Report: Cyber-Physical Systems Summit

1. Introduction

The Cyber-Physical Systems (CPS) Summit was held in St. Louis, Missouri on Thursday, April 24 and Friday, April 25, 2008, at the end of the first CPS Week multi-conference.[1] Over 80 participants came to the CPS Summit to discuss the long-term scientific vision for this new field. The utmost importance and urgency of CPS for US industrial competitiveness has been highlighted by the August 2007 Report of the President's Council of Advisors on Science and Technology (PCAST) presenting a formal assessment of the Federal Networking and Information Technology R&D (NITRD). PCAST concluded that the Federal NITRD Program needs to be rebalanced and placed CPS as one of the top for priorities for substantial federal research investment.[2]

In response to this urgent national need, we were asked by the National Science Foundation to bring together academic and industrial leaders from a broad range of disciplines to help delineate: (a) a far-reaching and compelling vision for future cyber-physical systems; (b) the key technical challenges and the new scientific foundations required for cyber-physical systems; and (c) elements of an effective research program that will assure the success of the CPS vision. The CPS Summit built on the results of a series of NITRD and NSF Workshops exploring trends and key aspects of this emerging area. The breadth and depth of industrial attendees at the Summit (National Instruments, Toyota, Microsoft, NEC Labs, Honeywell, United Technologies, Rockwell Collins, Crossbow Technology, BAE Systems, MathWorks, Johnson Controls, Lockheed Martin) illustrates the importance of this area to industry.

This report describes in detail the discussions held during the summit. The next section frames the problem. Section 3 describes grand visions and grand challenges that could be used to drive research in cyber-physical systems. Section 4 describes scientific and technological foundations that need to be built for this new field. Finally, Section 5 proposes possible elements of and strategies for a bold CPS research initiative.

2. Motivation

The CPS Summit was predicated on an emerging consensus that we need to formulate a 21st-century notion of cyber-physical systems—systems in which computing, broadly construed, interacts with the physical world. Traditionally, one of the most visible and successful integrations of cyber and physical systems have been computer-mediated control systems. In today's applications, cyber systems perform sophisticated sensing and decision functions that go far beyond simply closing dedicated feedback control loops. For example, in the DARPA desert and urban challenge vehicles on-board cyber systems acquire data and process information from a broad and diverse set of sensors and cameras to determine the vehicle location, understand the terrain, and detect the locations of other vehicles, people, obstacles, signs, etc. Although there are many such sophisticated embedded computing systems today, there few examples of true CPS exhibiting a

[1] http://www.cpsweek.org/

[2] http://ostp.gov/pdf/nitrd_review.pdf, p. 31.

seamless, fully synergistic integration of sensing, computation, and control with physical devices and processes.

The principal barrier to developing CPS is the lack of a theory that comprehends cyber and physical resources in a single unified framework. Despite the 70-year history of computing (dating back to the early electronic computers of the 1940s) and the 200+-year history of engineering physical systems that are receptive to control (starting from Watt’s governor), the fields of computer science and control theory have remained largely separate, both technically and culturally. This separation extends to virtually all domains where computers interact with the physical world. Methods for designing computer systems and physical systems are based on simplifying assumptions about each other that limit the range of systems that we can build. On the one hand, computer engineers and scientists do not know how to translate requirements for physical systems, such as stability, into computational requirements on performance, power consumption, etc. On the other hand, control and signal processing theory abstract computers largely as infallible numerical devices. This simplification ignores many important aspects of computing, such as increasingly larger timing variance due to caches and energy management and increasingly higher software error rates caused by complexity. Simplifying assumptions are also made about communications. Initial designs assume zero-loss, zero-delay communications, while neither occur in the wireless, low-power, shared, rapidly changing systems used in most CPS. The viability of future CPS must also address noise in measurements, inaccuracies in actuation, disturbances from the environment, and faults and failures in the computational process in a coherent, unified framework.

CPS clearly has a role to play in developing a new theory of computer-mediated physical systems. We believe that it will have the biggest impact, both in terms of scientific development and economic payoff, in large-scale distributed systems. In many systems, such as the national power grid and traffic control systems, both the plants and the computers for monitoring and control are physically distributed. In such systems, the dynamics of the distributed computing platform and the distributed plant interact in ways that determine the overall operation of the system, but are as yet poorly understood. Recent advances in computer/communication technology have made it easier to more tightly couple distributed physical systems than was previously imaginable. This can lead to more efficient management of large-scale systems such as the power grid. But it also gives rise to emergent behaviors that limit our ability to build predictable large-scale networked systems.

We believe that CPS will emerge as the critical technology that underpins every major industry in the United States. Exactly how the field will unfold is unclear: as a new discipline, an extension of systems engineering, integrated into existing disciplines, a new form of computer engineering, etc. In whatever way the field develops, we believe that CPS will transform the way we understand the relationship between computing and the physical world.

3. Grand Challenge Problems

The discussion at the Summit of grand challenge problems focused on 15-to-20-year visionary applications and how to differentiate those applications from current state of the art or near-term research problems and from research directions outside CPS. Where possible the vision for how to approach these problems was also considered, though it was noted at several points that differentiation from what we can envision now is also crucial. Key to the motivation for these grant challenges is a strong societal need.

The discussion focused on areas of potential interest to NSF. A number of applications were identified that could present grand challenges for CPS, including:

- redefining mass production, e.g., components built to looser initial tolerances that adjust to each other upon assembly;
- a worldwide sensing utility (with web browser for the physical world);
- food and clean water to everyone in the world;
- pervasive sensor networks that never need to have the batteries replaced;
- VLSI-like tool chains for autonomous robots;
- climate monitoring and control for global warming;
- make it possible for elderly people to live independently;
- robots truly interacting with people (nurse bots actually examining people, robots playing soccer with people);
- context-aware personal tutors for all types of skills;
- automated farming using farmbots and sensor networks;
- 24/7 robotic personal assistant;
- optimal harvesting of solar and wind power;
- net-zero energy buildings and homes;
- energy efficient cars;
- persistent ocean-observing systems;
- integrated medical device systems: lab-on-a-chip to doctor-on-a-chip;
- tele-presence/tele-operation of assistive devices;
- neural prosthesis for people with paralysis using the control signals from the motor cortex;
- sensor-actuator networks at the cellular level—hundreds or thousands of sensors/actuators in explants and in vivo implants to advance our understanding of biology, and in the long run improved health care;
- zero medical deaths from metabolic disorders through ubiquitous monitoring and control.

From the many problems that were discussed, three main grand challenge problems areas were identified and developed more fully: future distributed energy systems, future transportation systems, and next-generation healthcare systems.

3.1 Future Distributed Energy Systems

Current energy generation and transmission consists of a decades-old transmission network delivering power radiating from largely fossil fuel generation. As fossil fuel usage decreases due to reduced availability and the concerns over global warming, there will be increased emphasis on renewables and on local energy generation and storage. Thus, energy systems of the future will become increasingly heterogeneous in the kinds of energy sources and their locations and capacities, and increasingly autonomous in terms of when participants draw or generate power (and how much) from a common inter-connected distribution grid. How to coordinate distributed and dynamically interacting participants and how to control the physics of the common grid will become an increasingly significant problem as current trends toward adding solar panels to roofs (e.g., Sun Energy California project) and windmills to farms, and even back yards, continues. These trends are also likely to intersect with the potential economic incentive to recoup personal investments in energy source technology through generation as well as consumption of power. In addition to the technical considerations of power transmission physics, control, and energy storage, the social effects of the increasing national interest in energy technologies are thus likely to be significant as well.

The technical challenges that must be overcome to address the emerging and future situational contexts for power generation and distribution are both significant and hard. For example, we must enable self-sustaining power cooperatives at the local to regional scales with minimal external energy dependence and stable and fluid power generation, storage, and distribution in the contexts of (1) energy management as a collection of distributed communities rather than a radial power generation and distribution strategy, (2) inter-connecting such communities via the existing power grid infrastructure, (3) managing the effects of connecting such established communities to each other and to traditional power grid participants at the regional or national scale, and (4) injecting new sensing/actuation technologies (e.g., PMUs/FACTS, video, etc.) into the existing power grid at the regional or national scale.

A concrete example of a grand challenge problem in this space is to design and deploy (either intact or through transformation of the existing power grid) an electric energy system infrastructure that (1) is flexible in allowing heterogeneous participants to supply or draw energy dynamically, (2) is secure against inappropriate manipulation and is resilient against accidental or intentional damage, (3) doesn't black out or introduce dangerous power surges, (4) degrades gracefully in the face of damage or unmanageable conditions (e.g., over-supply or over-demand), (5) is efficient in real-time (e.g., 25% or better improvement), and (6) supports seamless evolution of its architecture, control, and implementation as unforeseen future technologies enable further evolution of the grid. Existing test beds at the national labs appear likely to be extensible in place with additional adaptation, control, and other "intelligent" elements, to achieve new outcomes like: (1) fine-grained integrated and context-aware optimization of efficient energy generation, storage, transmission, and use, to achieve specific (e.g., 50%) reductions in waste, carbon footprint, costs, etc.; (2) multi-regional and/or multi-national coordination, pricing, and control of power supply and demand; (3) "just enough" generation and use of energy; (4) computational definition and management of energy needs, costs, carbon footprint, etc.

In order to cope with continually increasing demands (especially from air conditioning, but soon with pluggable electric vehicles) the system capacity will either need to be built up -- or made smarter. Smarter is where CPS can play a major role. Today's power grid power runs one-way downhill from generation sites to customers. The system is designed for peak power loads, which are far larger than the average usage. Being able to lower peak usage can save very large amounts of capital spending and lower risks of system failure. In order to lower peak loads, the grid needs to be able to be two-way and have intelligent sensor/decision systems throughout that can make detailed decisions (change thermostat settings, postpone defrosting the refrigerator, charge or draw charge from an electric vehicle) and to handle power generated by customers (wind, solar panels, etc.). Expected power needs and costs need to be predicted, based on weather, time and day of the week, energy futures costs, etc., and all this needs to be coordinated on a highly distributed system.

The existing civil infrastructures such as the current power grid also need to be transformed. This issue cross-cuts several of the grand challenge problem areas, and warrants consideration as broadly as possible, including end of lifetime of infrastructure (e.g., power grid is around 75 years old), how to upgrade/retrofit/augment existing infrastructures with CPS policies and mechanisms, and tools for going beyond worst-case planning and design to achieve just-in-time, just-in-place, situationally exact solutions. An example project in this space is in-place transformation of the existing electric energy generation/distribution system (including the power grid), without loss of reliability. For example, as the power grid is augmented with cutting-edge FACTS devices, so that eventually 5% then 10% then 25% and then 50% of power grid consists of facilities that don't exist in the state of the practice today, what becomes different about the new grid, and what does CPS add? What picture emerges as we chart progress towards goals including cost reductions, reliability increase as the grid is transformed?

3.2 Future Transportation Systems

The nation's transportation infrastructure is significantly overstressed. Roadway corridors and air traffic corridors continue to have increased congestion resulting in lost productivity and wasted fuel usage. Tremendous expense is required to increase roadway traffic throughput. Congested terminal air traffic results in increasing delays with no real way to expand the airspace. Improving efficiency holds promise in making better use of existing infrastructures. Thus, future transportation systems need to incorporate both manned and unmanned vehicles of different granularities (e.g., single person vs. mass transportation vehicles) and mixed criticality (e.g., freight vs. passenger vs. emergency vehicles moving within, and possibly between, air and ground spaces). This domain shares goals similar to the goals of the future distributed energy systems grand challenge problem: flexibility of use, security, stability, safety, graceful degradation, efficiency, and seamless evolvability.

These systems will enable bulk transport capabilities beyond what we can imagine, with higher throughput in already congested corridors, personal/package latency four-times better than is available today, etc. Realizing this vision will require significant research advances in hard/soft real-time control and optimization, composable hybrid control architectures that support mixed initiative run-time safety guarantees, and practical techniques for verification and validation of safety, particularly in the face of the likely complexity of aeronautic, automotive, and mixed aeronautic/automotive vehicles and of the infrastructures needed to ensure safety and reliability of the entire transportation system.

Example projects in this space include: (1) safe, certifiable air transport systems created with less than 50% of current resource demand and feasible in-process technology upgrade; (2) safe cyber-physical regional ground transportation systems with zero accidents and zero fatalities, lower personal commute times, higher total throughput, and lower maintenance costs; and (3) mixed transportation / tele-presence / tele-immersion / tele-operation systems, where the physical locations of the participants are flexible but can be optimized effectively, efficiently, and dynamically according to situational demands and constraints (e.g., providing cost-efficient state of the art healthcare in widely dispersed rural settings, or 50% greater reduction in costs for business meetings while increasing personal flexibility).

The scientific and theoretical underpinnings of this grand challenge problem stem in part from the fact that while people with a limited amount of training are allowed to drive cars, it currently requires highly trained pilots to fly planes. Both systems assume decisions are made according to established protocols, but the complexity and consequences of failure in the aviation setting are significantly more challenging. Another key issue regards both passenger and operator comfort levels with the systems and protocols involved. Human-centered design, implementation, and certification issues, need to be considered, including how we understand comfort, usability, and correctness, and how design, verification, and validation can incorporate that understanding. Scientific challenges include: personal and aggregate safety; emissions and energy consumption; personalization of travel; cost (design technology for optimization); availability and reliability; social acceptance and impact; risk and hazard models; and how to fit new technologies and vehicles within existing / maintained / evolved infrastructure.

As it is for the future energy systems grand challenge, the issues of how to upgrade the current infrastructure is also crucial for future transportation systems. How can we reconcile the every day transportation needs (e.g., best routes for personal travel vs. emergency vehicles), with how to respond safely (and perhaps optimally) to exceptional situations such as a terrorist attack, an earthquake, or a structural failure? How can we address lifecycle mismatches – e.g., from civil infrastructures lasting 100 years, to vehicles lasting 10 years?

Potential stakeholders in such a transformative effort run the gamut from individual vehicle technologies (e.g., a car that never fails – Toyota and other auto makers), to vehicle-to-vehicle systems (e.g., a safe personal air vehicle network – Boeing and other aerospace companies), to national or international infrastructure (e.g., maintaining a CPS-enabled infrastructure – US Department of Transportation, Federal Highway Administration).

3.3 Health Care Grand Challenge

Certifiably safe automation of healthcare diagnosis and delivery is a grand challenge problem that has high social relevance, as the entire population uses the current health care system and is thus a stakeholder in its improvement. Current trends, including demographics and escalating costs of equipment and procedures, are putting the ability to provide reasonable-cost, high-quality healthcare services at significant risk both nationally and internationally. This grand challenge focuses on how to achieve affordable radical improvements in quality of care through (1) certifiably safe medical device systems on-site, (2) remote multi-tasked diagnosis and prescription services that give the greatest benefit possible from the investment of healthcare personnel (both time and expertise), (3) certifiably safe remote intervention and delivery procedures accomplished through those devices and services, and (4) a transformation of medical training, practice, and infrastructure that enables and advances this approach. An essential question is whether (and if so how) CPS can make more effective (i.e., safe and radically better) and efficient (i.e., less costly) use of scarce healthcare resources including technology, expertise, and finances. By removing technology barriers, CPS has the potential to enable "destabilizing" transformative benefits to the populace that include social acceptance and regulatory effectiveness, in addition to the core ambition of better quality of care at lower cost.

A concrete example of a problem domain for exploratory projects, which already is receiving attention within the CPS community, is the certifiably safe medical device system composition problem that has been discussed at a recent workshops on High Confidence Medical Device Software and Systems (HCMDSS) and Medical Device Plug-and-Play (MD PnP).[3] Research challenges include (1) how CPS can help to remove improper interactions among devices and personnel (with resulting reductions in mortality and morbidity), (2) how devices and personnel could be co-certified for specific procedures in specific contexts, and (3) how diverse devices and IT resources can be interlinked safely and effectively.

A grand challenge goal would be zero medical deaths from metabolic disorders through ubiquitous monitoring and control. Ubiquitous monitoring and control subsumes sensor networks but includes things like ubiquitous monitoring and control of metabolic disorders and structures. Type II diabetes is a good example, and might include active monitoring of sugar levels in urine (via smart toilets), active monitoring of blood sugar and insulin release (via implantable devices), wireless communication of data with doctors and central health administrators (via wireless hubs at home), data collection and statistical analysis (via centralized databases), dosing updates (also via wireless), communication with doctors for off-line analysis, and even active alerts sent to the patient and/or health-care providers.

4. Scientific and Technological Foundations

This section presents several scientific and technological challenges for CPS identified by the Summit participants. The discussions were rich and varied, with many themes recurring in different contexts. The many issues are tightly inter-related and could be organized in many different ways. The following paragraphs attempt to capture the many issues and insights by organizing the notes from the discussions into a set of primary topics. These topics are presented below in alphabetical order so as to avoid any suggestion if relative importance by the order in which they are presented.

4.1 Compositionality

Compositionality that cuts across the heterogeneous cyber and physical aspects of CPS is a major scientific challenge. Separation of concerns can be achieved by defining flexible interfaces that support reprogrammability, leading to more adaptive compositions that may address legacy issues. Another challenge related to compositionality is modeling and predicting performance of composition. Basic research in compositionality leads to reduced challenges in system integration of both subsystems and systems-of-systems. Applied research is needed on open systems that protect proprietary intellectual property while supporting flexible interconnections between components from current and future vendors. An outcome of this research could be "plug & play" cyber-physical components that integrate seamlessly.

[3] http://rtg.cis.upenn.edu/hcmdss/index.php3 and http://rtg.cis.upenn.edu/hcmdss07/index.php3

Many tools and approaches exist for creating components and composing them. There are a large number of models, languages, and notations that exist, however, many of which are most appropriate only for particular problem or areas. No complete solutions exist for CPS. A compositional theory for CPS must be created that can integrate synchronous, asynchronous, continuous, and discrete models. The theory must account for computation that involves time, location, memory requirements, cost, energy and security requirements. Uncertainty and the realities of the physical world must also be considered, including uncertainties arising from wireless networks that will provide the communication infrastructure for many cyber-physical systems. From the theory it must be possible to analyze the aggregate behavior of a complex and large CPS. A new notion of correctness must be developed that includes safety and possibly a concept of asymptotic behavioral properties. Since CPS will exist in and among humans, they will be open systems giving rise to a high degree of continually evolving characteristics. Correctness of a system requires runtime monitoring and verification as the system evolves. Analyzing such systems will be a major new undertaking.

4.2 Distributed Sensing, Computation and Control

It is not clear we have sufficient paradigms for making distributed control, sensing, and communication, in safety and time-critical CPS. Key problems include collecting adequate information and asserting control in a distributed environment. For example, when do samples need to be collected for distributed decisions? What information needs to be collected? Where should the computations be performed? Traditional models of sensing, computing, decision making, and acting are based on low-latency environments in which the data is fresh and the actions are executed immediately. In a networked cyber-physical system, latency can create uncertainty in the decision making process; latency can create serious problems if actions are executed (become effective) too late. Current uncertainty decision models do not support the traditional control loop model, let alone a distributed control environment.

CPS may exhibit dynamic and explicit global coupling, cascading actions with adverse effects, with only partial information available at the time scales within which actions to stop such cascades are needed. New science and theory is needed regarding how communication can facilitate safe operation, failure interlocks, and adaptation. Integration of information and actions across time (with understanding of uncertainty at different scales) is also essential.

Many CPS are inherently distributed (and increasingly other traditionally centralized ones are becoming more so) so that one can’t rely on a notion of a monolithic “control room” in such systems. This in turn leads to a lack of centralized measurement and decision making, that, when combined with the locality of physical effects and with propagation of those effects at small time scales, motivates new science and theory on time-bounded establishment of context, malleable dependence structures, and distributed cyber-physical control.

4.3 Physical Interfaces and Integration

Contacts with the physical world characterize an essential feature of many CPS. Systems with physical contacts include wheeled or legged mobile sensor platforms, medical catheters and laparoscopic instruments, nano-sized cell and protein manipulators, and human exo-skeletons, all of which can do tasks requiring mechanical work—either autonomously or by augmenting human capabilities. Physical issues also arise in the cyber infrastructure, such as wireless networks.

Foundational research supporting the development of CPS with contact include a hierarchy of models of physical contact with varying levels of resolution and degrees of freedom (e.g., soft tissue, hard surfaces, slippery or sticky), the mathematical properties the models in the hierarchy, the corresponding algorithms with known performance properties, and the couplings of the contact models to the cyber and other physical components of the system. This knowledge must then be applied in the design process to characterize co-stability of a cyber and physical system and to predict its cost and performance.

4.4 Human Interfaces and Integration

An important barrier to operator-mediated CPS stems from the need to interface the modeling, control, and adaptation of CPS with human influence and perception. Time scales for system adaptation and human interaction need to be meshed. Multi-scale modeling needs to be integrated at human response time scales with appropriate on-line interlocks to avoid interference between finer-grained and human time scales. Predictive specification of performance needs, modal responses, and other approaches also need to ensure congruent local context, so that incongruent actions by different people (“friendly fire”) are avoided. How to express local context, make systems aware of it, and incorporate it with both the objectives of users and the behavior of system in its environment remain important open problems for CPS. Mixed initiative systems in which the CPS are expected to operate semi-autonomously raise additional issues of transfer of authority, tolerance of operator neglect, and common views across system modes that may vary at finer time scales that human perception.

Traditional HCI and human-centered computing are cyber physical in the sense that the human is the physical part. Where CPS might hope to break new ground is not in the conventional way - how do we successfully marry what people are good at with what computers are good at, in order to help people accomplish tasks - but rather by moving up a level and asking, for example, how we achieve certain systems "ility" properties (reliability, security, scalability, safety, etc.) when humans are part of the system. This requires understanding humans beyond biology, i.e. as agents with intention, purpose, and behavior. Right now, computing is embedded in artificial devices (e.g. mobile phones). Once GPS and various kinds of sensor data become available to those devices, the applications will change. For example, if a person is about to step in front of a truck, their cell phone might emit a very loud noise, and their shoes might somehow constrain their feet.

Human backup for CPS failure could be realized. CPS will make extensive use of autonomous control, and may not continue to expose manual interfaces in some cases, despite the need for potential intervention by human operators in the event of anomalous system behavior. For example, pilot/autopilot interactions increasingly appear in multiple CPS domains (e.g., skid auto-steering in automotive applications). New science and theory is needed to define safe hand-offs between human and cyber-based control, cyber-physical interlocks, protocols for mixed initiative inter-operation, and maintaining the operator’s mental model and appropriate skepticism. Designers of human-computer systems need to consider additional issues such as: tradeoffs between thinking (computation), talking (communications), and moving (control) in terms of power consumption vs. performance. In order to tackle these problems, a unified model is needed.

4.5 Information: From Data to Knowledge

CPS will be used for a wide variety of applications where the main intent of these systems is to develop knowledge and use that knowledge for major improvements in the particular application domain. For this to occur, it is necessary for the physical and cyber layers to interact in a synergistic fashion. The effect of sensor properties, including calibration, context information, and real world uncertainties, must be accounted for in sensor fusion algorithms and higher-level information creation. Knowledge of the reliability and trust of the data sources are necessary. New theories and models are required that capture the raw-data-to-trusted-knowledge dependency chain of processing. It is also valuable to feed back assessments of the created knowledge to the physical layer. For example, to increase knowledge belief it may require an increase in sensing rate or activation of additional sensors or sensor types at the physical layer. Most current techniques for knowledge acquisition do not act across all layers or operate as adaptively as will be required by open CPS.

Mining of data streams in real time is significantly different than knowledge acquisition from static data. More than simple throughput is required, though that is a challenge in itself. We need to better understand how to keep and manipulate several representations of data, some closer to the sampled data and some more abstract, so that we can efficiently search and browse real-time data. We also need history- and information-aware data repositories that allow us to retrieve meaningful information in a more direct/effective/targeted way.

We would like to go beyond safety, reliability, trust, verification, formal tools, programming languages, etc. to applications that require more of a focus on the basic functionality: new perception, planning, or control techniques. e.g., use of smart cameras.

4.6 Modeling and Analysis: Heterogeneity, Scales, Views

CPS are composed of components that exhibit massive heterogeneity. Components may have different notions of time, across different scales. In a heterogeneous, physically-aware CPS, feedback can occur through both the cyber and physical environments. If it occurs through the physical environment, then everything changes in terms of feedback control models. Heterogeneous CPS models of time, discrete or continuous, do not work together. Similarly, event-driven modeling tools (for asynchronous systems) and time driven models (for synchronous systems) do not work together, though both are applicable and likely to be useful. Since the role of time and events will be critical in CPS, new hybrid models will be needed. Another significant challenge will be the integration of multiple scales of temporal and spatial resolutions within a heterogeneous, physically-aware CPS.

Most of computer science deals with computational models that focus on abstract notions of time on either real or abstract machines. Many formal and practical properties can be derived with these models. However, for CPS new models of computation are required that explicitly address new observables: time, location, energy, memory size, cost, and uncertainties of sensing data. Such models can integrate the cyber and the physical, thereby more accurately representing physical systems. Since CPS will transcend many orders of magnitude in many dimensions (time, space, energy, etc.) any new models must allow for composition at multiple scales.

New abstractions and models are needed for cyber-physical control. CPS control must consider uncertainty in cyber "plant" on which controller runs, and from that needs to reason about and compensate for uncertainty in controller operation based on the potential effects of the cyber part. Research needs on this topic include tractable verification based on uncertainly; control/computing co-design for safety; and feedback-based approaches to uncertainty modeling and compensation that address the problem that bounds on uncertainty may not be known a priori.

The physical aspects of CPS cannot always be abstracted away. The uncertainty added by sensors, wireless communications, noise, and mobility often affect the cyber design, implementation, and performance. Real-time and safety requirements are often paramount. The impact of the physical layer via these properties is so profound that when building CPS, errors and loss (of sensing, communications, and nodes) are natural and common. Solutions must be robust and operate in this type of non-deterministic, probabilistic, and delay-induced environment.

In CPS, the trajectories of physical (voltage, distance, etc.) and cyber (memory, CPU, communication, etc.) states are coupled, so that it is essential that CPS be designed in both cyber and physical domains to account for each other's timing and physical trajectory, errors of approximation, potential implementation flaws, etc. New science and theory is needed for design in the face of both cyber and physical imperfection. What imperfections do we expect in software? In hardware? In control devices? How do platform/control/software interactions influence reliability? Loosely coupled networks of physically coupled systems, in which system reliability can be achieved even with unreliable parts, motivate development of theory for predictable and "stable" adaptive systems.

A further need for CPS is to support concurrent engineering of evolving systems, in which people can add constraints and requirements in local context, which then need to be visible in other relevant contexts, and need to be integrated within those contexts e.g., multi-view modeling with context-aware risk analysis and mitigation. In this approach, construction of the system is never finished.

Large-scale systems need better models of the trade-offs between system level efficiency and local efficiency. Several Federal agencies responsible for large-scale systems---FAA, DOE, etc--- have done studies, established test beds, etc., but further investigation is needed in collaboration with those groups. Most engineers and computer scientists would agree that our ability to design very complex large-scale systems is not as strong as we would like, particularly when we consider very long-lived systems. While we can analyze particular use cases, we would like to better understand tipping points and other emergent behavior; more importantly, we would like to understand how to design systems to avoid undesirable versions and encourage desirable versions of emergent behavior.

4.7 Privacy, Trust, Security

The discussion of key scientific and theoretical challenges for CPS began with the topics of security and privacy. CPS raises new issues in these topics because physical systems reveal information, there are limits on what information can be hidden, and new kinds of physical and cyber-physical attacks are possible. New science and theory needed for CPS on these topics include design principles for resilient CPS, threat analysis vs. hazard analysis, theories of cyber-physical inter-dependence, and examination of the possible role of gaming of different layers of the system.

CPS will create an enormous amount of data. The physical aspects of CPS will create new and difficult privacy problems. In particular, vast quantities of data will come directly from sensors and RFID tags. Devices pose new problems such as easily revealing location and time as well as the identities of individuals due to physical-layer fingerprints that exist with most transceivers. Sophisticated inference techniques using this vast amount of data will make maintaining privacy a very difficult issue. New research is needed to create formal models that can specify CPS privacy requirements. To date, privacy models focus on permanent user records and not real-time sensor data. With well understood semantics provided by these formal models, new research can prove various privacy properties of systems. This will place privacy on a scientific basis not yet achieved. It is anticipated that new mathematical theories will be needed to hide private information coming from real-time sensor streams. Current statistical techniques have focused on static data.

CPS are also susceptible to additional security attacks beyond those found in cyber systems. This includes jamming the communications, physical tampering, overhearing and many more. The limited capacity of the nodes makes providing security of CPS extremely challenging with a need for “lightweight” (in terms of overhead) solutions. Most solutions for the cyber world require significant memory and execution time.

In CPS, physical and cyber elements motivate different models of trust so that erroneous behavior is detected and human operators maintain appropriate skepticism during system operation. New science and theory is needed to define cyber-physical inter-confidence and trust maps, CPS context dependent trust models, and ground truth detection capabilities (based e.g., on real-world physical limits).

4.8 Robustness, Adaptation, Reconfiguration

CPS will operate in very dynamic environments. Thus, it is imperative that small changes in operating assumptions do not lead to uncertain outcomes for the cyber-physical system. Contrary to well-established methods of robust control, which can handle modeling uncertainty, we need new notions of robust system design that address uncertainty at the cyber level (computing decision or scheduling choices), and are resilient with respect to massively uncertain/untrusted data and structural/topological uncertainties and reconfiguration. CPS will also need to be reconfigurable and adaptive to overcome faults in both physical and cyber levels, e.g., to enable the management of fault containment in power grid-like scenarios.

CPS uncertainty bounds may not be known a priori. Families of related designs may need to be certified together, both to select among them and to leverage analysis.

Techniques for identifying, categorizing, and selecting appropriate recovery trajectories for CPS failure modes are needed, along with system science and design processes that offer guarantees of evolution toward “good” designs that can increasingly account for discovered faults and changing uncertainty measures. The need to evolve CPS designs also raises the question of incremental deployment of control systems in extant contexts (e.g., roadways) as CPS features are added, modified, and extended.

4.9 Software

Traditional programming languages for embedded systems do not provide native support for many of the scientific challenges identified for CPS. New programming languages and computational paradigms are needed to provide structure to manage the anticipated complexity and scale of CPS. Moreover, languages will need to be linked to semantic and performance models as part of a correct-by-design synthesis methodology.

In CPS, adding software-based control may introduce cyber-physical coupling because the software depends on a plant (the hardware platform) that runs it. New control theory is needed to reason about the behavior of the software plant model integrated within the system model, analog consequences of software-in-the-loop behavior, quantization of nonlinearities, and scheduler jitter in the system dynamics.

One goal of this theory should be physical/controller design against software faults. At the present, we have no foundational links between computational faults and control-theoretic error conditions. To solve these problems, we may have to rethink the very notion of what computation is (the Turing-Church thesis), the notion of algorithms at the center, the foundations of programming languages (semantics, expressiveness), models of concurrency (process algebras, nondeterministic transition systems, etc.), etc. No widely used programming language has any element of time in its semantics is a consequence of these foundations. We can optimize for aggregate, average-case performance only because timing properties are irrelevant to correctness of execution of any program in any modern programming language (by "correctness" we mean relative to the semantics of the language, not relative to what is needed in the application).

To facilitate these advances, it is very important for the CPS community to develop a "vertical integration" framework in which various techniques (for modeling, analysis, programming, communication, OS, networking, etc.) can be integrated to provide a holistic design-and-implementation solution for CPS.

4.10 Verification, Testing, and Certification

We need to develop new theories of correctness for CPS that allow new “correct-by-construction” approaches: property preserving transformation of existing and new systems, CPS requirements through specification through design through implementation, correctness validated by testing assumptions (rather than by attempting to test everything). We also need methods for reasoning about the co-stability of cyber and physical domain features: degrees of freedom in physical design, degrees of freedom in cyber design, and coupling of cyber and physical design assumptions. Verification of these properties will be particularly challenging for open systems and systems based on wireless communications.

Composition (“on beyond compositionality”) that deals with masking faults, can detect divergence of system behavior from specification, and can reduce and control interaction complexity appropriately is needed. These concerns offer a possible philosophical revolution for CPS: "on beyond correctness" in which we need to quantify the degree to which systems are x-tolerant (not just a 0 or 1 determination) that includes proof of properties and potential behaviors vs. risk (stochastic, etc.) rather than only existence of potential modes of failure.

Another important barrier that must be overcome for certification and verification of CPS is the fundamental uncertainties about what hazards there are to different kinds of CPS infrastructure, and how those hazards should be represented. For example, compositionality is always probabilistic in Civil Engineering, with "high confidence of low probability of failure in the given environment" as the objective of verification and validation.

Assumption validation is a particularly important barrier for CPS, and new approaches are needed to ensure that CPS models can: represent physical constraints and situational assumptions accurately (what should be captured, how to specify, what are degrees of freedom); ensure completeness of treatment and robust handling of changes; and can represent effects of both cyber and physical domain assumptions in the system model (e.g., to avoid problems like the Arian-5 rocket buffer overflow issue and the Hubble lens issue). Another important challenge is how to capture and check design assumptions automatically, and within the design process to invalidate assumptions when they no longer can hold or when they are no longer relevant. This raises a notion of lifelong correctness through verification maintenance as well as design and implementation maintenance, which is essential to achieving X-tolerant systems – tolerance to tampering, adversaries, environmental changes, etc.

Europe has embraced formal verification of embedded systems to a much greater extent than the U.S. has. While verification, synthesis, and testing of general software systems remains a grand challenge, embedded systems provide much more structure than general software systems, which may result in much more tractable problems. Formal methods offer one avenue of achieving high confidence in CPS, but formal models must also be linked to performance models, which is not feasible in large scale systems today. In all likelihood, the complexity of CPS will result in the need to develop evidence-based methods of asserting correctness, which should be combined with correct-by-design synthesis methods, as well as bridging the gap between testing and verification.

4.11 Societal Impact

The impact on society and the impact of society on the resulting engineered CPS cannot be under-estimated. Each grand challenge discussion noted the need for social acceptance of new systems. The social science aspect must be treated on a uniform basis with the computer science and engineering aspects of the CPS so that rigorous models of these social systems become a critical aspect of design, verifiability, validation, operation, privacy, trust, and fault tolerance.

5. Architecture of a CPS Research Initiative

A national initiative on CPS is needed with participants from industry, academia, and government laboratories and agencies. CPS needs to build high-cost high-risk systems to study, with an initiative structure based on desired outcomes, and incentives for participation across the board. The greatest outcome for CPS would be that the current

US competitiveness crisis (in real-time and embedded systems and beyond) would be addressed, through the creation of a vibrant new research field of CPS with high societal impact and ongoing technology advance as a result.

The National Science Foundation has an important role to play in both the technical and cultural aspects of this new field. Clearly, NSF can develop new programs that will encourage proposals aimed at developing fundamental aspects of CPS, research infrastructure, etc. The proposal process is a key part of building a new culture of CPS that spans the not only computer science and engineering, but all of the NSF directorates. In addition to single-investigator grants, NSF can support group projects at several levels of scale:

- Small teams (2-3 investigators) can look into specific topics at the boundaries of fields, such as control and security.
- Large teams (4-6 investigators) can pursue larger studies that delve into several aspects of CPS. These efforts may also build systems to test out new theories.
- Centers can provide an umbrella for comprehensive CPS research efforts, build research infrastructure that can support the entire research community, and involve industry in problem definition and technology transfer.

Given the large scope of this new field, we expect that other Federal agencies (as well as state funding) will contribute to the development of CPS. NSF’s budget is not large enough to support all of the research required to build CPS to a mature field, nor is it chartered with technology development *per se*. Equally important, different agencies can bring new points of view into the development of this field. CPS will change the way that industry works, how government develops infrastructure, and how the military defends the United States. Agencies can cooperate to develop the core principles of CPS while developing sub-specialties within CPS that will support their own requirements.

The European Union’s Artemis effort is clearly aimed at the same fundamental problems as this proposed U. S. effort in CPS. Artemis provides much more funding than we can hope to receive in the U. S. through NSF alone; it also closely and effectively couples academia and industry. The U. S. needs to respond appropriately with the adequate levels of funding, integration of industry, and multi-agency initiatives across the science/engineering spectrum. An important difference between the US and Europe is that European governments do not depend on industry to do all of the commercial research. Government labs do commercial, as well as military, research. The U.S. is doing less basic research in military projects. Trickle down technology, which gave us the Internet, microwave ovens, and Teflon, is not working any longer.

A significant challenge that CPS proposals must address is breaking down barriers between disciplines; different disciplines have fundamentally different thought processes. CPS crucially needs to promote and support multi-disciplinary work through: (1) grants big enough to support pairs, teams, etc.; (2) interactive program management that encourages research teams to avoid group-think and to teach each other their disciplines; (3) themes designed to bring multiple disciplines together; (4) industry stakeholders to support activity beyond initial investment by NSF and other government agencies; (5) new opportunities for spin-off capabilities and products through industry participation (e.g., like the NSF GOALI program but extended to other agencies); and (6) community-wide access to platforms and infrastructure through which new CPS problems can be discovered, along with solution approaches to existing known problems.

Ways should be sought to incorporate experts from the social sciences. Social scientists think in fundamentally different ways than computer scientists and engineers. CPS needs to find a common semantic basis with which to communicate across multi-disciplinary divides (a scientific Rosetta stone?).

Education also needs to be addressed. Education is extremely important in terms of creating a workforce capable of dealing with the CPS that we propose building. The current status of U.S. education is poorly prepared to train the next generation of CPS engineers and scientists. An NSF-sponsored initiative should have substantial impact at all levels of US educational system, ranging from secondary school to graduate level instruction.

Rethinking education requires an emphasis on engineering at the secondary school level as well as a revised emphasis on integrating computer science with the traditional engineering sciences at both the undergraduate and graduate levels. Revision of undergraduate and graduate curricula should be encouraged to prepare students to meet the demands of the CPS research priorities and roadmaps for such revisions should be established. Forums to exchange ideas and approaches to curriculum changes should be created and wide dissemination of the results should be encouraged.

At the graduate level, we expect research projects in CPS to directly result in new graduate courses whose materials can be used at other universities. These curriculum efforts could be aided by the creation of consortia that develop and share material, providing an early audience and motivation for improving material. The participants believe that summer schools, such as the ones commonplace in Europe, would help train both faculty and graduate students in this emerging field. As the science and technology of CPS mature, we expect to see a new generation of undergraduate courses that simultaneously consider engineering and computing issues. Some of this material should ultimately make its way to technology courses in high schools, replacing traditional (and often dated) microprocessor-based systems design courses with some basic elements of CPS technology.

Appendix – Summit Organization and Participants

NSF Sponsors

Dr. Jeannette Wing	Assistant Director, CISE Directorate, NSF
Dr. Taieb Znati	Director, CNS Division, NSF
Dr. Usha Varshney	Director, EECS Division, NSF
Dr. Helen Gill	Program Director, CNS Division, NSF
Dr. Scott Midkiff	Program Director, EECS Division, NSF

Organizing Committee

Bruce Krogh	Carnegie Mellon University
Edward Lee	University of California at Berkeley
Insup Lee	University of Pennsylvania
Al Mok	University of Texas at Austin
George Pappas	University of Pennsylvania
Raj Rajkumar	Carnegie Mellon University
Harvey Rubin	University of Pennsylvania
Alberto Sangiovanni-Vincentelli	University of California at Berkeley
Lui Sha	University of Illinois at Urbana Champaign
Kang Shin	University of Michigan at Ann Arbor
Jack Stankovic	University of Virginia at Charlottesville
Janos Sztipanovits	Vanderbilt University
Wayne Wolf	Georgia Tech
Wei Zhao	Rensselaer Polytechnic Institute

CPS Summit Participants

Prof. Tarek Abdelzaher, University of Illinois at Urbana Champaign
Prof. Rajeev Alur, University of Pennsylvania
Dr. Hugo Andrade, National Instruments
Prof. Panos Antsaklis, University of Notre Dame
Prof. John Baras, University of Maryland at College Park
Prof. Calin Belta, Boston University
Prof. Michael Branicky, Case Western Reserve University
Dr. Ken Butts, Toyota
Prof. Roy Campbell, University of Illinois at Urbana Champaign
Prof. Christos Cassandras, Boston University
Prof. Ed Clarke, Carnegie Mellon University
Prof. Munther Dahleh, MIT
Prof. John Doyle, Caltech
Prof. Nikil Dutt, University of California, Irvine
Prof. Stephen Edwards, Columbia University
Prof. Magnus Egerstedt, Georgia Tech
Prof. Deborah Estrin, University of California at Los Angeles

Prof. Eric Feron, Georgia Tech
Dr. Alessandro Forin, Microsoft Research
Prof. Robert Gao, University of Massachusetts at Amherst
Prof. Chris Gill, Washington University of St. Louis
Dr. Helen Gill, National Science Foundation
Prof. Steve Goddard, University of Nebraska at Lincoln
Prof. Susan Graham, University of California at Berkeley
Prof. Carl Gunter, University of Illinois at Urbana Champaign
Dr. Aarti Gupta, NEC Labs
Prof. Rajesh Gupta, University of California at San Diego
Prof. Jiawei Han, Univ of Illinois at Urbana Champaign
Dr. Walt Heimerdinger, Honeywell
Prof. Jessica Hodgins, Carnegie Mellon University
Prof. Seth Hutchinson, University of Illinois at Urbana Champaign
Prof. Marija Ilic, Carnegie Mellon University
Dr. Clas Jacobson, United Technologies
Dr. Ray Kamin, Rockwell Collins
Prof. Takeo Kanade, Carnegie Mellon University
Ms. Frankie King, NCO/NITRD
Dr. Ralph Kling, Crossbow Technology
Prof. Xenofon Koutsoukos, Vanderbilt University
Prof. Bruce Krogh, Carnegie Mellon University
Dr. Sri Kumar, BAE Systems
Prof. Steve LaValle, Univ of Illinois at Urbana Champaign
Prof. Edward Lee, University of California at Berkeley
Prof. Insup Lee, University of Pennsylvania
Prof. John Lehoczky, Carnegie Mellon University
Dr. Michael Lemmon, University of Notre Dame
Dr. Matthew Mason, Carnegie Mellon University
Prof. Nicholas Maxemchuk, Columbia University, IMDEA Networks
Dr. Paul McLaughlin, Honeywell
Prof. Bruce McMillin, Missouri University of Science & Technology
Dr. J. Michael McQuade, United Technologies
Dr. Scott Midkiff, National Science Foundation
Prof. Al Mok, University of Texas at Austin
Dr. Pieter Mosterman, The MathWorks
Prof. Jose' Moura, Carnegie Mellon University
Prof. Klara Nahrstedt, University of Illinois at Urbana Champaign
Prof. George Pappas, University of Pennsylvania
Dr. Youngchoon Park, Johnson Controls, Inc
Prof. Raj Rajkumar, Carnegie Mellon University
Prof. Ari Requicha, Univ. of Southern California
Dr. Harvey Rubin, University of Pennsylvania
Prof. Bill Sanders, Univ of Illinois at Urbana Champaign
Prof. Alberto Sangiovanni-Vincentelli, University of California Berkeley
Prof. Andreas Savvides, Yale University

Prof. Doug Schmidt, Vanderbilt University
Prof. Lui Sha, U of Illinois at Urbana Champaign
Prof. Kang Shin, The University of Michigan
Prof. Dan Siewiorek, Carnegie Mellon University
Prof. John Stankovic, University of Virginia
Prof. Bozidar Stojadinovic, University of California at Berkeley
Prof. Janos Sztipanovits, Vanderbilt University
Prof. Paulo Tabuada, University of California at Los Angeles
Dr. Peter Tufano, BAE Systems
Prof. Jeffrey Trinkle, Rensselaer Polytechnic Institute
Dr. David Waltz, Columbia University
Mr. Ben Watson, Lockheed Martin
Dr. Jeanette Wing, National Science Foundation
Prof. Brian Williams, Massachusetts Institute of Technology
Prof. Kensall Wise, University of Michigan at Ann Arbor
Prof. Wayne Wolf, Georgia Institute of Technology
Dr. Feng Zhao, Microsoft Research
Prof. Wei Zhao, Rensselaer Polytechnic Institute
Dr. Ty Znati, National Science Foundation