

Cyber-Physical Systems

Executive Summary

Prepared by the CPS Steering Group

March 6, 2008

Introduction

The integration of physical systems and processes with networked computing has led to the emergence of a new generation of engineered systems: Cyber-Physical Systems (CPS). Such systems use computations and communication deeply embedded in and interacting with physical processes to add new capabilities to physical systems. These cyber-physical systems range from miniscule (pace makers) to large-scale (the national power-grid). Because computer-augmented devices are everywhere, they are a huge source of economic leverage.

While most people think of “computers” as PCs and “computing” as browsing on the World Wide Web, most of the computers in the world are components of cyber-physical systems. The share of electronics in the cost of final products has been increasing dramatically, according to a recent study¹ conducted by the European Commission. This study argues that in automotive, avionics/aerospace, industrial automation, telecommunications, consumer electronics, intelligent homes, and health and medical equipment, electronics will reach 53% of the cost by the end of the decade. The depth of this change is well-reflected by the trends in automobiles: in 1990 the percentage of value in an automobile’s worth of electronics was 16%; in 2003, it reached 52% and by the end of this decade, it will grow to 56%. This shift of the center of gravity in networking and information technology (NIT) has been recognized by an August 2007 report² of the President’s Council of Advisors on Science and Technology (PCAST), which calls for the restructuring the national priorities in NIT research and development and placing CPS on the top of the list.

Why do we need so many cyber-physical systems? Generally speaking, embedded computers allow us to add capabilities to physical systems that we could not feasibly add in any other way. An early example of the successful marriage of computers and mechanical systems is the automobile. The advent of computer-controlled automotive engines is critical to both fuel-efficient and low-emission cars. It is unlikely that any car could be sold in the United States today that did not make use of computers to meet its fuel efficiency and pollution mandates. Clearly, the technological advantages brought about by marrying computers and physical devices have broad impact on the economy and society. By merging computing and communication with physical processes and mediating the way we interact with the physical world, cyber-physical systems bring many benefits: they make systems safer and more efficient; they reduce the cost of building and operating these systems; and they allow individual machines to work together to form complex systems that provide new capabilities.

¹ Study of Worldwide Trends and R&D Programmes in Embedded Systems in View of Maximising the Impact of a Technology Platform in the Area. Prepared for the European Commission, November 18, 2005

² Leadership Under Challenge: Information Technology R&D in a Competitive World. An Assessment of the Federal Networking and Information Technology R&D Program. President’s Council of Advisors on Science and Technology (PCAST), August 2007

Many of the embedded systems-related studies and efforts in the past have focused on the challenges the physical environment brings to the scientific foundations of NIT. However, the full scope of this change has much more breadth and depth than a restructuring inside NIT; it is a profound revolution that turns entire industrial sectors into producers of cyber-physical systems. This is not about adding computing and communication equipment to conventional products where both sides maintain separate identities. This is about merging computing and networking with physical systems to create new capabilities and improve/maintain product quality. CPS has extraordinary significance for the future of the U.S. industry. There is much more at stake than extending our leadership in NIT to an exploding new market segment. Falling behind in the foundations of CPS may render our scientific and technological infrastructure obsolete, leading to rapid loss in our competitiveness in major industrial segments including automotive, aerospace, defense, industrial automation, health/medical equipment, critical infrastructure and defense. Whether we recognize it or not, we are in the midst of a pervasive, profound shift in the way humans engineer physical systems and manage their physical environment.

This document serves as a brief introduction to our current understanding of cyber-physical systems: why they have emerged now; how they will affect not only industry but society as well; the challenges and scientific agenda for this emerging field; and some recommendations on the ways that industry, government, and academia can cooperate to precipitate these changes. Our analysis and recommendations are in full concurrence with the main conclusion of the August 2007 PCAST report that presented a formal assessment of the Federal Networking and Information Technology R&D (NITRD). We agree that the current U.S. leadership position in NIT will not translate into leadership in the future. The profound shift toward CPS brings NIT to the front line of a new technological revolution. A new national effort is required to capitalize on this opportunity and to establish U. S. leadership in this area.

Methodology

Our conclusions have been developed using inputs from the following National Workshops initiated by the National Science Foundation between 2005 and 2008. The detailed program, contributions and conclusions of these Workshops are available on their respective web sites.

- [National Workshop on "High Confidence Medical Device Software and Systems \(HCMDSS\)", June 2 - 3, 2005, Philadelphia, PA.](#)
- [National Workshop on "Aviation Software Systems: Design for Certifiably Dependable Systems", October 5-6, 2006, Alexandria, TX.](#)
- "Aviation Workshop" Report attached below.
- [NSF Workshop on "Cyber-Physical Systems", October 16-17, 2006, Austin, TX.](#)
- "CPS Workshop" Report attached below.
- [National Meeting on "Beyond SCADA: Networked Embedded Control for Cyber Physical Systems";, November 8-9, 2006, Pittsburgh, PA.](#)

- "Beyond SCADA" Workshop Report attached below.
- [National Workshop on "High-Confidence Software Platforms for Cyber-Physical Systems \(HCSP-CPS\)"](#), November 30 - December 1, 2006, Alexandria, VA.
- [NSF Industry Round-Table on Cyber-Physical Systems](#), May 17, 2007, Arlington, VA.
- ["Joint Workshop On High-Confidence Medical Devices, Software, and Systems \(HCMDSS\) and Medical Device Plug-and-Play \(MD PnP\) Interoperability"](#), June 25-27, 2007, Boston, MA.
- [National Workshop on "Composable and Systems Technologies for High-Confidence Cyber-Physical Systems"](#), July 9-10, 2007, Arlington, VA.

Besides the NSF Workshops, we considered the conclusions of several studies and new programs initiated in other countries, particularly the following studies undertaken in the European Union, Japan and South Korea.

- European Union: [Advanced Research and Technology for Embedded Intelligence and Systems \(ARTEMIS\)](#)
- European Union: [Smart System Integration \(EPoSS\)](#)
- Japan: While information about government initiatives are not available, the level of interest is well demonstrated by the fact that the world's largest tradeshow and conference in embedded systems is organized in Japan ([Embedded Technology 2008](#)) with 26,646 total registrants and 10,000 conference attendees.
- South Korea: [Korean IT Industry Promotion Agency \(KIPA\)](#) - Embedded systems programs

Importance of CPS for the United States

It is critical that the United States cultivate and maintain a lead in the design of cyber-physical systems. Cyber-physical systems are an emerging trend around the world because of fundamental technological and economic forces. Cyber-physical systems are the primary area where disruptive technologies emerge that create new industries and rearrange the status quo in entire industrial sectors. As a global leader in NIT, the U.S. is well positioned to capture the initiative and gain technology advantage in CPS. This advantage does not translate immediately into advantage in CPS, however. This is well recognized by our competitors. For example, between 2007 and 2013 the ARTEMIS Program in EU will invest seven billion in mid-2007 dollars in R&D to achieve "world leadership in intelligent electronic systems" by 2016.

A series of technological and economic drivers have aligned themselves to reshape industry. Many of these are moving fast and the U.S. can maintain its lead only by moving quickly. If the U.S. allows other countries to develop capabilities in cyber-physical systems that we lack, then our industrial competitiveness will suffer long-term economic harm that will be very difficult to reverse.

Technological and Economic Drivers

The last two decades have brought a digital revolution that has been transforming industry. This change is not a matter of choice; it is driven by fundamental, long-term technological and economic trends that we expect to continue with or without our active participation.

- The decreasing cost of computation, networking, and sensing provides the basic economic motivation for adopting NIT in every industry and application. Moore's Law, which mandates an exponential growth in computing power, has brought us extremely sophisticated computers at consumer electronics prices. The same trends have vastly improved communication and sensing. Computers and communication have become the "universal system integrator" that keeps large systems together. They enable the construction of cyber-physical system infrastructures that run on a national or global scale: the national power grid, the air traffic control system, the national transportation network.
- A variety of social and economic forces will require us to use our national infrastructure more efficiently. Building new roads and power lines, for example, both costs money and introduces environmental and local impacts that we may not want to absorb. By operating vehicles at closer spacings and tolerating smaller margins in the power grid we can use more of the available capacities. By monitoring the conditions of roadways, we can apply maintenance where it is needed, both improving utilization and reducing unnecessary maintenance.
- Environmental pressures will mandate the rapid introduction of technologies to improve energy efficiency and reduce pollution. As we have seen in automobiles, improving energy efficiency and reducing pollution simultaneously is possible only by using the complex control systems that embedded computers allow us to build.
- As the national population ages, we will need to make more efficient use of our health care systems, ranging from facilities to medical data and information. This requires shift toward increased automation in treatment management, in-home health care delivery, and the use of standardized treatment protocols.

Industrial Competitiveness

The technological and economic drivers are creating an environment that both enable and require a range of new capabilities. Progress during the past decade has produced early examples of a new generation of systems that rely on cyber-physical technology:

- Airplanes and automobiles that are environmentally friendly and energy efficient.
- Integrated, self optimizing transportation systems and vehicles that are able to interface with them.
- Advanced health care via increased automation, integrating smart devices, and providing safe access to electronic medical records.
- New biotechnology via engineering closed-loop biological systems.
- Quality, safety and efficiency of our national infrastructure.

Appearance of these new capabilities creates a competitive pressure for all players of the global economy. However, answering these pressures requires significantly more than inserting new technologies into traditional industries. We can compete effectively and capitalize on the power of combining computers and physical systems only if we rethink mass production for the 21st century.







In the late 19th and early 20th centuries, advances in materials (steel) and locomotive power (electricity, internal combustion engines) drove a host of advances that rippled through society to create the modern industrial society that we take for granted today. In order to create the cars, appliances, etc. that are mainstays of everyday life, an entire culture of mass production has been developed. This culture ranges from the nature of the components that we use (bolts and nuts, motors, etc.) through the training that industrial workers receive all the way to the organization of companies and industries.

Cyber-physical systems provide a host of opportunities, some of which will be obvious to the end user and some of which will not. An example of a possible way in which CPS may invisibly change the nature of industry is in the changing nature of a *component*. As outlined in the book *The Machine That Changed the World* (Womack et al.), car manufacturers before Ford built hand-crafted cars, not two of which were alike. Their methods could not scale to large quantities; increasing volume was critical to lowering the price of the automobile. Twentieth-century mass production has relied on standardization not only to reduce cost but also to make it possible to complete basic industrial tasks. When we use CPS components, it is entirely possible that some components could adapt themselves automatically to the other components in their assembly, which inevitably changes the way in which these CPS-enabled components are designed and manufactured. To take full advantage of cyber-enabled physical systems, we need to rethink industry and manufacturing for the 21st century:

- We need a new vocabulary of components that we use and re-use to build complex systems. Since reusable components are the foundation for commoditization that drive the formation of supply chains, the resulting restructuring will be far reaching in all industrial sectors.
- We need new methods for designing and testing systems in which physical characteristics are determined (or influenced) computationally and the other way around. Consider, for example, the way that we work on cars today as compared to how mechanics worked on cars 30 years ago. In the 1970s, mechanics were just that---using wrenches and screwdrivers as well as their own eyes and ears to adjust how well the car ran. Today, the only way to deal with modern computer-controlled engines is to plug it into another computer and deal with it through a screen, mouse, and keyboard.
- We need a workforce with new skills that are adapted to cyber-physical systems. As the complexity of produces is increasingly concentrated in software and electronics, the traditional disciplinary boundaries need to be realigned in all levels of training and education.

Opportunities Created by Cyber-Physical Systems

Cyber-physical technology can be applied in a wide range of domains, offering numerous opportunities in products:

Sectors	Opportunities	
<i>Transportation</i>	Aircraft that fly faster and further on less energy. Air traffic control systems that make more efficient use of airspace. Automobiles that are more capable and safer but use less energy.	
<i>Defense</i>	More capable defense systems; defense systems that make better use of networked fleets of autonomous vehicles.	
<i>Energy and Industrial Automation</i>	New and renewable energy sources. Homes, office, buildings and vehicles that are more energy efficient and cheaper to operate.	
<i>Health and Biomedical</i>	In-home healthcare delivery. More capable biomedical devices for measuring health. New prosthetics for use within and outside the body. Networked biomedical systems that increase automation and extend the biomedical device beyond the body.	
<i>Agriculture</i>	Energy efficient technologies. Increased automation. Closed-loop bioengineering processes. Resource and environmental impact optimization. Improved safety of food products.	
<i>Critical Infrastructure</i>	Highway systems that allow traffic to become denser while also operating more safely. A national power grid that is more reliable and efficient.	

In summary, we need to invest in cyber-physical systems because they are the foundation for the 21st-century industry. Societal needs, as well as competitive pressure from other countries that have dramatically increased their investment in CPS related research, mandate rethinking our national R&D investment strategy. Investment in CPS underpins all design and manufacturing industries as well as health care, agriculture, and infrastructure. By investing in underlying technology of cyber-physical systems, we make United States industries more competitive across the board.

Science and Technology Challenges and Recommendations

In order to discuss the science and technology challenges facing cyber-physical systems, let us describe CPS more precisely. A CPS is a system:

- in which computation/information processing and physical processes are so tightly integrated that it is not possible to identify whether behavioral attributes are the result of computations (computer programs), physical laws, or both working together;
- where functionality and salient system characteristics are emerging through the interaction of physical and computational objects;
- in which computers, networks, devices and their environments in which they are embedded have interacting physical properties, consume resources, and contribute to the overall system behavior.

Although the 20th-century science and technology has provided us with effective methods and tools for designing both computational and physical systems, the design of cyber-physical systems is much more than the union of those two fields. Traditionally, information scientists have had only a hazy notion of the requirements imposed by the physical environment attached to their computers. Similarly, mechanical, civil, and chemical engineers have viewed computers strictly as devices executing algorithms and ignored the physical properties of the embedded computing platforms. To the extent that we have designed cyber-physical systems, we have done so in an ad hoc, one-off manner that is not repeatable.

A new science of cyber-physical system design will allow us to create new machines with complex dynamics and high reliability; it will allow us to be able to apply the principles of cyber-physical systems to new industries and applications in a reliable and economically efficient way. Progress requires nothing less than the reintegration of physical and information sciences---the construction of a new science and technology foundation for CPS that is simultaneously physical and computational.

Shortcomings of the Current Science and Technology Infrastructure

The past fifteen years provided ample evidence that the separation of information science and physical science has created a divergence in scientific foundations and technologies that has become strongly limiting to progress in the design of CPS. For example, dominant abstractions in programming languages typically avoid the explicit representation of time and other aspects of physicality, lumping together all the related physical design considerations in the category of “non-functional” requirements (such as timing, power and reliability). On the physical side, although engineering increasingly relies on computer based implementations, systems science has developed and evolved abstractions that largely neglect salient properties of computing and communication platforms (such as scheduling, resource management, network delays) and considers those as secondary implementation issues. The resulting barrier between systems and computer science has kept the research communities apart, splitting education into isolated disciplines and resulting in compartmentalized design flows that ultimately lead to major difficulties and failures as complexity has increased.

Current industrial experience tells us that, in fact, we have reached the limits of our knowledge of how to combine computers and physical systems. If we continue to build systems using our very limited methods and tools, not only will we do so very inefficiently, we run the risk of endangering many people with unsafe and unpredictable systems. These shortcomings range from technical limitations in the foundations of our understanding of cyber-physical systems to the way we organize our industries and educate engineers and scientists that support cyber-physical system design. Listed below are indicators for the need of new foundations that show it is doubtful that they can be solved by simply combining existing theories and methods from either side.

Composition is a technical foundation of all engineering disciplines---it helps us manage complexity, decrease time-to-market, and contain costs. The feasibility of component-based system design depends on two key conditions: *compositionality* – meaning that system-level properties can be computed from local properties of components – and *composability* – meaning that component properties are not changing as a result of interactions with other components. Lack of compositionality leads to brittleness, that is, systems that don’t behave well outside of a small operational envelope and that are hard to maintain.

The foundations of component-based design are well understood and successfully applied in many engineering and computer science disciplines, such as digital logic in computer engineering, linear dynamics in control engineering, or process algebra modeling some aspects of distributed computing, to name a few. The common feature of all successful compositional design frameworks is *homogeneity* in terms of the properties composed and the semantic framework used in modeling.

Cyber-physical systems are inherently heterogeneous not only in terms of their components but also in terms of essential design requirements. In addition to functional

properties, CPS are subject to a wide range of physical requirements, such as dynamics, power, physical size and to systems-level requirements, such as safety, security and fault tolerance. This heterogeneity does not go well with current design methods and practices for several reasons. The most important principle used in managing multi-objective design problems is *separation of concerns* (in other words, defining design viewpoints). Separation of concerns works if the design views are orthogonal, i.e., design decisions in one view do not influence design decisions in other views. This is clearly not working in CPS due to the complex interactions across different system layers and design views that are frequently not modelled.

Design automation is embodied in computer-aided design tools, but our current limited understanding of how to build tools for the design of cyber-physical systems has led to an inefficient, application-by-application organization of design automation companies. Today's application-oriented organization of the design automation industry has led to expensive, monolithic tool chains that support only a narrow range of problems. These stovepiped tool chains do not scale to the large, heterogeneous systems that we need to build using CPS techniques.

CPS products are heterogeneous systems comprising multiple types of physical systems and multiple models of computation and communication. This heterogeneity makes design flows increasingly product specific – which is bad news for design automation. It implies that increasing design complexity goes hand-in-hand with increasing product specificity, and this means a decreasing and highly volatile market for specialized design tool chains - a prohibiting factor for the emergence of a viable tool industry. This is a serious problem because lack of effective, specialized design automation tools limits design productivity and keeps the all important time-to-market high in the most aggressively changing industrial sectors.

If computers are universal devices, then we should be able to leverage the tools and processes used to design embedded computing systems across the broad range of cyber-physical systems, from medical devices to air traffic control systems. Design approaches have similarity in all engineering disciplines: design flows include precisely-defined layers of abstractions. The emerging design is expressed in the form of abstract models that are analyzed and used for verifying required properties and the models are incrementally refined, composed and transformed until sufficient details are reached for generating (manufacturing) the engineering artifact. Today's standard is to use different processes and tools, however, for different applications. This means that much money and time is wasted in duplicative efforts. It also means that many tools are not developed because a single stovepiped application does not provide a large enough market to justify investment in tool development.

System integration is the elephant in the room of large-scale system design. We would like to think that we can specify components, design them independently, and then easily plug them together to create complex systems. Unfortunately, history is full of examples that show we cannot reliably integrate complex components into complex systems. Our current technology cannot provide predictability for partially compositional properties,

which is a common situation in all large scale system development. We are so far from a science of system integration that we generally consider this a management problem divorced from science and engineering; systems integration is almost totally absent from the engineering and computer science curricula.

Certification is a key problem for safety critical systems. Today, we only know how to certify computer-based systems by first building them and then testing them (sometimes to the point of destruction). Current certification methods, by-and-large, resort to standards (such as DO-178B) that impose requirements on the development process and require testing-based evidence. This approach does not scale and is inherently incomplete. Consequently, the cost of certification of complex systems (airplanes, medical equipment, etc.) is cripplingly high and extremely hard to bind in the preliminary design phase.

The cost and risk pushes companies to stick with older technologies and to test systems “until the money runs out”. The desired solution would be compositional certification. This would make it possible to certify the certification of system components (physical and cyber) separately without the need of re-certifying them after they are composed into systems. Compositional certification works well in physical systems (using reliability analysis) but breaks down in CPS. It is a generally accepted position that lack of compositionality in CPS certification has its roots in system design and development technologies. Existing compositionality properties of physical systems and compositionality results for software (e.g. assume-guarantee reasoning methods) lose applicability when cyber and physical systems are composed.

Security and privacy are necessary for both economic security and quality of life. Cyber-physical systems open up new threats---physical systems can now be attacked through cyberspace and cyberspace can be attacked through physical devices. Consider, for example, an attacker who modifies the software in automobile engine controllers so that all cars stop at a certain time; such an attack would cause huge accidents and would also impede the emergency vehicles destined to respond to the accidents.

One example of the way that cyber-physical systems change the rules of dependable system design is in the nature of the faults. In physical systems, we generally assume *parsimony*---that physical failures are independent and that multiple, simultaneous failures are unlikely. Computers and networking added to physical systems introduce *reverse parsimony*, however. Cyberattacks can be massive, introducing simultaneous or precisely coordinated failures. Distributed denial-of-service attacks, for example, rely on using large number of “zombie” computers to attack a target in such a way that the system under attack cannot easily recover. Even when failures are inadvertent, replicated software can cause systematic failures that are not common in purely physical systems. Furthermore, today’s computer systems do not allow us to distribute computer-based control in ways that preserve reliability; network services that have been designed to support multimedia and entertainment are not adequate foundations for the development of secure cyber-infrastructure. Our knowledge of the types of failures that can occur in

cyber-physical systems is totally inadequate; our knowledge of how to protect ourselves from these failures is even farther behind.

Education is another shortcoming due to today's stovepiping of engineering versus computer science. We are failing to train a new generation of cyber-physical systems workers at all levels---Ph.D. researchers, design engineers, and maintenance personnel. Today's academic requirements and organizational structures lead workers and students away from obtaining the set of skills required to be successful CPS experts.

The currently dominant forms of **government/industry/academic interactions** have severely limited our effectiveness to answer the challenges of emerging cyber-physical systems technology. This new area, by its very nature, requires experimental research and challenge problems generated by industry and government labs. The Federal government classifies research into categories---6.1 for basic research, 6.2 for applied research, 6.3 for development. This classification and the related restrictions virtually guarantee that academic and industrial groups work in isolation. This makes technology transitioning excruciatingly slow and practical validation of academic research results very rare.

Science and Technology Challenges

Although cyber-physical systems cover an extremely wide range of application areas, we can create a new scientific and technological understanding of the interactions of information processing, networking and physical processes. This new science of cyber-physical systems will have broad application, aided by specializations and additions for particular application domains. This new science of CPS will allow us to design systems more economically by sharing both abstract knowledge and concrete tools. It will also allow us to design more dependable cyber-physical systems, since we can apply best practices to the entire range of cyber-physical applications.

In order to meet the challenge of cyber-physical system design, we need to create a new systems science foundation and new technology infrastructure that is not wholly derived from either computer science or engineering. This new discipline will merge fundamental concepts from each and will inject new ideas of its own. The creation of this new science and technology of cyber-physical systems can be guided by the following major challenges:

1. We need to **realign abstraction layers in design flows**. Computational abstractions need to include physical concepts, such as time and energy. Abstractions developed for describing physical dynamics should be extended to capture uncertainties of implementation platforms, such as network delays, finite word length and round-off errors. These changes in abstractions layers will allow the synthesis of computations with physical properties and physical system dynamics that are robust against implementation uncertainties.
2. We need to develop **semantic foundations for composing heterogeneous models and modeling languages** describing different physics and logics. We need to introduce or develop mathematical frameworks that make semantics not

- only mathematically precise, but also explicit, understandable and practical for system developers as well as tool developers.
3. We need to develop new understanding of **compositionality in heterogeneous systems** that allows us to take into account both physical and computational properties. This new view of compositionality will allow us to create large, networked systems that satisfy essential physical properties and deliver the required functionality in a reliable way.
 4. Cyber physical systems will have properties for which achieving full compositionality would be expensive or impractical. Development of technology for achieving **predictability in partially compositional properties** is a hard problem that must be addressed.
 5. We need a science and technology **foundation for system integration** that is model-based, precise, and predictable. Transforming system integration from a high risk engineering practice into a science-based engineering discipline is a huge challenge that will require close collaboration between industry and academy.
 6. We need new theories and methods for **compositional certification** of cyber-physical systems. We must be able to compose CPS components into a large CPS system in such a way that we can reuse the certification of the components as evidence in certifying the larger system. Certification should rely more on verification and less on testing.
 7. We need a new infrastructure for **agile design automation** of cyber-physical systems. As new application domains of cyber-enhanced physical systems appear, we must be able to rapidly adapt our existing tool base to help us design those systems. If we must wait for tools to be created before we can jump into a new industry, we risk losing the lead to competitors who can use either agile tool chains or massive amounts of labor to work through design problems.
 8. We need to develop new **open architectures for cyber-physical systems** that will allow us to build national-scale and global-scale capabilities. These architectures should be defined by policies controlling their evolution, invariants that need to be maintained - and not by static structures so that they can be more easily adapted to different operational conditions.
 9. We need architectures and tools that allow us to build **reliable CPS systems** from unreliable components.
 10. We need architectures and tools that allow us to build **resilient CPS systems** that can tolerate malicious attacks from either the cyber or physical domains. These architectures should leverage open systems technologies to reduce design times and increase confidence.

Addressing the challenges above requires a research agenda that crosses current disciplinary boundaries and offers benefits to stakeholders that span a broad range of application domains. We believe that general applicability is not contradictory to the need for establishing a focused science and technology program.

The following table summarizes how these challenges map to important examples of cyber-physical systems (■: very important; □: important; ○: not as important):

	New abstraction layers for design	Semantic foundations for composing models	Composition platforms for heterogeneous systems	Predictability under limited compositionality	Foundation for system integration	Compositional certification	Agile design automation	Open Architectures	Reliable systems from unreliable components	Resiliency to cyber attacks
<i>Aerospace</i>	■	■	■	□	■	■	■	○	■	■
<i>Automotive</i>	■	■	■	□	■	○	■	□	□	■
<i>Defense</i>	■	■	■	□	■	■	■	■	□	■
<i>Energy</i>	■	□	■	■	■	○	○	■	■	■
<i>Biomedical</i>	■	□	■	■	■	■	□	□	■	■

Summary of Recommendations

In order to ensure that the U. S. reaps the benefits of leadership in cyber-physical systems technology, we recommend the establishment of a National Initiative for Cyber-Physical Systems (NICPS) based upon the following goals:

- Becoming a leader in the science of cyber-physical systems.
- Creating an educational infrastructure that imparts knowledge of cyber-physical systems at all levels of education.
- Becoming a leader in the application of cyber-physical systems science to industry.
- Building the world's most advanced system of cyber-physical infrastructure that will efficiently tie together the nation's economy.

We recommend several actions to promote both the development of a science of cyber-physical systems and its dissemination into industry and society at large:

- A coordinated research program that encompasses multiple agencies in order to develop the basic scientific principles of cyber-physical systems as well as new ways to integrate them into industrial practice.
- A public/private partnership that includes investment by both industry and government.

- New organizational and funding framework for multidisciplinary and multi-institutional academic-industry centers that tightly integrate academic excellence with industry-strength challenges and enable long-term, large-scale, project-driven R&D.
- Transformation of our educational system to teach students at all levels who understand the principles of cyber-physical systems appropriate to their jobs.
- A new system of incentives for academic research that measures and rewards both academic excellence and technological impact and encourages universities to become more open and flexible to an ongoing realignment of disciplinary boundaries.
- A coordinated research agenda for government, academia, and industry that is organized around the fundamental challenges posed by cyber-physical systems.
- International collaborations that leverage the U.S. position as a magnet for talent.

This agenda will require significant and sustained investments of effort and money by government, academia, and industry:

- Investment in K-12 education is needed to train students with improved computational skills and the ability to integrate those skills with the traditional vocational arts.
- Colleges and universities will need to revamp curricula to include CPS concepts. All engineering and computer science students should receive some training in cyber-physical systems. We must also create a new cadre of cyber-physical systems specialists who work on both architectures and methodologies.
- Universities need to produce at least 1,000 new Ph.D.s in cyber-physical systems in the next decade. In order to insert this new discipline into the broad range of industries that it affects, we will need huge numbers of highly trained engineers who can pave the way for lower-level workers with CPS skills.
- Individual industries must invest heavily to understand how to leverage CPS technologies to create advanced systems. Some of this investment will be pre-competitive and could be spread across all industry players; other parts of the investment will be specific to companies and proprietary.
- Government agencies will need to invest in cyber-physical infrastructure that operates more efficiently and reliably. This investment will both serve as an example of CPS technologies and will enable further use of cyber-physical techniques in other parts of the economy.
- Government, industry, and academia should co-operate to create test beds to develop and prove CPS components, along with centers that will both develop and disseminate key CPS technologies.

Members of the Steering Group

- Bruce H. Krogh, Carnegie Mellon University
- Edward Lee, UC Berkeley
- Insup Lee, University of Pennsylvania
- Al Mok, UT Austin
- George Pappas, University of Pennsylvania
- Raj Rajkumar, Carnegie Mellon University
- Lui Raymond Sha, UIUC
- Alberto Sangiovanni Vincentelli, UC Berkeley
- Kang Shin, University of Michigan
- Jack Stankovic, University of Virginia
- Janos Sztipanovits, Vanderbilt University
- Wayne Wolf, Georgia Institute of Technology
- Wei Zhao, RPI

List of Related Documents

A set of presentations from workshops held to develop the cyber-physical systems agenda can be found at CpsNCO Web site

(<http://ike.ece.cmu.edu/twiki/bin/view/CpsNCO/WebHome>).