432, Apr. 1980.