Response to RFI for Five-Year Strategic Plan for the Federal NITRD Program

Q. What do you imagine as the future in terms of desired NIT capabilities?

While NIT advances over the past two decades have brought about tremendous improvements in individual productivity and organizational efficiency, their impact on the societal infrastructure – for transportation, health-care, electricity/water distribution – has been marginal. There exists a tremendous potential for substantial reengineering of the societal infrastructure – related to movement, distribution, awareness and participation – that can be made possible by the emerging NIT capabilities. The interdependencies between these systems and their adaptabilities provide opportunities for resource conservation, better quality of life and improved societal dialogue & conduct. Crude examples based on extrapolation of what we see today include intelligent highway systems (versus single, independently-operated vehicles), dynamically configured and integrated intensive care or emergency transport units (versus separately-operated sensors or support systems), and intelligent home and industrial environments (versus traditional heating, lighting, and air conditioning subsystems). Much more is possible when the new technological underpinnings have been devised.

But these advances will not happen as a matter of entirely proliferation of commercial technologies. To be sure there are fundamental challenges: scientific, societal and economic. Scientifically, a whole new engineering discipline is needed to deal with phenomena at the intersection of logical and physical processes, societal needs for policy, equity and fairness (especially as infrastructure investments involve use of public funds.)

Q. What roles do you imagine for the NITRD Program and for the academic, commercial, international, and other domains in achieving that future?

The funding for S&T has undergone dramatic changes due to changes in the agency priorities. Even as the changes in individual funding sources have been gradual their cumulative impact is beginning to challenge the fundamental structure and assumptions underlying the federal (and increasingly state) outlay for supporting S&T research. One of the major pillars that has come under increased stress and questioning is the NSF directed research as rooted in the NRC act, and going back to its genesis in NDRC of 1940. The original concept leading up to formation of "Research University" was based on policy pillars spelled out in the act. The most important among these was the role of NSF in directing *unsolicited* research.

Current realities make such unsolicited research as one of the least attractive options for individual researchers; in many subject areas this is not even an option. In this environment, NSF as a research funding institution faces dramatic challenges in creating programs that are forward looking and truly reflective of the original spirit behind the NRC act that directly lead to the successful "research university" in the post world-war era. This success was in no small measure attributed to the leveraging effect of capability-seeking S&T 6.2 funding by DARPA as well as successful consortia funding such as Sematech and SRC. All of these institutions have undergone changes with the primary effect being the removal of their amplifying effect on the NSF/DOE sponsored research. NSF now faces the real danger of creating programs that are likely to mislead – rather than benefit from – the tremendous diversity of S&T talent in the nation; by creating artificial priorities and programs that do not reflect the true state of knowledge which is increasingly tied to the practice.

NITRD has a special role. NITRD has the capability to take a broader view of the state of knowledge and technology, and thus advise individual agencies of not only the opportunities that they should pursue, but also the moral imperatives emerging from such a broader view. As it does so, it must strive to overcome the barriers and mindsets regarding knowledge created across institutions, across industry and academia. In particular, it is no longer true that universities or a group of academic working directly in areas related to IT accurately represent the state of the art in knowledge; or that they know the best opportunities for advancing the state of knowledge. Indeed, even as industrial labs have changed, the state of the art – and visibility into such knowledge – has shifted to small-scale enterprise and their backers as in the leading edge venture capital community. This visibility into VC community is important because as a group they would not support investments necessary to advance societal infrastructure, and yet they provide – through their investments – important technology pieces that dramatically reduce the costs in retooling infrastructure through commercialization. Further and specifically in the context of advancing the state of the knowledge in *cyber-physical systems (CPS),* the programs should also seek to meaningfully engage public-works agencies, such as state DOTs, water resources management. These will be crucial to technology transition efforts into practice.

Q. Technical challenges that should be addressed

There are many scientific challenges underlying the theory and engineering needed to build CPS systems. Many of these have been spelled out in related reports by NSF organized workshops. To be specific, I will focus on one such challenge. Current CPS systems lack the ability to capture spatial information – information related to the location of actions as well as the use of location information in defining actions. While geographical location information can be 'stored' in various forms and at various levels, semantic support to use this information at various levels of the system implementation is severely lacking. When building CPS, this limitation manifests in many ways: from inadequately specified CPS functionalities and their validation to a lack of any guarantees related to availability or unavailability of computational resources as a function of location. Cyber-physical systems are at one level embedded sensor systems: they react to and manipulate spatiotemporal sensory information. *Yet these lack models and methods to capture such information, validate their behavior, performance against timing and spatial requirements.* This presents a fundamental barrier to the scaling and use of cyber-physical systems to societal-scale applications because a whole host of constraints, from energy, power, bandwidth, processing to resource availability, simply rule out the use of all the sensors at all times. To achieve the goal of semantic support for location and time at all levels, we need to address the following technical problems: (a) How do we capture location (and timing) information into CPS models that allows for validation of the logical properties of a program against the constraints imposed by its physical (sensor) interactions? (b) What are useful models for capturing faults and disconnections within the coupled physical-computational systems? How can we reason with these models to define the notion of system availability? (c) What kind of properties that can be verified, and assertions that can be ensured in applications that make use of both physical (real) time as well as location information? Do these propositions require direct algebraic support for location? How best these location and timing aware assertions can be validated? (d) What programming model is best suited for CPS applications utilizing dynamic behaviors? Are there any specific operating system or 'middleware' services that can ease the task of building such applications, and doing so reliably? (e) What are the metrics to measure effectiveness of physically-coupled embedded systems? How do we characterize operational efficiency with measures that take into account spatial information?

To answer these questions, our shared research challenges span the choice of abstractions (models and methods), programming models that use location information, to infrastructural support for location (virtualization, location determination/validation support ser-vices, etc). Promising projects to address these challenges will span programming, formal methods, distributed and embedded real-time systems, and whole host of disciplines that come together in building sensors and sensor network applications.