

Towards an Architecture for Distributed Cyber-Physical Systems

Tarek Abdelzaher, Associate Professor, UIUC, Biography link: <http://www.cs.uiuc.edu/homes/zaher>

Introduction

An expanding frontier for computer scientists lies at the intersection of the logical and physical realms. This direction is motivated by three fundamental trends that have great implications on the future of computer science as a discipline. These trends are:

- Moore's law. It implies increased miniaturization and reduced cost of hardware, leading to its gradual proliferation.
- The widening human/machine bandwidth gap. While computing becomes faster and more pervasive per Moore's law, the human information processing capability and human population evolve much slower. Over time, this leads to an increasing gap between the ability of computing devices to collect information and the ability of humans to consume it. Therefore, human participation in information collection and processing will become increasingly more marginal and at an increasingly higher (supervisory) levels of abstraction. Computing devices will collectively need to become more autonomous to avoid the human bottleneck. This implies that they will need to have their own means of interaction with their environment with progressively less human mediation. In other words, not only will computing be more autonomous but also increasingly more embedded.
- The cost of "lack of communication." There will always be a cost to the lack of communication or knowledge. For example, the quality of optimization and decision making can always be improved by interconnecting relevant information sources. This leads to a fundamental tendency for (i) global interconnection of the proliferating future autonomous, embedded devices and (ii) reducing their communication cost.

The above three fundamental trends herald the emergence of omnipresent, increasingly autonomous, globally interconnected networks of embedded devices with their own means of interaction with the physical world. Applications may include (i) disaster response networks, (ii) complex real-time control networks (such as power-grid control), (iii) environmental sensing networks to explore environmental phenomena at an unprecedented spatial and temporal granularity, and (iv) personal sensor networks. A fundamental question that arises is one of architecture. Namely, what are the reusable middleware components, network protocols, and system services, at different levels of abstraction, that can form a layered, distributed software architecture for cyber-physical computing? Is the existing networking infrastructure sufficient? Are current distributed middleware frameworks (such as web services) a good match to the needs of cyber-physical applications? To answer these question, we explore the deficiencies of current software systems and networks (when applied to cyber physical systems), highlighting new challenges that emerge in the envisioned future.

Fundamental Limitations and Important Research Challenges

Globally interconnected networks of embedded, possibly mobile devices offer a variety of challenges to address the deficiencies of current systems. These are described below.

- **Networking challenges:** The current Internet architecture is not designed to accommodate myriads of physical data sources, actuators, and distributed computing elements. The present architecture is optimized for point-to-point communication. Indeed, the main abstraction of TCP is a reliable point-to-point connection. In contrast, future autonomous networks of embedded devices will be centered around data fusion as the fundamental network purpose. This is a direct consequence of the integration of an increasing number of embedded devices

(such as sensors) with the Internet and the increasing need to elevate information abstractions to bridge the human/machine gap as mentioned above. Data fusion as the fundamental network purpose will likely require a different protocol stack that might take inspirations from distributed database design, data mining, control, and sensor networks. A fundamental challenge is therefore to redefine network architecture in a way that optimizes it for distributed information synthesis and retrieval as opposed to point-to-point communication.

- **Computing challenges:** New models and paradigms are needed for distributed cyber-physical computing. Past distributed computing paradigms (and current middleware frameworks that implement them) have attempted to abstract away distributed communication, providing support for location transparency and hiding communication details. In contrast, future cyber-physical computing paradigms will need to (i) abstract distributed, massively concurrent interaction with the physical world, and (ii) represent the external environment conveniently for the programmer. We call computing that satisfies the above requirements, *environmentally immersive computing*. For example, a future object-oriented (environmentally immersive) programming system might export the abstraction of logical objects that correspond to physical entities in the environment. Such objects should encapsulate the state of the corresponding physical entities. This encapsulation needs to hide complex details of state estimation in a distributed noisy environment, as well as complex distributed protocols for physical object tracking and disambiguation. New distributed middleware and underlying theoretical foundations are needed to support such paradigms.
- **Programming language design challenges:** Future cyber-physical applications will be increasingly geared towards data fusion. In the context of application development, at least two types of programmers will emerge; (i) those who perfect the distributed data fusion and signal processing algorithms, and (ii) those who write complex distributed, largely autonomous, decision and control systems that use high-level information (i.e., the outputs of data fusion) to manipulate networked objects. Current programming languages are not suitable for either purpose. Ideally, abstractions of future programming languages for cyber-physical computing should revolve directly around environmental elements (such as physical objects, events, activities, and data sources). These languages must have explicit support for the representation of uncertainty as a first class abstraction. For example, a typical object in an object-oriented language can either exist or not. In contrast, an object that represents an external entity (such as tracked vehicle) might be associated with a confidence level that describes how likely it is that this object truly exists in the external world. New programming languages must be designed to support these needed abstractions.
- **Software engineering challenges:** The increased interconnection of heterogeneous information networks and the increased reliance on distributed information processing offer more opportunities for functional and timing errors due to integration of progressively larger numbers of components into complex information processing networks. Current software engineering practice is greatly concerned with improving system robustness. However, it is geared more towards centralized or clustered systems. Future massively distributed embedded systems offer new challenges in software engineering and real-time guarantees that arise because of distribution. Ensuring timeliness and robustness of such systems in the presence of bugs, software integration errors, and complex timing interactions will be one of the fundamental concerns of cyber-physical computing.
- **Data management challenges:** The increased autonomy of future systems implies that data mining and machine learning techniques will play an increasingly important role in such

systems to identify data patterns, learn context, detect complex distributed signatures of events of interest, and generally act without human assistance. This has important implications on the design of network protocols and programming abstractions. Reusable tools will be needed to deal with data management functions in lossy, possibly mobile, poorly-structured environments.

- **Privacy and security challenges:** Privacy concerns must be met and security must be assured in future cyber-physical systems. The deep interaction with a distributed physical environment both increases the risks (e.g., the potential physical damage due to a security breach) and offers new opportunities (e.g., the use of physical data to authenticate nodes or detect intruders). Privacy and security mechanisms must be an integral part of the future architecture and not an after-thought.

Promising Innovations and Abstractions: Towards a Layered Architecture for CPS

Taking lessons from the success of the Internet and given the increasingly distributed nature of the envisioned, embedded, autonomous, cyber-physical systems, a layered architecture is needed to structure the development of future systems. Wide-area testbeds such as GENI might prove instrumental for experimentation with next generation cyber-physical computing protocols. The functionality of the cyber-physical system distributed “protocol stack” will be broader than that of typical communication stacks. It may include the following innovations and abstractions:

- **An information-centric protocol stack.** Future networks will focus around elevating information abstractions along communication pathways (from sensors to consumers). Hence, raw data that enters the network will be gradually converted to actionable high-level application-specific information that may be used for immediate action or learning. Scalable support will be needed for network storage and in-network processing. The emphasis on data storage and fusion as integral network functions may require rethinking fundamental network protocols such as those for congestion control. For example, in a network of storage nodes that perform processing on live and stored data, a critical resource becomes storage size (and CPU power) as opposed to only communication bandwidth. “Congestion” should therefore consider depletion of these resources, which is not necessarily the same as bandwidth over-utilization.
- **Environmentally-immersive middleware.** Services may be needed for easy composition of complex data fusion rules implemented by the network for higher-level applications. At least two different common data fusion purposes will need to be supported. The first is to identify and logically represent physical environmental objects from distributed sensory signatures. The second is to collect and track distributed object state. New object-oriented middleware is needed to support objects that represent entities in the physical world.
- **Integration analysis tools.** Tools (e.g., schedulability analyzers) and common architectural mechanisms will need to be developed to enforce and analyze distributed end-to-end system properties of complex cyber-physical systems such as their timing, security, reliability, and robustness to failures.

The development and integration of the above protocols, middleware services, and tools into a common architecture with layered, reusable, secure, fault-isolating components is a main necessity towards cost-efficient distributed cyber-physical system design and development. A 5-year milestone might be a GENI-based implementation of such an architecture. A 10-year milestone could be translating it into a common state of practice on a new Internet and its supporting services.