Cyber-Physical Systems: Position Paper

Paulo Tabuada
Electrical Engineering Department
University of California at Los Angeles
Tel. 310 794 4266, e-mail tabuada@ee.ucla.edu

***Abstract*— In this paper we discuss some of the technical challenges that need to be addressed in order to interface and manipulate the physical word. We also make some considerations regarding the requirements of cyber-physical systems.**

I. General considerations

Cyber-physical systems (CPS) will soon redefine how we perceive and interact with the physical world. Using commercially available hand-held devices we will be able to observe, change and even customize certain aspects of the physical environment that traditionally were beyond reach. An amusing but illustrative example is the following. Sport cars manufacturers partially justify the high prices of their models with the ”unique driving experience” that such cars provide. The characteristic driving sensations that each one of these cars affords were slowly developed over the years by experimenting with many different mechanical designs and engineering solutions and by incorporating the feedback of many test pilots. From an engineering perspective, this expensive ”unique driving experience” can be described by the response of the car to the stimuli produced by the road, by the driving conditions and by the driver. Moreover, such response can be recreated by using a ”programmable car”, equipped with a myriad of networked sensors and actuators monitoring and regulating the suspension, brakes, engine, assisted steering, etc. By downloading appropriately designed software into the ”programable car” we can emulate the ”unique driving experience” produced by any of our preferred sport cars. We can even combine aspects of different existing models thus customizing physical characteristics that were traditionally bound to different mechanical designs. This example attempts to illustrate one of the desirable characteristics of the next generation of CPS: the possibility to decouple, within certain limits, the physical environment or substratum from what is perceived by the end user.

In the remaining paper we will speculate about the necessary abstractions for locally physical but globally virtual CPS. We will pay especial attention to the extent that the physical substratum will shape the abstractions as well as the physical network of embedded sensors, actuators, computing and communicating devices.

II. Topological abstractions of the physical world

CPS usually comprise a network of physically distributed embedded sensors and actuators equipped with computing and communicating capabilities. Although each individual device is fairly inept at monitoring or regulating the physical substratum, the coordinated action of the individual network nodes has the potential for unprecedented capabilities.

A. Fundamental limitations

At present, the necessary coordination between networked embedded devices is hindered by high data[1] flows requiring an infrastructure supporting large bandwidths and quickly depleting the energy resources of individual nodes when communication is performed wirelessly. Although one typically envisions network elements communicating solely with neighboring nodes, the physical substratum dictates that measurements of physical quantities made at a given spatial location may be needed at a different and, more often than not, physically distant location. Moreover, when

[1] In Section III we will discuss some ideas on how to reduce data flows to information flows.

physical properties of the environment need to be regulated, the required data flows must consistently transverse the network at high rates in order to meet the stringent control objectives of safety critical applications. The important observation to be made is that the notion of locality induced by the physical substratum is not necessarily compatible with the notion of locality induced by the network of sensors and actuators. The later is usually identified with physical proximity while this is not necessarily the case for the former.

B. *Important research challenges*

A fundamental research challenge for the future of CPS is to understand how we can adapt the notion of locality induced by a network of embedded systems to the notion of locality induced by the physical substratum. One can conceive two different ways in which this adaptation can be done: physically and virtually. At the physical level one would judiciously chose the location, computing and communication capabilities as well as energy reserves of network nodes in order to handle the required network data flows more efficiently. At the virtual level, one can change the topological characteristics of the network by changing transmission power, medium access control, communication protocols, etc. Both of these approaches would require a model describing which physical locations should be considered as being close (or being part of the same locality) because frequent exchange of information is required. This model, that we shall call a topological[2] abstraction, would depend both on the physical substratum and on the desired services to be provided by the CPS.

C. *Promising innovations and abstractions*

A promising research direction is the computation of these topological abstractions explaining how CPS should interface, monitor and regulate the physical substratum. Although some recent progress has been made in the computation of discrete abstractions of the continuous physical world, most notably in the area of Hybrid Systems, issues of spatiality and locality have not yet been addressed. Equally important is the identification of the relevant notions of composition for these topological abstractions. Many CPS are formed by large numbers equal or similar physical systems as is the case in transportation networks, power networks, automated farms, etc. It then becomes crucial to understand how the the notion of locality of a large collection of similar physical systems can be obtained from the notion of locality of each individual physical system.

[2]The way in which constituent parts are interrelated or arranged.

D. *Milestones for the next 5 to 10 years*

1) Sound notions of topological abstractions;
2) Algorithms for the extraction of topological abstractions from formal models of the physical environment and the specification;
3) Integrate different abstraction layers for CPS with topological abstractions.

III. Self-organizing, distributed, in-network computation

Current research on CPS emphasizes the use of networked embedded systems as distributed sensing and data gathering devices. The natural next step is to consider *actuation*, thus moving from a *passive* framework, where information is extracted from the physical world, to an *active* framework, where information is sensed, processed and used *within the network.*

A. *Fundamental limitations*

When CPS are used to monitor or regulate the physical environment large data flows are required to transverse the network of sensors and actuators. These flows pose fundamental constraints on the performance and capabilities of CPS as they require a large bandwidth and are responsible for a quick depletion of energy reserves when communication is performed wirelessly. Mitigating this limitation requires one to regard CPS as distributed information processing devices that transform sensed data into information as we next describe.

B. *Important research challenges*

Regulation and control functionalities provided by CPS require the measurement of physical quantities at spatially distinct locations, the computation of feedback control laws based on these measurements and physical actuation based on these computations. Currently, control loops are

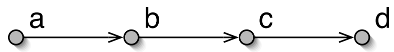

Fig. 1. Linear network illustrating in-network computation of feedback control laws.

implemented by transmitting all the sensed information to the network nodes where the actuators are collocated and where the control law is actually computed. This approach overly strains CPS since relaying the measured information requires, in general, the cooperation of several sensor nodes between the sensors and the actuators. We argue that a considerable part of the transmitted data is in fact irrelevant for the control task being solved and is thus placing an unnecessary burden on the CPS.

C. Promising innovations and abstractions

In typical regulation and control loops the number of actuators is much smaller than the number of state variables that need to be measured in order to compute the feedback control law. Consider for example the linear network in Figure 1 where the measurements of nodes a, b and c are used to compute a feedback control law determining the physical actuation provided by an actuator collocated with node d. If each measurement is encoded using n bits, and since each node transmits its data to the node on its right, we see that node a transmits n bits, node b transmits $n + n$ bits and node c transmits $n + n + n$ bits. If we now further assume that the control law is the linear function $\lambda_a x_a + \lambda_b x_b + \lambda_c x_c$ where λ_a, λ_b and λ_c are constants and x_a, x_b and x_c are the measurements of nodes a, b and c, respectively, we can reduce the data flow through in-network computation. Each node i would simply measure x_i, compute $\lambda_i x_i$, add this value to what has been transmitted by its left neighbor and relay the result to its right neighbor. This scheme would make node a transmit $\lambda_a x_a$ to node b, would make node b compute $\lambda_b x_b$ add it to $\lambda_a x_a$ and transmit the resulting n bits coding $\lambda_a x_a + \lambda_b x_b$ to node c and so on. It is not difficult to see that each node only transmits n bits and node d is not required to do any computation. This simple example illustrates that through in-network computation, sensed data can be processed by extracting only the information that is relevant for the control task being solved. It remains a challenge to be addressed in the future how to determine which amount of computation is performed by each network node for general nonlinear control laws and network topologies in a self-organizing manner thus catering for robustness with respect to changes in the environment and network topology.

D. Milestones for the next 5 to 10 years

1) Transform data flows into information flows;
2) Self-organizing in-network computation over dynamic network topologies.

IV. Brief Biography

Paulo Tabuada was born in Lisbon, Portugal, one year after the Carnation Revolution. He received his ”Licenciatura” degree in Aerospace Engineering from Instituto Superior Tecnico, Lisbon, Portugal in 1998 and his Ph.D. degree in Electrical and Computer Engineering in 2002 from the Institute for Systems and Robotics, a private research institute associated with Instituto Superior Tecnico. Between January 2002 and July 2003 he was a postdoctoral researcher at the University of Pennsylvania. After spending three years at the University of Notre Dame as an Assistant Professor he joined the Electrical Engineering Department at the University of California at Los Angeles.

Paulo Tabuada was the recipient of the Francisco de Holanda prize in 1998 for the best research project with an artistic or aesthetic component awarded by the Portuguese Science Foundation. He was a finalist for the Best Student Paper Award at both the 2001 American Control Conference and the 2001 IEEE Conference on Decision and Control and he was the recipient of a NSF CAREER award in 2005. He co-edited the volume Networked Embedded Sensing and Control published in Springer's Lecture Notes in Control and Information Sciences series.

His research interests include modeling, analysis and control of real-time, embedded, networked and distributed systems; geometric control theory and mathematical systems theory.