NEW VISTAS IN CIP RESEARCH AND DEVELOPMENT: SECURE NETWORK EMBEDDED SYSTEMS

Report of the NSF/OSTP Workshop on Innovative Information Technologies for Critical Infrastructure Protection

September 19-20, 2002

Workshop held in Leesburg, Virginia at the National Conference Center

Report prepared for
Dr. Helen Gill, CISE, National Science Foundation

by the conference organizers
Shankar Sastry, Jack Stankovic, and Janos Sztipanovits

With the assistance of the participants of the workshop (listed in Appendix C)

Acknowledgements

A big of vote of thanks are due to Ms. Erica Morrison for her organization of the event and Ms. Carmen Whitson for her recording of the minutes of the workshop.

Special thanks are due to the President's Office of Science and Technology and to the National Science Foundation for sponsoring the workshop and participating in it: We acknowledge especially the contributions and participation of Dr. Peter Freeman, Dr. Helen Gill, Dr. Carl Landwehr, Dr. Tom Greene, Dr. Mari Maeda, Dr. Darlene Fisher, Dr. Frank Anger, Dr. Kamal Abdali, Mr. Steve Mahaney, Dr. Darlene Fisher the National Science Foundation Directorate for Computer and Information Science and Engineering. We also acknowledge the participation of Mr. Richard Russell, Deputy Director of Office of Science and Technology Policy for technology, Ms. Sharon Hayes, and Dr. Simon Szykman. We owe Mr. Mark Leblanc from Office of Science and Technology Policy a special debt of gratitude for his help in planning this meeting. Additional thanks are due to the US State Department and Mr. Stan Riveles in particular for his help with organization. Dr. Sam Varnado from Sandia National Laboratories also has a special debt of gratitude in organizing the presentation of the talks on inter dependencies. Dr. Massoud Amin from the Electric Power Research Institute is to be lauded for his efforts resulting in a stimulating panel on protection of power systems. Dr. Art Pyster, Dr. Marshall Potter, Dr. Feisal Keblawi and Dr. Ernie Lucier coordinated an excellent series of presentations from the Federal Aviation Administration.

Dr. Doug Maughan of the Advanced Technology Office, Dr. Doug Schmidt of the Information Exploitation Office and Dr. Sri Kumar of the Information Processing Technology Office of The Defense Advanced Research Projects Agency provided a great deal of valuable input in terms of Department of Defense program plans, as did Dr. Steve King from the CIP program of the Office of Secretary of Defense. We also thank participants from the Office of Naval Research, and Department of Defense.

Finally we thank Dr. Cita Furlani and Dr. Sally Howe from the National Coordination Office for Information Technology Research and Development for their participation in the event.

Executive Summary

Long-range research in information technology is crucial to Critical Infrastructure Protection. Today's weak infrastructure is due in large part to the fact that traditional approaches to Digital Control Systems (DCS) and SCADA have not been brought up to the standards of modern information technology. The techniques commonly employed are ad hoc combinations of Proportional Integral Derivative (PID) control and Discrete Event Control. These typically are rudimentary designs focused on control of independent subsystems and provide only limited supervisory and coordination capability. However, today's systems are increasingly coupled and interdependent. The fundamentals of reliable infrastructure have not been adequately worked out for complex networks of highly-interacting subsystems, such as the power grid and the airspace-aircraft environment. These are complex, often dynamically reconfigured, networks. The primary challenge for future generations of these systems is to provide increasingly higher efficiency, while assuring joint physical and logical containment of adverse effects. This is the research agenda of secure network embedded systems.

This NSF/OSTP workshop on September 19th, 20th 2002 began with a number of plenary presentations and contextual discussions of issues in the area of information assurance and survivability, critical infrastructure protection and networking. Two infrastructures, power and air transportation, were highlighted as exemplars to focus on. Several break out sessions were organized to draw out a research agenda to support the most critical needs. An important backdrop to the workshop was the Draft National Strategy to Secure Cyberspace, which was released for comment on September 18th, 2000 - the day before the workshop by the Presidential Critical Infrastructure Protection Board.

The technology recommendations of our workshop call urgently for new research and development targeted in the following areas (details of the subtasks in the areas are in the report)

1. ***Information Assurance and Survivability***
2. ***Secure Network Embedded Systems***
3. ***Validated Modeling, Simulation and Visualization of Critical Infrastructure Systems and their Interdependencies***

This workshop report develops recommendations on the questions of how to speed up technology transitions of the research into the stakeholder critical infrastructures.

This report does not aim to develop specific program recommendations for the inter-agency funding of programs in the three areas listed above. However, the group felt that it was important that research programs be formulated urgently to begin in FY 2003 by both traditional funding agencies for research: the National Science Foundation, Defense Advance Research Projects Agency, Department of Defense, National Institute for Standards and Technology, and others along with stakeholder agencies like the Department of Energy, the FAA, the Transportation Safety Administration, Department of Commerce, Department of Treasury and other agencies in concert with the establishment of the Department of Homeland Security. The problems are urgent and large. The community is unusually strongly motivated and industry is present at the table to begin a series of very exciting public private partnerships.

Table of Contents

1. Introduction

Long-range research in information technology is crucial to Critical Infrastructure Protection. Today's weak infrastructure is due in large part to the fact that traditional approaches to Digital Control Systems (DCS) and SCADA have not been brought up to the standards of modern information technology. The techniques commonly employed are ad hoc combinations of PID and Discrete Event Control. These typically are rudimentary designs focused on control of independent subsystems and provide only limited supervisory and coordination capability. However, today's systems are increasingly coupled and interdependent. The fundamentals of reliable infrastructure have not been adequately worked out for complex networks of highly-interacting subsystems, such as the power grid and the airspace-aircraft environment. These are complex, often dynamically reconfigured, networks. The primary challenge for future generations of these systems is to provide increasingly higher efficiency, while assuring joint physical and logical containment of adverse effects.

Increasingly, autonomous but cooperative action is demanded of constituent elements. Examples include the technology needed to support aircraft in high-capacity airspace, enabling the execution of parallel landing patterns under terminal area control. A deregulated power grid draws new market participants. These new players may produce highly variable efficiency, potentially adverse environmental effects, and they may pose hazards to system-wide stability. This trend towards autonomous, cooperative action will continue, with the demands of current and next-generation systems for open, interoperating, and cooperating systems. The achievement of a satisfactory level of interoperable functionality is both enabled by, and dependent upon, advances in information and control infrastructure for coordinated operation. Furthermore, entirely new capabilities, such as networks of devices for pervasive sensing and actuation are becoming viable, and the control and communication technologies for their effective use must be fully developed and integrated into distributed infrastructure systems.

Although reference frequently is made to the next generation of technologies as "intelligent agent" systems or self-healing or self-reconfiguring or autonomic systems, this terminology conceals a complex of carefully integrated systems and software concerns. There is no panacea; services must be carefully engineered from the ground up in order to safely support a façade of highly autonomous action. Advances in software and information technology have improved the potential for a better substrate for future, more reliable infrastructures. For example, progress is seen in promising areas such as hybrid discrete and continuous control; open and object-oriented software technology for distributing cooperative interacting computations; optical networking that is more resilient against electronic intrusion; real-time and embedded systems research at various scales; Quality of Service ("QoS") management and accommodation of variability in networking protocols at all levels; fault isolation analysis and implementation; fault tolerance mechanisms; and improved technologies for intrusion detection, encryption, and key management. Model-based design technologies are improving our understanding of vulnerabilities and interdependencies among systems through simulation and other predictive methods. However, many gaps exist: security for wireless networks; combined networking modalities for dependable, real-time embedded systems; improved integration of authentication and management of authority in human and autonomous systems. Foremost is the need for ongoing assimilation of high-confidence information technologies into the Nation's infrastructure. These issues must be addressed intentionally and systematically. An integrated strategy is required to achieve all aspects of cyberinfrastructure to meet the Nation's needs for high-confidence information- and software-centric control of physical systems, and for secure information technology to support its networked and interdependent cyberinfrastructure.

The workshop began with a number of set point presentations and contextual discussions of issues in the area of information assurance and survivability, critical infrastructure protection and networking. Two infrastructures: power and air transportation were highlighted as exemplars to focus on. Several break out sessions were organized to draw out a research agenda to support the most critical needs. All the participants felt a need for urgency to engage in the problems and participate not only in the research but also in

the technology transition to products that find their way into the hands of the stakeholders.

An important backdrop to the workshop was the Draft National Strategy to Secure Cyberspace, which was released for comment on September 18th, 2002 - the day before the workshop - by Mr. Richard Clark and Dr. Howard Schmidt of the Presidential Critical Infrastructure Protection Board. A summary of that report (given in Appendix A) was made available to participants. While the primary thrust of the strategy is operational and focused on securing cyberspace quickly for CIP, this group felt it important to develop a research roadmap to support this strategy.

The organization of this report is as follows:
Section 2 has the proceedings of the workshop (supported by copies of the presentation briefs that were made available to the writers of the report). Section 3 has the technology recommendations. In summary, the technology recommendations call for research and development targeted in the following areas:

1. ***Information Assurance and Survivability***
2. ***Secure Network Embedded Systems***
3. ***Validated Modeling, Simulation and Visualization of Critical Infrastructure Systems and their Interdependencies***

Finally, Section 4 comments on the questions of technology transitions of the research into the stakeholder critical infrastructures. This section also provides some comments on the linkage to the national strategy to secure cyberspace.

This report does not have specific recommendations for the funding of interagency funding of programs in the three areas listed above. However, the group felt that it was important that research programs be urgently formulated to begin in FY 2003 by both traditional funding agencies for research: the National Science Foundation, Defense Advance Research Projects Agency, Department of Defense, National Institute for Standards and Technology along with stakeholder agencies like the Department of Energy, the FAA, the Transportation Safety Administration, Department of Commerce,

Department of Treasury in concert with the establishment of the Department of Homeland Security. The problems are urgent and large. The community is unusually strongly motivated and industry is present at the table to begin a series of very exciting public private partnerships.

2. Workshop Proceedings Summary

Opening Plenary Session

Peter Freeman, the CISE Assistant Director at the National Science Foundation kicked off the workshop and advocated for short term, intermediate term and long-term research. He emphasized the need for the scientific community to help mission agencies with specific needs in the short term as well as longer term. We are entering a period not unlike the cold war period with a set of hard problems to be worked on for the next 25 years.

Richard Russell, Deputy Director for Technology at the Office of Science Technology and Policy in the White House emphasized that the development of a science base for Critical Infrastructure Protection Research was a global issue. He pointed out that the role of OSTP, working with the Office of Homeland Security, is to articulate and coordinate R & D through an interagency working group in a manner analogous to the current Networking and Information Technology Research and Development (NITRD) group. He emphasized that OSTP would provide technical advice to homeland security agency. The OSTP role was to interact with all agencies on a continuing basis, set up working groups and to assure that good work and ideas in government gets translated to cohesive agenda and get funded.
Private industry holds 85% of the Critical Infrastructures in the nation. This highlights the need for public private partnerships to transition the developed technologies and research into the private sector. He exhorted the scientific community to enhance scientific research by understanding what others in community are doing and looking for opportunities to collaborate in CIP. Hard problems abound in the areas of intrinsic technology problems for IT systems, user expectations, interdependency analysis, intrusion detection, vulnerability assessment for providing framework, anticipate new attack scenarios and risk cost analysis.
Russell concluded by emphasizing the need for collaboration among many groups to make substantial progress, and pledging the support of OSTP in representing the concerns at the national S & T level.

Tom Cabe of the Critical Infrastructure Protection Board gave a brief introduction to the national strategy, which had been released for comment on September 18th, 2002. The details of this plan are reviewed in the Appendix to this report.

Peter Freeman emphasized that the focus of workshop was on secure network and embedded systems – particularly SCADA systems that are critical to all critical

infrastructure systems. He pointed out the need to explore interdependency of cyber structure (Internet) with IT centric physical systems (power, telecommunications, transportation, etc.). He viewed its role in considering the research gaps to identify gaps in software and IT that provide support for monitoring and protecting information systems. He directed the group to provide a roadmap for R&D activities in this specific area of embedded systems, with strategic goals being to empower digital systems and SCADA systems to protect their cyberspace. He emphasized that vital research and new technology for encryption and authentication capability for SCADA systems was needed. Mark Leblanc the Executive Officer for the PCCIP committee in OSTP and an officer of the subcommittee for CIP R&D re-emphasized the need for a strong road map effort and a need to identify steps for security emergency preparedness and areas of collaboration across the agencies.

Doug Maughan of the Advanced Technology Office at DARPA spoke about the Infosec Research Council and its efforts at developing a list of hard problems. These are posted at the web site www.pccip.ncr.gov: The Infosec Research Council is a government-sponsored group founded in 1996 to facilitate collaboration among government researchers and funding agencies. Its findings do not impact agency funding. Rather, they serve as coordination mechanisms between government program managers at DARPA, NSF, DoE, and other agencies to coordinate research. The IRC has bi-monthly meetings and presentations attended by DoD, DoE, NIST, NSF, FBI, NSA, OSD, AFRL, ARL, NRL, ONR, CIA, and NRO. The web site for the IRC is http://www.infosec-research.org. The IRC hard problems list was very instructive and we have tapped into it in our own assessment of what remains to be done. See the recommendations section of this report for this group's discussion of the CIP hard problems list.

Charge for the Workshop

Shankar Sastry gave the charge for the workshop. He began by chronicling the evolution of warfare from Pearl Harbor to terrorist attacks on 9/11 moving on to possible cyber threats to our most critical infrastructure. He made a case for the use of the term asymmetric threat to account for attacks on the soft commercial underbelly of the country. As outlined by President Bush the prime concern in protecting critical infrastructure and key assets are: Intelligence and warning, border and transportation security, defending against catastrophic threats and emergency preparedness. He emphasized that it was important to envision public private partnerships for critical infrastructure protection that harness national economic strength. In particular, the entrepreneurial energies of the venture community need to be leveraged in the solution. He emphasized that it was important to develop some distinctions between information assurance and survivability and critical infrastructure protection, since this had not been traditionally addressed.
The evolving infrastructure will include networked embedded systems (of which SCADA/DCS were the initial exemplars), which needed to be protected. In the state of the art in networked embedded systems, current methods have reached their limits in terms of scale and complexity. New work needs to be done to bring together the fields of security, networking and embedded systems and software. The development of sensor

networks with their ability to ubiquitously monitor the environment and infrastructures provides a promising set of new opportunities for protecting our infrastructure.

In addition to protecting today's infrastructure it was important to evolve the infrastructure in new directions to meet new technological challenges and provide new functionalities such as in chemical plants/nuclear power generation, automated highway systems, medical devices. A key characteristic of critical infrastructures is the need to be able to operate through attacks. Despite significant progress in intrusion detection and protection using firewalls, VPNs and other technology solutions, current systems continue to be unable to operate through attacks. New approaches to tolerating attacks include the use of diversity, redundancy, decentralization, detection and repair of damage. Further, long term research on biological models of tolerating attacks may shed light on the development of these schemes.

Sastry concluded with a list of policy concerns, which he encouraged the participants to keep in mind:

1. Ambiguity Between US National Interest and Multinational Corporations
2. Governance/Management of Critical Infrastructure Systems (Power Grid, Internet, Telecommunications)
3. Technical Standards including PKI
4. Security/Reliability Best Practices
5. Defining and Defending Privacy
6. Conflict Between National Regimes (Privacy, Cultural, Economic, Security) and Network Imperatives
7. Structures for Government – Industry Joint Operations
8. Information Sharing
9. Response/Restoration
10. Legal and Liability Systems for the Cyber Age
11. Human Interface Challenges
12. Managing a dynamic infrastructure with new generations
13. Intelligence about the future of conflict

Network Embedded Systems

Panel chaired by Helen Gill, National Science Foundation
Panelists:
- Janos Sztipanovits, Vanderbilt University
- Doug Schmidt, DARPA
- Mary Maeda, NSF

Helen Gill focused attention on a) identifying research gaps to secure Critical Infrastructures, b) thinking about a long term view of evolving architecture, and c) addressing the need for a community spanning areas of networking, embedded software and security. Although traditional research communities have grown in each one of these areas especially rapidly, there is a need for a unified framework with a clear short term, intermediate term and long-term agenda.

Janos Sztipanovits from Vanderbilt University gave a strong case for the use of networked embedded systems and software control: Embedded information processing is becoming the primary source of innovation in civilian and military systems. The new wave of inexpensive MEMS-based sensors and actuators and the continued progress in photonics and communication technology will further accelerate this trend. Systems will become increasingly "information rich," where embedded monitoring, control and diagnostic functions penetrate more deeply, with smaller granularity, in physical component structures. Given this trend, the separation of physical and information processing architectures is not sustainable. Strong mutual interdependence requires their fusion at fine levels of granularity, i.e. the distribution of information processing among physical components. The coordinated operation of distributed embedded systems makes *embedding, distribution, and coordination* the fundamental technical challenge for embedded software.

Research in networked embedded systems increasingly enables "fine-grain" fusion of physical and information processes. Soon, we will be able to build dependable, real-time, distributed, embedded applications comprising 10^2-10^6 computing nodes. The nodes will be networked; their operation is coordinated and dynamically reconfigured as a response to changing physical conditions and modes of operation. The nodes include physical and information system components coupled by sensors and actuators. Closed loop interaction between physical and information system components is an essential feature of networked embedded system applications. It differentiates the IT components (computing and networking) in this area from general, ubiquitous computing directions. Examples for emerging applications include MEMS based control and health management of complex mechanical systems, coordinated operation and control of large groups of physical objects (cars, weapons), and smart structures.

Major challenges in networked embedded systems technology include *coordination, system synthesis* and *security*. *Coordination services* include fault tolerant, self-stabilizing protocols for time, data exchange, synchronization, and replication in large, distributed, real-time systems. *Synthesis services* provide time-bounded solution for complex, distributed constraint satisfaction tasks required for dynamic reconfiguration of applications. *Security* services offer protection against denial of services attacks using physical and computational attack strategies, unauthorized access to the computation and communication fabric of applications, and violation of system integrity by merging unauthorized nodes and communication channels in networked embedded system applications. These services are crucial to making aggregate behavior of large networked embedded systems predictable and dependable despite local failures and upsets. The services need to be designed to be optimizable for specific applications and underlying distributed computing platforms and execution contexts. The application and computing platform specific optimization of service packages will require automated composition. Support of partitioning is essential despite critical and non-critical applications sharing the same fabric.

Doug Schmidt of DARPA spoke of the necessity for developing open computing platforms for distributed real time and embedded applications. One of his key concerns was the need to have reliable or available or high confidence platforms for the infrastructure of the future. He laid out a series of challenges in middleware and components for complex distributed embedded systems. The challenges are to develop standardization, open standards and open source methods for building embedded software (using as model the spectacular advances in computing and networks in the recent past) for improving software development quality, productivity and assurance. Schmidt recommends government and the research community take responsibility for meeting the long term challenges of developing new open standards and methods.

Mari Maeda of NSF gave an overview of strategic directions in networking research. Optical networking, the internet and IP technologies, and wireless networks have gone through a period for explosive growth with incredibly high data rates now possible for the backbone. The reconfigurable all-optical network is an exemplar of a survivable architecture with fast provisioning and service set up. Also emerging wireless technologies such as ultra wide band radio networks (with their attendant single chip radios), ad-hoc multi-hop low power low cost sensor networks are new enablers for critical infrastructure monitoring and protection solutions. Several key issues remain to be addressed including interconnectivity, diversity of interconnection, layer abstraction for logical networks and overlays and complexity of the core networking protocols, as are network control and management. Robust networking is the key ingredient for critical infrastructure protection.

Models and Analysis of Interdependencies

Panel chaired by Sam Varnado, Sandia National Laboratories
Panelists:
- Steven Wicker, Cornell University
- Miriam Heller, National Science Foundation (could not attend)
- Linda Nozick, Cornell University

Sam Varnado of Sandia Labs spoke about the difficulties in identifying physical threats, cyber threats and systems interdependencies in complex infrastructure systems. Interdependencies are frequently introduced for reasons of convenience, data sharing and synergies of operations during normal modes of operation. However, under attack or degraded modes of operation, these same interdependencies magnify the consequences of disruptions of one infrastructure on others. One example is the use of networking for real time management of contracts in deregulated energy market places. The increased reliance of the Internet for bidding and contracts on power makes the power systems operations vulnerable to cyber attack. Other such examples include the use of networked resources for reading, configuring and health management of SCADA/DCS from centralized offices rather than the field. Varnado discussed the role of the National Infrastructure Simulation Center (NISAC) in developing modeling and simulation capabilities for developing validated models of interdependencies of critical infrastructures. In this context, there was extensive discussion of the need to protect

sensitive vulnerability information while making available realistic interdependent networks for red-teaming and blue teaming by the research community. Varnado emphasized that NISAC was expected to move into the soon to be created Department of Homeland Security and was looking for research partnerships with the academic and research community in all areas.

Steve Wicker of Cornell University discussed the evolution of telecom infrastructure over 150 years as a context for lessons to be learned about robustness in design. The telecom sector is one of the most heavily interdependent sectors since severed telephone fiber can affect many sectors such as mercantile exchange, air traffic control, power, and all the other critical infrastructures. Modeling interdependencies has three components: static representation of network states, dynamic representations of network trajectory, and decision theory. Tools for analysis need foundations for developing analysis tools, design tools, and operational tools.

Linda Nozick of Cornell University described her research on modeling and analyzing infrastructure networks using graph models. The goal of the work is to understand network robustness and to optimize investment decisions to increase network reliability. The graph models capture probabilistic information: they characterize interconnected networks with probability distributions of link capacities and correlations between link capacities. Graph models enable analysis of a) the probability that demands are met, b) the probability distribution for demands which are met and c) the investments that might improve those probabilities. At the present phase of the research, gas and electric networks and their interdependences have been modeled and analyzed. The current conclusion is that uncertainties, correlations and interdependences create complex systems behaviors. There is a need for modeling environments, which allow exploration of the design and operating choices and drive analysis and simulation tools.
The follow-on discussion has brought about interesting issues on the effects of deregulations and obtaining data for interdependency modeling. While regulated monopolies could react and perform successfully under pressure, companies in open deregulated environment may not be interested on spending money on overall robustness. The highly competitive environment also makes it difficult to get real data on security and robustness revealed. Nozick concluded that new legislation would probably be required to resolve these problems.

Overview of Critical Infrastructure Systems: Power Grid and SCADA

Panel chaired by Massoud Amin, EPRI
Panelists:
- Jose R. Gracia, Tennessee Valley Authority
- Robert Hutchinson, Sandia National Laboratories
- Marija Ilic, Carnegie Mellon University
- Robert Thomas, Cornell University

Massoud Amin introduced the panel discussion by reviewing key numbers characterizing the problems of the power grid:

- Electric power ranked number one for impact on society by NAE. The national power grid includes 861,199 MW, 517,116 miles of transmission lines and carries $224.5B revenue.
- The dominant trend is increased stress on all aspects of the system. Power grid is being operated closer to the edge. Capacity margin has shrunk to 10-15% from 25% in 1980. Demand for electricity grew 32% while new capacity grew only 16-17%. Annual growth of consumption is 2.1% nationally while transmission capacity expands much slower.
- Under these circumstances the role of information technology and CIP are very important. Unfortunately, R&D expenditures in the power industry comprise less than 0.3%, placing this industry in the bottom 20 among all industries. Characteristically, the nation spends more on dog food research than on electricity.

Our key challenge is updating aging infrastructure to meet today's needs, including national security concerns.

Jose Garcia from Tennessee Valley Authority gave an overview of security challenges facing operators of large power grids. TVA provides power for 80,000 square miles. Responsibilities include power control and telecommunications. Operators don't know where problem is going to be, and can't protect against everything. Since control centers are in basements or bunkers of buildings, Supervisory Control and Data Acquisition Systems (SCADA) are very important. SCADA systems take information out of plants and bring it into control centers so operators can make their decisions. Human performance is very important in this environment where mistakes are unallowable. Deregulation and move to the free market has had tremendous consequences. The power control system currently operates in way it was not intended: Power system information is now available through the Internet; the number of transactions is increasing across TVA power grids; more complex, interconnected electric systems are required by the market; power infrastructure is owned by number of companies.
Within this context, critical inter-relationships need to be scrutinized. Among these challenges is the new imperative of protection against simultaneous cyber- and physical terrorism.

Robert Hutchinson form Sandia National Labs discussed cyber vulnerabilities and security for process control systems (PCS). The US industry increasingly relies on PCS-s for proper operation. This makes vulnerabilities of SCADA systems a very important security issue. Currently, there is no authentication of origination of commands. Frequently, software updates to PCS-s are made through SCADA network, which are connected to the Internet. Since the risk environment is changing, technology needs to respond.
There are several important research topics to make PCS-s more secure: ability to test and model effects of cybersystems on these networks; development of secure PCS architecture; development and adoption of security standards and conformance testing, in wired and wireless networks. Intrusion detection systems need to be extended to PCS-s. Further, we need to have a better understanding the implications of wireless network security.

Marija Ilic from CMU outlined a control engineering approach for complex infrastructures. The current trend is to move to highly decentralized systems. This trend is supported by the network-centric distributed IT infrastructure. There is a strong need to develop a control engineering perspective on the operation of these emerging complex dynamic systems. We need tools that are flexible and reliable. Metrics for measuring and evaluating reliability and flexibility are crucial. Researchers also need to understand how to model these systems for robust feedback control.

Bob Thomas of Cornell discussed vulnerabilities of the power grid. As a result of restructuring in the power industry, there are many more players and devices, significantly increasing heterogeneity and connectivity. Current systems were not designed for this environment. Consequently, they are becoming more exposed and vulnerable. Protecting the grid requires avoidance, assurance, detection, and recovery / restoration. Reliable and secure operation can be provided if we move from assessment and avoidance architectures to command and control architectures, which respond in real time to real-time collected data. The ongoing restructuring in the industry depends on movement of data (markets, metering, billing, etc.). It means that operational reliability becomes more dependent on better monitoring and control, which involves more mission critical communication and requires computer systems with all information security implications. The importance of this issue is clearly demonstrated by the fact that widespread failure of interconnected large-scale complex networks is almost always initiated by a failure in the electric power system.

Additional Plenary Talks

Protecting the Federal Aviation Administration from Cyber-Attack: Art Pyster
Although the air traffic control system is not challenged by cyber attacks, the risk is huge to secure operation. Every day about 2 million people fly, while 60,000 tons of cargo moved. The Nation's air space is controlled by 500 FAA-managed ATC towers and includes 10,000 airports in the country. There are 180 low altitude radar control systems and 20 enroute centers for controlling high altitude traffic.
The threat has emerged differently than expected. Previously, security concerns concentrated on hackers. Now, there is more awareness of the potential danger of cyber attack occurring simultaneously with physical assault.
The FAA follows a multi-prong strategy to secure safe airspace management.

- First, there is a continuous updating and revision of strategy, policy and guidance. This activity includes updating threats and vulnerabilities to determine focus for mitigation and updating security requirements according to the vulnerabilities of systems and attention to day-to-day administrative issues – passwords, patches, etc. The amount of data from monitoring compliance is enormous. The challenge is to analyze this data and make decisions based on data. The FAA needs research and tools to look at massive data and understand when systems are under attack – particularly a distributed attack.
- The second activity seeks to harden individual system and network elements. The requirements are the following: no one can take over system; elements can be

isolated to avoid viral spread; and key elements have backup to avoid service disruptions.

- The third activity focuses on coupling between different elements. An approach similar to the 5-layer enterprise security model is followed. A particular problem is that physical security and cyber security have different issues and cultures, which makes the communication between communities difficult.
- The fourth area of activities focus on awareness and execution of actions: training for awareness and training people to do the right thing. Every new system is certified to be secure. Internet access points are tightly managed so that the ATC system never touches the Internet (rather, it runs as private network). The ATC system is constantly monitored to ensure that new functionality does not add hooks to the Internet. Anti-viral software and firewalls throughout the network prevent spreading. Physical security is strictly maintained. Smart card technology replaces ID badges throughout agency and contractors, and control access to computer systems. The FAA, NSA and NIST are working on common criteria to define security requirements for ATC systems. This includes defining enterprise wide PKI infrastructure.

The ATC not used to the current, rapid pace of technology changes. There are cultural challenges for rapid updates of security technologies. We need more specific ways to monitor policy compliance, from automated compliance to security policies. We do not yet have the ideal architecture defined, and there is a need more insight for large, complex networks.

Introduction to Airspace Management Issues: Shankar Sastry

Shankar Sastry began by giving an overview of the current organization of the National Airspace System with its 22 regions and 257 sectors and its centralized dispatch from the Airline Operating Centers and more than fifty TRACONs. Technology trends in communication, navigation and surveillance (CNS) are predicted to lead to better capacity for airline traffic while simultaneously bringing new vulnerabilities. The economic impact of disruptions in the air traffic sector is enormous. Sastry presented data showing the wide spread macro economic impact caused by the drop in air traffic after September 11th attacks. Presenting the perspective of Professor John Hansman of MIT, Sastry emphasized the need to introduce more of a systems approach to assessment of vulnerability, redundancy and countermeasures with a cost benefit analysis. Jamming and spoofing are critical to flight management and navigation on board aircraft and inner loops on Air Traffic Control. Decentralization of Air Traffic Control, Traffic Collision Alert Systems (TCAS), and levels of redundancy for communication and radar interruptions were desirable to guard against attacks.

Vulnerabilities that might be introduced on board aircraft by the introduction of Internet connections on aircraft, as well as the networking of radar data, are cause for concern. While it was likely that SCADA systems still constitute a vulnerability to air traffic management and control, more work is needed to identify the specific vulnerabilities.

Aviation Safety and Efficiency: Marshall Potter

Dr. Marshall Potter, the Chief Scientist of the FAA spoke about safety, security and system efficiency on the National Airspace System (NAS). The NAS system, operated by the FAA, involves 67 million operations a year involving many stakeholders and partners with varying levels of trust. It has safety as its number one goal. In the next 10 years, a 52 % increase in long haul passengers and a 132 % increase in commuter traffic are expected within the US and a 100% increase is expected internationally. Given the 30% increase in expected workload, there is an opportunity to update the infrastructure and make it more secure in the process. Additionally, there will be a demand for new services, such as increased sharing of real-time data. Creating a strategic plan for the upgrade includes: cyber security, e-government, and business value. Cyber security takes the largest part of budget. E-government will make it easier to deal with community regarding data management and web-based systems. Availability and integrity are the highest priorities. Confidentiality is not a primary concern, unlike in the DoD. The key functions of the FAA in operating the NAS system are protection, detection, response, and recovery. The FAA has done work in protecting and detecting, including monitoring networks. Response and recovery mechanisms have not gone through same development and need to be addressed.

Potter differentiated the FAA focus in long-term research from DoD research priorities. The FAA addresses three new programs –

1. Real-time intrusion protection, detection, response and recovery.
2. Integrity and confidentiality in the mobile environment.
3. Trustworthy systems from untrustworthy system with untrustworthy actors.

Security With Privacy: Doug Tygar

Doug Tygar of UC Berkeley reported on the details of a recent ISAT study, "Security with Privacy," that he had conducted with Edward Felten for the Information Awareness Office at DARPA. Tygar talked about the challenges in developing strategies to integrate the maximum amount of security with the maximum amount of privacy - and the need to build a privacy simulator using synthetic data that is representative. The study came up with 11 recommendations as follows:

- *Policy recommendations*
  - Citizen advisory board to inform & shape policy
  - Support research on privacy laws & policy options
- *Technology challenges*
  - Accurate labels for derived data
  - Formal language for expressing privacy rules
  - Simulator for testing policy alternatives
  - Privacy toolbar
  - Tamper-evident distributed audit
- *Fundamental research topics*
  - Privacy & human factors
  - Distributed information flow security
  - Advanced crypto protocols
  - Adaptation

Report on the Digital Pearl Harbor Project: Richard Hunter and French Caldwell
Richard Hunter and French Caldwell of the Gartner group spoke about the Gartner group study called Digital Pearl Harbor commissioned by the Naval war college. The purpose of this study was to study disruption of a sample cyber attack on U.S. designed to cause significant shift in balance of power and assess damage to power grid, telecommunication system, networking systems, financial services systems. While several industry sectors have run their own attack scenarios, with the exception of the financial sector, there has not been a synergistic attack involving all sectors simultaneously.
The study identified key points in each area that would have leverage. Potential vulnerabilities were SCADA and operational control systems. Coordination of a cyber attack with a physical attack would cause more damage to the telecom switching systems and transatlantic fiber cables. In telecom and power grid, attacks were difficult to effect and would require insider cooperation. Information on how to pull attacks together was readily available (for example, where critical overseas cables were located). In the financial services system, the key vulnerability is exchange of transactions. The key conclusions of the Digital Pearl Harbor project are that a coordinated attack would cause significant disruptions and bring down the entire network with devastating impact.

North America Power Grid: Jose Garcia
Jose Garcia of the Tennessee Valley Authority spoke of the vulnerabilities of the electric power infrastructure system. In addition to protecting physical assets associated with the power network, the main vulnerability is in the area of SCADA systems and voice information exchange. The key technology challenges here are to provide real time authentication and to move away from QoS guarantees to real-time guarantees needed for real-time operation. Garcia also made a case for modeling, simulation and visualization tools to assess risks and then to develop an assessment process. The vast diversity of different power plant designs in the country made standardization difficult. While government and other programs have tried to assess the vulnerability of critical infrastructures, there is a sense that such assessments focus on piecemeal attacks (those without strategic objectives of the kind that a determined adversary is likely to mount). For assessing interdependencies' vulnerability for a strategic attack, an open / unclassified platform is needed.

Concerned Scientists for Cybersecurity: Sami Saydjari
Sami Saydjari from SRI International reported responses from a group of leading computer scientists, Concerned Scientists for Cybersecurity, to the recently released National Strategy to Secure Cyberspace. He advocated for defense against cyber attack to be a national priority with increased allocation of government resources. Counter strategies to national attack need to be developed by government in a top-down, technical-driven approach, rather than the current bottom-up and market driven approaches. He emphasized that current technologies are insufficient to defend against cyber war. More extensive R&D is needed and should be focused on engineering solutions. Concerned Scientists for Cybersecurity also advocated for increased government subsidization of the efforts addressing critical infrastructures and transition of technology, since private companies may have difficulty making business cases to defend against nation states.

Critical Infrastructure Systems: Air Traffic Management & National Airspace Systems

Panel moderated by Mr. Feisal Keblawi, FAA
Panelists:
Tim Wallace, FAA
David Sharp, Boeing Phantom Works
Edward Lee, Berkeley
Chip Meserole, Boeing Commercial

Feisal Keblawi reviewed the most challenging issues in R&D for air traffic management, highlighting interdependencies among the systems involved as a key problem. The ATC system is highly complex and highly redundant. This redundancy provides the basis for a system fail-safe capability. However, human interaction is a critical component and must be accounted for (in both positive and negative ways) in technological solutions.
The industry sees trends for higher level of connectivity between subsystems, the push for free flight, higher degrees of automation, the need for collaborative information among different users of the ATC system, and complex interdependencies emerging. New research is required to meet these challenges and address these issues. Another key R&D issue is to develop results for a balance between security and protection, detection capabilities, and response & recovery. Intrusion detection systems cause false alarms, requiring manual intervention, which impacts the operational budget. We need to minimize false alarms without missing genuine security breaches.
It is also difficult to assess the FAA's success in security and protection. Cost is a major issue, and the FAA needs solutions that reduce life cycle cost. The costs must also deal with operations & maintenance, which incurs costs higher than the capital cost. Hence, we need to develop new metrics, measurement techniques, and cost models.
We also must understand what the true threats are. Given those threats, how much insurance is enough? We need trust models that apply to systems of systems. The technological challenge includes the need for administrative and management tools that work in separate domains, but work together for assess different levels of risk for the integrated system.

Tim Wallace discussed The North America Surveillance Plan to distribute surveillance information and protect the integrity of that data. A goal is to create a common airspace picture. This sharing of data must be done in a manner where requests for information can be authenticated. We must also have solutions that preserve information so that it is not compromised. Studies are required that identify what data is appropriate for release and what vulnerabilities exist in the system. Another research question is how to manage security vulnerabilities given that the system is becoming more "open." Mr. Wallace suggested the need exploit the expertise of private industry in solving these problems.

David Sharp noted that operational complexity of large scale embedded systems is increasing at greater rates than in the past. To accommodate this increase in complexity the focus has been on product & process technologies, on safety and security, on using

diverse functionality, and developing new solutions for highly dynamic environments. Software is a key ingredient that contains multiple criticalities. Boeing has developed Bold Stroke to address many of the software issues and is using this technology as a transition conduit for research. Another technology challenge is certification. Certification is focused on the careful control of software – getting it right. For example, the 777 cost $1B to certify the software systems. As systems scale in size, there is a greater resource management challenge. It is necessary to share computational resources. At the same time it is necessary to extend the capabilities along multiple dimensions, e.g., greater timeliness, quality, security, power savings, and reliability. Solutions that can apply across multiple aircraft will make those products more affordable. However, it must be verified that transferring solutions this way will be safe and will work. It is also necessary to be able to do predict how it will work. A key research problem is dealing with systems of systems of systems. For example, for such large enterprises we require external dependability and internal dependability to build attack resistant systems. In the future, it may be necessary to move towards higher automation (for example, moving from trusting the pilot to trusting the software). New research to support an integrator of large-scale systems would be very valuable. If such a capability existed, safe, secure, heterogeneous systems that include trusted and untrusted parts could be deployed. The system would be comprised of different levels of safety criticality and involve new and legacy subsystems.

Edward Lee advocated the use of control algorithms for aircraft that have a property of localization of safety envelopes. To achieve such algorithms requires a research agenda that includes model-based design, on-line models, and mode changes to impose safety envelopes. Solutions will be based on both centralized and decentralized concepts, possibly semi-autonomous actors. The result could be fully automatic flight control. On the other hand reducing pilot authority could be dangerous, making it difficult to respond to emergencies. Consider an example: GPS can be jammed. It is necessary to have a backup to GPS in the aircraft. This can be done with safety envelopes. However, these localization-based safety envelopes must be bulletproof. Any software in flight control is subject to certification. This is one reason why model-based designs may prove beneficial. It is also a research problem on how to retrofit older aircraft with safety envelopes.

Chip Meserole identified that one key problem related to the system of systems architecture of ATM is that these systems have to be safe and secure and at the same time provide a great capacity. The purpose of aviation is to provide capacity. FAA must provide safety. One goal is that all players have access to information needed to make decisions. This makes you more vulnerable, but safer and more secure if information is used correctly. It is necessary to have information when the emergency happens, i.e., have the correct knowledge to react to it. This is referred to as getting the right data to the right place at the right time with security. The common information network is geographically distributed. We have to manage authentication and access control of the databases, so as to permit a vast variety of users access to data relevant to doing their job. We need to put a comprehensive system in place, but it is difficult to take one approach. How do we manage it? Where should research efforts be employed? Risk assessment,

vulnerability, budget analysis, security policy, auditing, and implementation are all key research problems.

Networked Embedded Systems

Panel participants were moderated by Steve Wicker of Cornell University and included
Anish Arora, Ohio State
Bhaskar Krishnamachari, University of Southern California
William Merrill, Sensoria Corporation
Dave Nicol, Dartmouth College
Jack Stankovic, Univeristy of Virginia

Steven Wicker of Cornell University asserted that while the direct benefits of power and communication to our nation's citizens are clear, the positive externalities resulting from the existence of these networks are not always appreciated. The public switched telephone network, for example, acts as a nearly free market in long-haul telecommunications capacity. This has in turn allowed for the development of subsidiary products and services, such as data communication for air traffic control, commodities exchanges, and the Internet, that would otherwise have been much more expensive, or even too expensive altogether. The same process is also occurring at a deeper level. Almost any modern enterprise depends upon computer networks, and it is common to insist that the components of those networks use commercially accepted standards. This COTS requirement is shared by new technology efforts in medical settings, in the military, in the nation's air traffic control systems, banking and financial systems, disaster response systems, etc. In effect, such systems must run "over the Internet", even if they may not use the public Internet per-se.

Interconnections of these various networks are not completely understood, nor are the resulting opportunities and vulnerabilities. Although these are incompletely understood, the latter are often painfully evident. To complete this picture of opportunity and vulnerability from a partially understood network of networks, we must factor in the user. Though often ignored, the role of the user is often dominant.

We call for a combined research, education, and outreach program dedicated to the problems of Complex Adaptive Networks for Critical Infrastructure. This program will work to shed light on the scientific issues underlying existing complex networks, to develop methodologies for building new and better layered networks in the future, and to educate the next generation of business people, operators and even consumers both about risks and technical options. We must put tools in the hands of regulators so that those charged with developing policies can also monitor and enforce them. Through such activities, we expect to make real progress on the engineering, management, and economic problems inherent in the nation's operation and dependence on complex networks, and also to educate the communities at which our solutions are aimed.

Our task will not be simple and in some situations, what we propose would not even be possible. Complex adaptive networks (CANs) often exhibit an extremely high degree of structural complexity. They show great diversity in nodal type and in interconnections between the nodes. Perhaps most importantly, such networks evolve over time, and it can be very difficult to predict their dynamic behavior. Although conceived in isolation from each other, critical infrastructure networks are disturbingly interdependent.
Thus, particularly if we focus on existing infrastructure, the problems that arise can be intractable.

Anish Arora from Ohio State University discussed protecting critical networks from faults and intruders using self-stabilization techniques. Good solutions are lacking so he called for a new research program in protecting critical networks from faults and intruders. The proposed research program would accomplish the following: develop methods for designing and composing systems in the presence of unanticipated faults. He stressed that no matter how well planned designs are, new and unanticipated attacks will be launched on them. He also presented information on how self-stabilization is used in achieving various kinds of security properties.

Bhaskar Krishnamachari of the University of Southern California discussed the criticality and robustness issues in wireless sensor networks. New challenges include the large scale of these systems, their extreme energy limitation, their high failure rates, and the fact that they must operate unattended. To complicate the issues is the lack of a design theory. Phase transitions provide a possible new approach and basis for a design theory. Phase transitions may identify emergent behavior at some abrupt change in a global system property. These phase transitions may help identify vulnerabilities or connectivity problems. They may also help identify robustness implications. Security and robustness must be built into the system from the beginning and not added as an afterthought.

William Merrill of Sensoria Corporation discussed the capabilities and limitations of wireless networked embedded systems for infrastructure protection. Unattended, simply deployed, low power embedded sensing and processing elements may have significant infrastructure security benefits, particularly as the intelligence of these systems is increased to provide the system flexibility and complexity to autonomously identify and respond to evolving threats. Lessons learned from two DARPA systems, the Self Healing Minefield and an automated perimeter security system demonstrated during the Steel Knight combined arms exercise, as well as the mGate commercial Telematics platform were provided to illustrate the current capabilities of embedded systems, and the need for similar open system, and power efficient requirements for embedded autonomous systems to be successful in infrastructure protection.

David Nicol of Dartmouth College. Critical infrastructure systems rely upon distributed systems for data collection, monitoring, maintenance, and control. Authentication of the devices and data comprising such a control system is a paramount problem, as is the survivability of the network. Marianas is a peer-to-peer network comprised of devices that have hardware support for authentication. A Marianas network provides a backbone of survivable trust, a distributed trusted third party. Application areas being explored

include authentication distributed computation, distributed remote maintenance, survivable self-organizing PKI, and privacy/security in delayed binding applications (e.g. role-based email).

Jack Stankovic of the University of Virginia discussed the need for real-time data services in next generation sensor networks. The theme of the presentation was getting the right data, to the right place, at the right time with security. These types of services are required for many types of applications including large-scale command and control, power grid and air traffic control applications. Each of these applications can make use of sensor networks. In sensor networks, research has been done looking at various wireless architectures that are most suitable for real-time data services including using a model that mimics external storage, using a local sensor net storage mode, and finally a data centric storage model. The latter seems the most promising. Real-time data services are also needed for large-scale distributed systems that include the Internet. Here, new transaction protocols that support deadlines and fresh (timely) data access are required. This has often been translated into developing new concurrency control, scheduling, and commit protocols

3. Technology Recommendations

The group agreed upon three important areas of research and development in the short term, intermediate term and long term. They are as follows

3.1 Information Assurance and Survivability

While there has been support for an effort in Information Assurance and Survivability primarily at DARPA and recently at the National Science Foundation, and with some support from NIST, DOE, OSD and the NSA, an examination of the hard problems list of the Info-Sec Research given below shows that a great deal more needs to be done both in research and technology transfer in each of the areas.

a. Intrusion and misuse detection. Here the research is geared at providing system and network security managers with tools that can detect attempts to defeat system security from both without and from within (insider attacks). The methods should be:
    i. Automatic
    ii. Predictive'
    iii. Have a low false positive rate
    iv. Identify the adversary

   The successes thus far have been primarily in the area of signature-based solutions but there is a great deal more that needs to be done to get these methods to have the attributes highlighted above.

b. Intrusion and misuse response the aim here is to provide system and network security managers with tools and techniques for responding to attack or misuse so as to identify, limit, and recover the damage done by

an attack and also investigate the origin and mechanisms of an attack. The attributes needed are:

i. Shared situational awareness
ii. Automated attack assessment and internal damage assessment
iii. Dynamic Reconfiguration
iv. Automated Counterattack

The bottom line here is that what is needed here is the method for correct attribution and retribution (with the appropriate policy to allow this).

c. Security of foreign and mobile code The aim is to provide users with the ability to execute software of unknown or hostile origin without putting sensitive information and resources at risk of disclosure, modification of destruction. The desired attributes are:
    i. Confinement of Access and Capability
    ii. Encapsulation of Code

   New methods are needed for protection against malicious mobile code and protection against malicious mobile code.
d. Controlled sharing of sensitive information. The aim here is to provide users with the ability to process extremely sensitive information including classified or compartmented information in open, networked environments while protecting the information from unauthorized disclosure. In the past this used to be viewed exclusively as a military issue, but it is increasingly important in civilian corporate and enterprise applications as well as in coalition and partnership scenarios. Attributes include:
    i. The ability to access and process information everywhere
    ii. Dynamic Authorization
    iii. Automated Data Tagging

   For Critical Infrastructure Protection to work this is needed both nationally and internationally. For areas such as civil aviation it was felt that this was critical even with nation states, which were not coalition partners.
e. Application security. The aim here is to provide tools and techniques to support the economical development of applications which enforce their own security policies with high assurance. Attributes include
    i. Security requirements beyond what the system provides
    ii. Assumptions and formal statements of trusted operating systems

   The feeling here is that in the short term this is usual and worthwhile, but in the intermediate term it is important to re-engage in a program of development of trusted composable high assurance trusted operating systems.
f. Denial of service. What is needed is to provide network and system components with the design and ability to help resist denial of service attacks. Desired attributes are
    i. Attribution and Retribution (identify, deter and eliminate) sources of attacks.
    ii. Modeling, measurement and analysis
    iii. DDOS Attack Detection
    iv. Infrastructure Attack Dissipation

The key feature here is the attribution and retribution across multiple legal and operational domains, and the ability to work through an attack (that is the ability to not crash the infrastructure while under attack).

g. Communications security. This is the ability to protect information in transit from unauthorized disclosure, and support for anonymity in networked environments. Issues here include improved cryptography, key distribution infrastructures and coalition issues of releasability and interoperability.
h. Security management infrastructure. This is the need to provide tools and techniques for managing security services in large networks subject to attack. Attributes include:
    i. Key Management Infrastructure Transparency and Interoperability
    ii. Secure Automated Network Configuration/Management
    iii. Better Authentication/Revocation
i. Secure Wireless Communications. This is the need to develop information security techniques and systems that are responsive to the special needs of mobile tactical hostile environments.
j. Secure Systems Composition. This is the need to develop techniques for building secure systems out of insecure components. This needs fundamental new research methods and techniques.
k. Metrics for Security. Measuring Levels of Assurance is the Achilles heel of acceptance and deployment of assurance technologies.
l. New and Emerging Challenges in Information Assurance and Survivability:
    i. Peer to Peer Collaboration
    ii. Security in Nomadic Computing Environment
    iii. Human Factors or Ergonomics in Security
    iv. Detecting and Limiting Data Infiltration
    v. Software Non-proliferation
    vi. Network Surveillance and Hygenie: Global Warning
    vii. Insider threat detection, monitoring, response

3.2 Secure Network Embedded Systems

Embedded Computing and Communication devices are becoming pervasive in our infrastructure. The so-called revolution in Ubiquitous computing and communications is happening quietly but is inexorably filling our surroundings with networked embedded devices. These systems bring a great deal of new functionality but also a number of vulnerabilities associated with them. There is a great need for *embedded software*, which is software operating with and controlling the physical world. The problem is hard because commercial B-to-B and enterprise software has only an idealized model of the real world. When one studies why there are cost overruns on every new procurement, it is largely due to the vast under-appreciation of the cost required to design, verify, validate, and certify the embedded software. While the DoD is a key stakeholder for embedded software and has been the lead agency in commissioning research in the area of

embedded systems, it is our thesis that the emergence of new vulnerabilities in SCADA, DCS and PCS are simply the bow wave of the realization of the tremendous vulnerabilities of embedded software in our critical infrastructures. Since for the most part critical infrastructures are privately owned, it is clear that the way to strengthen them is to have tremendous commercial innovations and technology transitions

1. Critical Infrastructures such as Power, Telecommunications, and Process Control which need a combination of distributed network embedded devices and ad-hoc wireless sensor networks.
2. Commercial avionics and automotive electronics (where it is predicted that the cost of the computers and embedded software will exceed the drive train, body, etc. by early 2003)
3. Consumer electronics such as PDAs, cell phones
4. Copier/printer and FAX machines,
5. Television and other media
6. Process control for chemical and industrial manufacturing processes.

As hardware gets commoditized and ubiquitously embedded in our emerging infrastructure, we must stay ahead by more rapidly introducing new embedded software functionality that exploits the hardware.

Four thrusts are critical to success in embedded software:

- *Automated design, verification, and validation*. Current embedded software design practices are stove-piped, with different engineers and software designers working sequentially in different domains. For instance, in the avionics domain a weapons software engineer works on networked fires, sensors, counter measures, etc., a guidance and navigation control engineer works on the flight dynamics, and a propulsion engineer works on the engine software. Each engineer has specific domain expertise, but seldom a clear understanding of hardware, operating systems, and networking issues traditionally implemented by the computer scientists and IT software development teams on the project. We need design practices that allow simultaneous design and propagation of constraints among these different domain specific design teams, which enable
    1. *Verified design*, in a mathematical or formal sense
    2. *Validated design*, in an engineering sense, and
    3. *Certifiable design*, to allow regulatory agencies to certify.
- *High confidence systems*. These are systems for human-centered automation, such as the monitoring and surveillance of critical infrastructures, civilian flight control systems, vehicle electronics, combat systems, and early warning networked defense systems. A key concern with these types of mission-critical systems is the fragility of their software and their ability to be compromised by security breaches and denial of service attacks. Important challenge areas in high confidence systems and software include:
    1. *Narrow-waisted middleware*. The tremendous success of the Internet was the standardization of IP protocols, which allowed for large variability in

the underlying transmission physical layers (optical, ATM, Ethernet, etc.) and diverse application layers, which used the same abstractions of the network because of the abstraction of the IP layer. We need to create narrow-waist middleware to allow for a diversity of lower level operating systems and networking protocols to present stable abstractions to higher-level application and service layers. The middleware should address multiple considerations in qualities of service, qualities of information assurance, etc.

2. *Security and composable operating systems*. We need operating systems of varied size footprints to support a wide spectrum of applications, ranging from PDAs to routers to servers, with modularity and assurance of multiple levels of security.
3. *Tamper-proof software*. One way of protecting hardware is to make it tamperproof. If superiority is encapsulated in embedded software, it must be made tamper proof as well.

- *Generative programming*. A fundamental difficulty in creating software for embedded systems is the large number of interdependent design concerns, constraints, and the massive amount of details that influence the structure and composition of the code. Even if the vast amount information used for the design, verification, and validation of embedded systems were captured in the form of models by design automation tools , the current relationship between the models and the code of the embedded software would be only loose and indirect. Generative programming is a new software paradigm that automatically manufactures highly optimized code from elementary, reusable implementation components using high-level design models by means of domain-specific configuration knowledge. A central issue in generative programming is the specification and synthesis of generators, i.e. programs that take high-level design models and produce efficient and correct implementations.
- *Intelligent Microsystems*. Intelligent Microsystems are a new class of highly adaptable, highly integrated components (micro-systems) with the ability to self-assess and adapt in real-time, optimizing their micro-level performance and providing new levels of macro-level functionalities to meet the needs of next generation of military sensor and weapon systems. Conceptually, intelligent microsystems can be thought of the inorganic equivalent of higher level living organism. These organisms have two levels of intelligence.

    1. At the microsystem level, the interaction with environmental factors is autonomous, i.e., decision-making is done with no interaction from higher level thinking (examples include pulling away from hot surfaces, adjusting pupil size to light conditions, or spiking hormone levels in response to fearful situation). In these cases the technical objective is to have the module or sub-system respond to and control its operation in the face of varying conditions (temperature, noise, available power, signal strengths etc).
    2. At the macrosystem level, how can modules and sub-systems be integrated together to create distributed functions (computing, sensors, actuators etc).

> Such a capability involves understanding how to reliably integrate and control multiple units of multiple types under a wide variety of operational conditions. This would not only allow complex, distributed systems, but would provide adaptability and redundancy (in event of non-functional units). The "intelligence" under such situations is analogous to the higher order thinking that is done to not only collect data and respond locally when necessary, but to analyze the data and make changes based not only on what has happened, but what may happen and to adjust accordingly.

In the post-PC era we are evolving into a world of ubiquitous or pervasive computation and sensing. In this era, we are surrounded by computational elements and sensors embedded in the environment around us. We are in the computational medium and surrounded by it. When this trend is fully implemented and deployed, it will have a dramatic impact on our emerging infrastructure, which is currently based on decision making in the face of poor or incomplete information. Key areas to be addressed in being able to harness this computational power are:

- *Oceanic databases*. In a world of distributed sensing and computation and data storage, it is critical to provide users with consistent and current views of the world. It is important that queries not be directed to specific locations but to ask for information and the "oceanic data base" flow appropriately to the query.
- *Secure collaborations*. Different levels of trust need to be dynamically determined and the ability to seamlessly collaborate across these coalitions is critical.
- *Natural user interfaces* including speech, gesture, vision and other modalities to allow for natural interaction with a pervasive computational environment.

The emerging integration role of software which was driven by the tremendous success of information technology, has resulted in the emergence of fundamentally new challenges specific to embedded software technology, namely:

(1) physicality,
(2) change, and
(3) variable structures.

"Physicality" means that embedded software must be composed to satisfy conflicting physical requirements, such as dynamics, reliability, noise, power constraints, etc. "Change" refers to the expectation that embedded software – being the glue that keep the platforms together – shall provide controlled flexibility to tolerate and manage changes in physical platforms. "Variable structure" is an emerging new requirement for networked embedded systems, whose structure dynamically changes during operation. These challenges are completely ignored by commercial software development due to the different roles that the dominant commercial information systems play. Consequently, commercial industry that will depend on embedded software (e.g. automotive industry) faces similar, potentially devastating problems in the future if the challenges left unanswered.

Major challenges in networked embedded systems technology include *coordination, system synthesis* and *security*. *Coordination services* include fault tolerant, self-stabilizing protocols for time, data exchange, synchronization, and replication in large, distributed, real-time systems. *Synthesis services* provide time-bounded solution for complex, distributed constraint satisfaction tasks required for dynamic reconfiguration of applications. *Security* services offer protection against denial of services attacks using

physical and computational attack strategies, unauthorized access to the computation and communication fabric of applications, and violation of system integrity by merging unauthorized nodes and communication channels in networked embedded system applications. These services are crucial to making aggregate behavior of large networked embedded systems predictable and dependable despite local failures and upsets. The services need to be designed to be optimizable for specific applications and underlying distributed computing platforms and execution contexts. The application and computing platform specific optimization of service packages will require automated composition. Support of partitioning is essential despite critical and non-critical applications sharing the same fabric.

3.2.1. Application Independent Coordination Services

Real-time coordination is a crucial problem in distributed control applications, which are the main drivers for using networked embedded systems. Distributed control requires complex, dynamic interactions among a large number of dynamically changing computational components. The foundations for the dependable implementation of these interactions are composable protocols for coordination services such as global time, general information exchange (consensus, agreement, membership, etc.) distributed synchronization, replication and replica determinism. The services must be application independent but customizable and must be developed for a wide range of distributed computation platforms. Selection of coordination services and computing platforms are based on the requirements of distributed control applications comprising tightly coupled physical and information system components. Ongoing research at DARPA and NSF will result in formally verified algorithms and code bases, composable micro-protocols, test results and application examples.

Examples for research in this area are:

1. Self-stabilizing solutions that guarantee eventual consistency and recovery in dynamic environment from arbitrary initial states in spite of rich class of faults. Extension of discrete self-stabilization approach to hybrid systems.
2. Parametric design of coordination services that allow optimization of generic solutions to application characteristics.
3. Solutions for achieving approximate consensus, approximate synchrony, non-uniform time bounds, hierarchical coordination. These solutions will help to limit the requirements for coordination based on locality and heterogeneity in physical interactions.
4. Probabilistic approaches for coordination services that can be adapted easily at run-time.

3.2.2. Time-bounded Synthesis

The size of networked embedded system configurations, their tight integration with dynamic, non-stationary physical processes and limitations in component reliability make the use of self-assembly, self-configuration, self-repair and other forms of adaptation mandatory. These capabilities mean that synthesis of control sequences, schedules, processing configurations, resource maps, etc. – usually performed at design time – are becoming part of real-time (i.e. time-bounded) networked embedded system operations. Independently from the technical details and peculiarities of different applications, the fundamental challenge in synthesis problems is that search-intensive algorithms (constraint processing, scheduling/planning, combinatorial optimization) can easily lead to computationally intractable problem instances. Recent breakthroughs in mathematics and computer science identified phase transitions in computationally hard problems, which separate the computationally hard and easy problem instances using simple order variable(s). New research explores the phase transition phenomenon to develop a new generation of transition-aware solvers for time-bounded synthesis. These solvers will use statistical methods to assess the hardness of problem instances and use this information to modify the problem if an intractable instance is found.

The following topics are examples for ongoing research efforts:

1. Extension of theoretical and experimental findings on phase transitions.
2. Statistical analysis methods for exploring problem spaces, and use the collected statistical data for assessing the criticality of actual problem instances.
3. Distributed anytime solvers and "transition-aware solvers" that generate solutions incrementally by working from simplified problem instances toward full problem, and use indicators to avoid hard problem instances.

3.2.3. Service Composition and Adaptation

Coordination services include distributed algorithms that are highly dependent on the underlying distributed computing model determined by the network topology, synchrony, failure model and characteristics of message services. Higher level services, such as distributed reset or consensus, include several interdependent layers, and need to satisfy highly application specific requirements that significantly influence the complexity of the algorithms. Scalability and dynamic properties of networked embedded systems make the development and use of a single, monolithic coordination service package unfeasible. Ongoing research addresses the development of fully automated design-time and run-time composition and adaptation of service packages, which are optimized to the actual requirements and computation, communication platforms. Automation of the composition and customization is crucial, because dependability and the need for behavioral assurances make verification of the composed services mandatory.

The following topics are examples for ongoing research efforts:

1. Rigorous modeling techniques for representing distributed coordination protocols, platform and application models and requirement models.
2. Use of deductive synthesis for deriving service package models.
3. Model-based generators for the automated generation and optimization of coordination service packages from the synthesized models.
4. Coordination service components with run-time adaptable parameters and methods for the coordinated, run-time adaptation of the distributed services.

3.3 Secure Embedded Sensor Networks

Embedded sensor networks hold the promise of facilitating large-scale, real-time processing in complex environments. Their application can help protect and monitor military, environmental, safety-critical, or domestic infrastructures and resources. For example, sensor networks can be deployed around remote infrastructure devices or areas (electric towers, for example) to detect intruders or damage and act to limit the damage. As another example, they can be used to monitor the nation's water infrastructure and detect biological or chemical attacks. Sensor networks can also be used as an emergency response system. For example, if an earthquake hits a city an immediate deployment of a city-wide sensor network can help locate survivors, injured people, gas leaks, fires, etc. and direct rescue crews and other aid to the right place at the right time. Emergency response also applies to the nation's critical infrastructure protection. If an attack is begun on an infrastructure, a sensor network can detect and act in real-time to assess and limit the damage of the attack. Such quick action may prevent a cascading set of failures.

In these and other vital areas, keeping the sensor network available for its intended use is essential. Security attacks such as denial of service can result in the damage to health and safety of people. Without adequate security mechanisms in place for sensor networks, these solutions will be limited to well controlled and confined applications and environments. This would negate much of the promise such systems hold. New research is required to develop security mechanisms for sensor networks. These mechanisms need to be inherently different from today's heavyweight solution because of the limited power and capabilities of individual sensor nodes.

For many sensor network applications, security is critical. Some face not only a harsh environment, but also active and intelligent opposition

- *Disasters.* It may be necessary to protect the location and status of casualties or infrastructure loses from unauthorized disclosure—particularly if the disaster relates to ongoing terrorist activities instead of natural causes.
- *Public safety.* False alarms about chemical, biological, or environmental threats could cause panic or disregard for warning systems. An attack on the system’s availability could precede a real attack on the protected resource.

Strictly speaking, although we usually use the term denial of service (DoS) to refer to an adversary’s attempt to disrupt, subvert, or destroy a network, a denial of service attack is any event that diminishes or eliminates a network’s capacity to perform its expected function. Hardware failures, software bugs, resource exhaustion, environmental conditions, or any complicated interaction between these factors can cause a DoS. Although attackers commonly use the Internet to exploit software bugs when making DoS attacks, here we consider primarily protocol- or design-level vulnerabilities.

An intrusion-detection system monitors a host or network for suspicious activity patterns such as those that match some preprogrammed or possibly learned rules about what constitutes normal or abnormal behavior. Sensor networks destined for harsh environments should already be designed to continue functioning in the presence of faults. This robustness against physical challenges may prevent some classes of DoS attacks. Fault-tolerance may mitigate even node subversion, and efficient protocols will limit opportunities for malicious waste of resources. Developers must, however, factor the complication of an intelligent, determined adversary into the design separately. For example, they can design sensors to withstand the effects of normal thermal cycles in a desert environment or to cope with transient irregularities in radio propagation. However, this will not be sufficient to thwart an attacker with physical access to the node, which can move or heat and cool the device at will. An adversary may possess a broad range of attack capabilities. A physically damaged or manipulated node used for attack may be less powerful than a normally functioning node. Subverted nodes that interact with the network only through software are as powerful as other nodes.

Layered network architecture can improve robustness by circumscribing layer interactions and interfaces. A clean division of layers may be sacrificed for performance in sensor networks, however, reducing robustness. Each layer is vulnerable to different DoS attacks, and has different options available for its defense. Some attacks crosscut multiple layers or exploit interactions between them. Table 1 lists the layers of a typical sensor network and describes each layer's vulnerabilities and defenses.

Table 1. Network Embedded System Layers and DoS defenses.

Network layer	Attacks	Defenses
Physical	Jamming	Spread-spectrum, priority messages, lower duty cycle, region mapping, mode change
	Tampering	Tamper-proofing, hiding
Link	Collision	Error-correcting code
	Exhaustion	Rate limitation
	Unfairness	Small frames
Network and routing	Neglect and greed	Redundancy, probing
	Homing	Encryption
	Misdirection	Egress filtering, authorization, monitoring
	Black holes	Authorization, monitoring, redundancy
Transport	Flooding	Client puzzles
	Desynchronization	Authentication

3.3.1. Physical Layer

Nodes in a sensor network use wireless communication because the network's ad hoc, large-scale deployment makes anything else impractical. Base stations or uplink nodes can use wired or satellite communication, but limitations on their mobility and energy make them scarcer.

Jamming

A well-known attack on wireless communication, jamming interferes with the radio frequencies a network's nodes are using. An adversary can disrupt the entire network with *k* randomly distributed jamming nodes, putting *N* nodes out of service, where *k* is much less than *N*. For single-frequency networks, this attack is simple and effective.

A node can easily distinguish jamming from the failure of its neighbors by determining that constant energy, not lack of response, impedes communication. Both effects have similar results, however, since constant jamming prevents nodes from exchanging data or even reporting the attack to remote monitoring stations. Even sporadic jamming can be enough to cause disruption because the data the network is communicating may be valid for only a short time.

The standard defense against jamming involves various forms of *spread-spectrum* communication. To attack frequency hoppers, jammers must be able either to follow the precise hopping sequence or to jam a wide section of the band. Code spreading is a jamming method that mobile-phone networks commonly use. Given that these abilities require greater design complexity and more power, low-cost, low-power sensor devices will likely be limited to single-frequency use.

In a large-scale deployment, an adversary is less likely to succeed at jamming the entire network, especially if only subverted sensors perform the jamming. In this scenario, a more appropriate response would be to call on the nodes surrounding the affected region to cooperatively map and report the DoS attack boundary to a base station.

Tampering

An attacker can also tamper with nodes physically, and interrogate and compromise them—threats that the large-scale, ad hoc, ubiquitous nature of sensor networks exacerbates. Realistically, we cannot expect to control access to hundreds of nodes spread over several kilometers. Such networks can fall prey to true brute-force destruction, but also to more sophisticated analysis. An attacker can damage or replace sensor and computation hardware or extract sensitive material such as cryptographic keys to gain unrestricted access to higher levels of communication. Node destruction may be indistinguishable from fail-silent behavior.

One defense involves tamper-proofing the node's physical package. Its success depends on

- how accurately and completely designers considered potential threats at design time;
- the resources available for design, construction, and test; and
- the attacker's cleverness and determination.

3.3.2. Link Layer

The link or media access control (MAC) layer provides channel arbitration for neighbor-to-neighbor communication. Cooperative schemes that rely on carrier sense, which let nodes detect if other nodes are transmitting, are particularly vulnerable to DoS.

Collision

Adversaries may only need to induce a collision in one octet of a transmission to disrupt an entire packet. A change in the data portion would cause a checksum mismatch at some other receiver. A corrupted ACK control message could induce costly exponential back–off in some MAC protocols. The amount of energy the attacker needs, beyond that required to listen for transmissions, is minute.

Error-correcting codes provide a flexible mechanism for tolerating variable levels of corruption in messages at any layer. However, these codes work best as counters to environmental or probabilistic errors. For a given encoding, malicious nodes can still corrupt more data than the network can correct, although at greater cost. The error-correcting codes themselves also incur additional processing and communication overhead.

The network can use collision-detection to identify these malicious collisions, which create a kind of link-layer jamming, but no completely effective defense is known. Proper transmission still requires cooperation among nodes, which are expected to avoid corruption of others' packets. A subverted node could intentionally and repeatedly deny access to the channel, expending much less energy than in full-time jamming.

Exhaustion

A naive link-layer implementation may attempt retransmission repeatedly, even when triggered by an unusually late collision, such as a collision induced near the end of the frame. This active DoS attack could culminate in the exhaustion of battery resources in nearby nodes. This attack would compromise availability even if the adversary expended no further effort. Random back–offs only decrease the probability of inadvertent collision, thus they would be ineffective at preventing this attack.

Time-division multiplexing gives each node a slot for transmission without requiring arbitration for each frame. This approach could solve the indefinite postponement problem in a back–off algorithm, but it is still susceptible to collisions.

A self-sacrificing node could exploit the interactive nature of most MAC-layer protocols in an *interrogation* attack. For example, IEEE 802.11-based MAC protocols use Request To Send, Clear To Send, and Data/Ack messages to reserve channel access and transmit data. The node could repeatedly request channel access with RTS, eliciting a CTS response from the targeted neighbor. Constant transmission would eventually exhaust the energy resources of both nodes.

One solution makes the MAC admission control *rate limiting*, so that the network can ignore excessive requests without sending expensive radio transmissions. This limit cannot drop below the expected maximum data rate the network supports, though.

Unfairness

Intermittent application of these attacks or abusing a cooperative MAC-layer priority scheme can cause unfairness, a weaker form of DoS. This threat may not entirely prevent legitimate access to the channel, but it could degrade service by, for example, causing users of a real-time MAC protocol to miss their deadlines.

One defense against this threat uses *small frames* so that an individual node can capture the channel only for a short time. If the network typically transmits long messages, however, this approach increases framing overhead. Further, an adversary can defeat this

defense by cheating when vying for access, such as by responding quickly while others delay randomly.

3.3.3. Network and Routing Layer

Higher layers may not require fully reliable transmission streams, but the network layer provides a critical service nonetheless. In a large-scale deployment, messages may traverse many hops before reaching their destination. Unfortunately, as the aggregate network cost of relaying a packet increases, so does the probability that the network will drop or misdirect the packet along the way.

Neglect and greed

One simple form of DoS attacks the node-as-router vulnerability by arbitrarily neglecting to route some messages. The subverted or malicious node can still participate in lower-level protocols, and may even acknowledge reception of data to the sender, but it drops messages on a random or arbitrary basis. Such a node is *neglectful*. If it also gives undue priority to its own messages, it is also *greedy*.

The dynamic source routing (DSR) protocol is susceptible to this attack. Because the network caches routes, communications from a region may all use the same route to a destination. If a node along that route is greedy, it may consistently degrade or block traffic from the region to, for example, a base station.

Using multiple routing paths or sending redundant messages can reduce the effect of this attack by making it necessary for an adversary to subvert more sensor nodes. Differentiating a greedy node from a failed node can be difficult, however, so prevention is safer than relying on detection.

Homing

In most sensor networks, some nodes will have special responsibilities, such as being elected the leader of a local group for coordination. More powerful nodes might serve as cryptographic key managers, query or monitoring access points, or network uplinks. These nodes attract an adversary's interest because they provide critical services to the network.

Location-based network protocols that rely on geographic forwarding expose the network to *homing* attacks. Here, a passive adversary observes traffic, learning the presence and location of critical resources. Once found, these nodes can be attacked by collaborators or mobile adversaries using other active means.

One approach to hiding important nodes provides confidentiality for both message headers and their content. If all neighbors share cryptographic keys, the network can encrypt the headers at each hop. This would prevent a passive adversary from easily learning about the source or destination of overheard messages, assuming a node has not been subverted and remains in possession of valid decryption keys.

Misdirection

A more active attack, *misdirection*, forwards messages along wrong paths, perhaps by fabricating malicious route advertisements. As a mechanism for diverting traffic away

from its intended destination, this DoS attack targets the sender. By misdirecting many traffic flows in one direction, the DoS attack can target an arbitrary victim.

In one variant of misdirection, Internet *smurf* attacks, the attacker forges the victim's address as the source of many broadcast Internet control-message-protocol echoes. The attacker directs all the echo replies back to the victim, flooding its network link. Among sensor network routing protocols, DSR is also vulnerable to this attack. An adversary can simply forge replies to route-discovery requests, including victims in the spoofed route.

A sensor network that relies on a hierarchical routing mechanism can use an approach similar to the *egress filtering* in Internet gateways, which can help prevent smurf attacks. By verifying the source addresses, parent routers can verify that all routed packets from below could have been originated legitimately by their children.

Black holes

Distance-vector-based protocols provide another easy avenue for an even more effective DoS attack. Nodes advertise zero-cost routes to every other node, forming *routing black holes* within the network. As their advertisement propagates, the network routes more traffic in their direction. In addition to disrupting message delivery, this causes intense resource contention around the malicious node as neighbors compete for limited bandwidth. These neighbors may themselves be exhausted prematurely, causing a hole or partition in the network.

Although nodes can detect a black-hole attack more easily than they can detect greed, neglect, or misdirection attacks, a black-hole attack is more disruptive. Other nodes with untainted knowledge of the network topology may suspect inconsistent advertisements.

Some solutions for these attacks include authorization, monitoring, probing and redundancy.

Authorization

One defense against misdirection and black-hole attacks lets only *authorized* nodes exchange routing information. Traditional wired networks with comparatively few routers often take this approach. Routers may use a public-key encryption infrastructure to sign and verify routing updates. Sensor networks place higher demands on scalability because every node is a potential router by design.

In addition to the computational and communication overhead, designers find that key management is difficult when using public-key cryptography in sensor networks. Nodes form ad hoc relationships upon deployment, they may be mobile, and additional nodes may replenish them during their lifetime. A centralized certification authority would create a single point of failure, greatly hampering the network's scalability.

Nodes can still be subverted with their key material intact. This vulnerability could give an adversary the unrestricted ability to construct valid routing messages, although threshold cryptography with share updating can protect against this possibility.

Monitoring

Nodes can also monitor their neighbors to ensure that they observe proper routing behavior. In one approach, the node relays a message to the next hop and then acts as a watchdog that verifies the next-hop transmission of the same packet. The watchdog can

detect misbehavior, subject to limitations caused by collisions, asymmetric physical connectivity, collusion, and so on. Watchdogs inform a quality-rating mechanism, also running at each node, which chooses the most reliable routes for message transmission in much the same way that certain flow-analysis procedures work.

Probing

A more active approach that does not require every node to participate tests network connectivity by *probing*. Networks using geography-based routing, such as Greedy Perimeter Stateless Routing, can use knowledge of the physical topology to detect black holes by periodically sending probes that cross the network's diameter. Subject to transient routing errors and overload, a probing node can identify blackout regions.

A distributed probing scheme can also work. To detect malicious nodes, probes must be indistinguishable from normal traffic. Otherwise, neglectful or greedy nodes could always choose to route probes correctly, escaping detection.

Redundancy

Redundancy can lessen the probability of encountering a malicious node. The network can send duplicate messages along the same path to protect against intermittent routing failure or random malice. If each message uses a different path, one of them might bypass consistently neglectful adversaries or even black holes. A more clever approach uses diversity coding to send encoded messages along different paths, but with lower cost than full duplication.

3.3.4. Transport Layer

This layer manages end-to-end connections. The service the layer provides can be as simple as an unreliable area-to-area anycast, or as complex and costly as a reliable sequenced-multicast bytestream. Sensor networks tend to use simple protocols to minimize the communication overhead of acknowledgements and retransmissions. Protocols that provide sequencing share many DoS vulnerabilities with the Internet transmission control protocol.

Flooding

Protocols that must maintain state at either end are vulnerable to memory exhaustion through *flooding*. As in the classic TCP SYN flood, an adversary sends many connection-establishment requests to the victim. Each request causes the victim to allocate resources that maintain state for that connection.

Limiting the number of connections prevents complete resource exhaustion, which would interfere with all other processes at the victim. However, this solution also prevents legitimate clients from connecting to the victim, as queues and tables fill with abandoned connections. Protocols that are connectionless, and therefore stateless, can naturally resist this type of attack somewhat, but they may not provide adequate transport-level services for the network.

One defense requires clients to demonstrate the commitment of their own resources to each connection by solving *client puzzles*. The server can create and verify the puzzles easily, and storage of client-specific information is not required while clients are solving

the puzzles. Servers distribute the puzzle, and clients wishing to connect must solve and present the puzzle to the server before receiving a connection. An adversary must therefore be able to commit far more computational resources per unit time to flood the server with valid connections. Under heavy load, the server could scale the puzzles to require even more work by potential clients.

This solution is most appropriate for combating adversaries that possess the same limitations as sensor nodes. It has the disadvantage of requiring more computational energy for legitimate sensor nodes, but it is less costly than wasting radio transmissions by flooding.

Desynchronization

An existing connection between two endpoints can be disrupted by *desynchronization*. In this attack, the adversary repeatedly forges messages to one or both endpoints. These messages carry sequence numbers or control flags that cause the endpoints to request retransmission of missed frames. If the adversary can maintain proper timing, it can prevent the endpoints from exchanging any useful information, causing them to waste energy in an endless sychronization-recovery protocol.

One counter to this attack *authenticates* all packets exchanged, including all control fields in the transport protocol header. Assuming that the adversary also cannot forge the authentication mechanism, the endpoints could then detect and ignore the malicious packets.

3.4 Validated Modeling, Simulation and Visualization of Critical Infrastructures and their Interdependencies

The case of electric power utilities is an important exemplar of the kind of research that is needed in this area. After several serious blackouts in the Northeastern United States in the sixties, major industry efforts were undertaken to develop efforts for preventing future occurrences. Some of these efforts have led to methods for differentiating the degree of system stress and to the development of various methods to be used under various degrees of system stress (so-called normal, alert, emergency and restorative system conditions). Currently system operators routinely rely on such tools, which differentiate among various operating modes. In some ways the work done by the power industry is a good exemplar of the kind of the modeling of infrastructures. We first provide a summary of current operating practices and describe challenges to the commonly made assumptions related to technology changes. Possible problematic areas and critical missing tools are described in context of the assumptions and the need for relaxing these in order to move the electric power industry into a highly efficient and reliable system architecture.

Current operating practices for managing electric power systems

To start with, system operators rely on qualitatively different tools under normal operating conditions than when the system is under stress. These practices greatly reflect the way in which the electric power system interconnections have evolved over time. Up until very recently, each electric power company planned to have adequate generation and transmission to provide its own customers under normal operation, while little exchange with neighboring companies took place for economic exchange. Regional studies were routinely performed to identify potential problems under serious equipment outages. The interconnections among subsystems (utilities within a power pool, power pools within a region, and alike) were built to ensure enough transfer capability from the subsystems to a subsystem in which an equipment outage took place. This practice has enabled savings through cooperation at a regional level because necessary generation reserve to supply power under emergencies was shared among several subsystems. This general approach has led to a well-established operating mode in both normal and abnormal conditions, which, in turn, could lead us to a premature conclusion that there is not much to worry about when it comes to operating power systems. One could further conjecture that if it hadn't been for recent complications under the industry restructuring, there would not be major questions on SCADA supported Energy Management Systems.

In normal operation, power systems are managed under various assumptions, which provide a framework for strong temporal and spatial hierarchies. At the highest level of each subsystem equipped with its own SCADA, generation is scheduled so that supply meets anticipated load demand. Both demand and transmission grid is assumed given. This is done in a feed-forward way with an objective of minimizing the total production cost. This is done at several rates, week ahead, day ahead, and 5-15 minutes ahead. This scheduling is done assuming power flow exchanges with the neighboring subsystems as agreed upon (despite that most lines have no direct flow control) and so that the worst-case outage in the entire region (comprising several subsystems) is insured against. In other words, the approach is preventive so that the worst-case scenario does not require any on-line corrective actions. This so-called (N-1) security criterion employed under normal conditions generally results in significant inefficiencies.

Deviations from anticipated system conditions are compensated in an automated way by several power plants participating in so-called Automatic Generation Control (AGC) and/or Automatic Voltage Control (AVC) schemes, the latter only implemented in some European systems. These schemes are generally recognized as secondary control level since their objective is to change the set points at primary (equipment) level controllers, in order to cancel power imbalances at each subsystem level caused by unexpected deviations in load demand. At a primary (equipment) level, watt regulators, automatic voltage controllers and power system stabilizers, respond to fast random deviations of local frequency and voltage from the values set by the secondary level control. These controllers are basic Proportional Integral Derivative (PID) controllers, whose tuning is often based on equivalencing the entire system as seen by the controller with a very simple network, whose parameters reflect typical conditions. Procedures for operating power systems under stress, on the other hand, are system specific. They are result of

combined off-line studies of what is perceived to be critical scenarios and the human operator's knowledge about the system response under major equipment failures. The fully automated protective relaying facilitates the response under stress, which is put in place to disconnect pieces of equipment as frequency and voltage vary in response to the triggering equipment failures. Several major system blackouts could be traced to the malfunctioning of protective relays; the malfunctioning is generally related to their logic, which is highly localized and non-adaptive to the changing system conditions. Methods for representing power systems and decision making under stress are practically nonexistent. The reasons for this are many, ranging from highly nonlinear hard to model dynamics in response to major equipment failures, through historic reliance on human knowledge about the specifics of particular subsystems.

These operating practices have recently been challenged by major technological changes and by the industry restructuring process. Both of these require major rethinking of the assumptions underlying temporal and spatial hierarchies summarized above. Generally, technological changes have created fertile ground for transitioning from highly passive transmission grid and end users, toward more responsive mode facilitated by small-scale distributed generation, load demand responsive technologies, and various switches which could be use for direct line flow control within the grid, as well as by the distributed network embedded systems. Similarly, industry restructuring is based on more decentralized decision-making by the end-users and power producers than current top-down feed-forward scheduling practices for the assumed load demand. It is becoming increasingly clear that transmission grid would become an active decision maker, and that it would, consequently, rely more and more on its own control and decision tools. As a result, once highly vertically integrated hierarchical operating practice is evolving into an architecture with many distributed decision makers, following their own objectives, and co-functioning within an electrical interconnection.

Major hidden problems in operating electric power grids according to current practices

To start with, the differentiation between normal and emergency operating conditions is questionable. System dynamics could be unstable either as a result of unusual generation, delivery, and consumption patterns and/or as a result of a large equipment failure under generation, delivery, consumption patterns for which the system was planned and designed. Therefore, it is generally not possible to follow qualitatively different decision rules in these two cases. What is needed, are highly adaptive methods, which adjust as the system conditions deviate from those for which the system was designed and basic tuning of controllers was done. As described in the Appendix, all automated controllers are currently designed to respond to very small changes around the anticipated conditions. Once this is no longer the case, their effects are hard to predict, and they could do more harm than good (the infamous example of several blackouts caused by the wrong logic of on-load tap changing transformers). Moreover, tuning of the primary controllers is typically done one at a time while representing the rest of the complex grid by a very simple equivalent, whose parameters could vary drastically under large deviations in

system conditions away from those assumed. It is very difficult to predict the controllability and observability of the actual power grid equipped with such controllers. There have been cases of various controllers ``fighting'' each other. This problem is likely to grow over time, as the transmission lines and end-users are responding at the same time as the power plants attempt to control their outputs.

Related to industry restructuring, system dynamics are affected by the control in response to technical conditions, as well as in response to economic signals, such as the price of electricity. At this point in time, no active R&D exist toward modeling power system dynamics driven by signals beyond strictly technical. Some major challenges to the electricity markets that ensure QoS required by the end-users could be posed by viewing dynamics of interest this way.

Finally, interdependencies between the short-term operating practices and longer-term system evolution with the right incentives for desired performance must be studied. The grid is generally designed to have considerable back-up capacity in case unexpected events occur. The cost of these back-up resources prior to industry restructuring was evenly distributed among all customers. In other words, the cost of managing uncertainties was not allocated to those causing uncertainties. As the industry moves forward, major work must be done to develop sustainable notions of value-based reliability and QoS . Without these, it will be impossible to provide mechanisms for differentiated QoS and reliability to those willing to pay for back-up services. It is possible to pose the problem of new technology- and industry restructuring-driven power system evolution as a discrete event-driven complex system dynamics. Depending on the specific industry structure and technologies in place, the degree of distributed decision-making will vary. Nevertheless, posing the problem this way opens enormous opportunities for assessing system performance and for designing multi-rate controllers in response to technical, economic and industry organizational conditions for achieving well-understood performance at various levels of the industry structure.

There is a need to develop test beds for the assessment of vulnerabilities and interdependencies of critical infrastructures. This would include testbeds for red teaming and response exercises by the research community. For obvious reasons, it is important that models of existing infrastructures not be made available. The task of developing test beds is a difficult and painstaking one but is one, which needs to be pursued. The FAA has taken such steps in the development of its testbeds at the FAA Research Center in New Jersey, but models of several sectors with interdependencies need to be developed for experimentation. The research problems here are not unlike those encountered by the Information Assurance and Survivability community in their search for experimental networking test beds. An exemplar of a proposal by NAI Laboratories to develop a national Distributed Denial of Service (DDOS) testbed is attached as Appendix E to this report. For reasons having to do with the sensitivity of this topic this particular recommendation was not debated fully at this meeting. The NISAC (National Infrastructure Simulation Center) has an important role to play in the development of such testbeds. There is a clear need for validated modeling and simulation of critical infrastructures, to facilitate an understanding of the propagation of attacks through

interdependencies between infrastructures and a sense of how to confine the effects of an attack to a part of a single infrastructure. The experiences that were presented by the power engineering community were especially valuable in this regard, since there has been extensive work since the NY blackouts of 1973 to develop alert mode and emergency mode operations of power systems. The work begun at the National Infrastructure Simulation Center (NISAC) at the Sandia National Laboratories is also valuable in this regard.

Our recommendations for investment in validated modeling, simulation and visualization of critical infrastructures falls into the following areas:

- **New Modeling and Simulation Tools Development for the Simulation of Hybrid Systems**: These are systems combining multiple models of computation, Some key difficulties with current approaches to simulation is the lack of tools to model continuous and discrete modes of computation and interconnections between them. The physical world is often best modeled by continuous time differential or partial differential equations. On the other hand, protocols, concurrency models and software are best modeled by finite state machines, Petri Nets, synchronous data flow or other such models of computation. Thus, complex interconnected systems such as infrastructure systems are best described by hierarchical systems of continuous and discrete models of computation. There are many technical difficulties in simulating such hybrid systems, since they do not have many of the desirable properties of continuous time or discrete time continuous state space models such as unique solutions, continuous dependence of solutions on initial conditions and robustness of the solutions to the identified model. Hybrid models of computation on the other hand are able to predict the kinds of cascading chains of events and propagations of degraded modes of operation across a complex system. The richness of such trajectories of hybrid systems makes them ideal candidates for use in the simulation of critical infrastructures. Domain specific tools for the infrastructure owners and operators would be important to develop.

- **Tools for the Assessment of the Level of Risk Posed to an Infrastructure by an Attack** on another infrastructure and response in the form of islanding interdependencies when under attack. Interdependencies grow into infrastructures from the desire to share information and services between different networked systems. They are added frequently for reasons of convenience during normal modes of operation of the infrastructure. However, under attack it is important to assess the threat condition present and then to dial down interconnections and interdependencies based on the level of the threat. Assessment of the threat level in an infrastructure is analogous to the network monitoring issues discussed under the title of Information Assurance and Survivability. Techniques for confinement of faults, including approaches such as islanding for fault isolation, and recovery from attack are important attributes of fault resistant networks and need to be developed.

- **Development of simulation test beds for red teaming exercises and response preparation and assessment.** It was felt that a good private public partnership with the stake holders, government labs and other entities like Sandia National Laboratories to develop simulation test beds which were not so sensitive as to be unusable by the open source community to use in red and blue teaming. The efforts of the Gartner group as well as those of numerous government studies of the vulnerabilities of critical infrastructures was to be lauded but it was felt that a greater level of engagement was needed by the open source community to study and develop countermeasures for key vulnerabilities.

4. Technology Transition Recommendations

The participants shared the concern about not only having a robust research strategy but also having robust public private partnerships as suggested by the National Strategy to Secure Cyberspace. The group had some specific points to discuss about how these public private partnerships could be better enabled:

Procurement by the Government as a lead customer—how can the Government be a good lead customer?

The suggestions offered included:

3.5 Assuring that there are no export controls for secure products.

3.6 Working to minimize a culture of Government Off the Shelf (GOTS); creating a culture of product certification. Also not all security products need to be in this category.

3.7 Using the Y2K experience to focus the attention of company CEOs and CFOs on the economic importance of cyber security.

3.8 Building on open standards since security problems are global and not only national.

3.9 Urging the government to explore the use of interoperable security technology between the government and commercial sectors to allow for the simultaneous development of commercial and government-only products. Although the Government would like to use Commercial Off the Shelf (COTS) in some cases more secure products (at greater cost) may be needed due to additional security concerns.

3.10 Promoting consistency in buying cyber secure products.

3.11 Developing and ensuring confidentiality, reliability and integrity as a critical part of secure systems.

1. **Is the strategy adequate for more determined attackers than just hackers?**

We agree that the draft document outlining the national strategy forms a useful guideline or roadmap toward a national strategy, but there is still much work ahead in forging a winning strategy to protect the nation's IT infrastructure from concerted cyber attack. Major improvements and extensions are needed to flesh out the national strategy, to raise the level of awareness about the topic, and to begin to set firm directions that the strategy might take.

In particular, the strategy does not adequately address the threat of a serious attacker such as a nation-state or well-funded terror group. There are major threats against U.S. cyberinfrastructure that could have overwhelmingly negative effects on the U.S. economy. The September 11th attacks caused $40.2B in insured losses, but the U.S. equity markets lost over $1T in value. A concerted cyber attack by a serious adversary could easily destroy as much or more value by shaking confidence in the U.S. economy, due to the extreme asymmetric leverage today's computer networks make available to a talented cyber adversary.

Issues to be addressed in greater detail include:

- Risk Assessment
- Modeling of Threats
- Attribution (and Retribution)
- Time Scale of Response and Notification
- Cyberpanel for Network Weather and threat levels on the network

2. Can public private partnerships be utilized?

The group was overwhelmingly in favor of public private partnerships for:

a. Investing in the development of research and prototypes previously funded in Government-sponsored research.
b. Testbeds for addressing scalability and interoperability of security solutions; for example, IPv6 with enhanced security, generalized peer-to-peer, grid computing, and other public private initiatives.
c. Federal funding models for ISACs. There was a feeling that not all ISACs are strong..
d. The need for a prototype model ISAC.
e. New models for early investment and co-investment with venture capital or corporate investors in development of security products.

3. Engagement Models with Stakeholders: How do secure products show up in infrastructures?

Establishment of best practices (akin to the Malcolm Baldridge/ISO) around specific business scenarios

a. Access Control
b. Remote Access

c. Wireless Access
d. Network Availability/Business Continuity
e. Network/Security management

4. **Privacy and Security**
   Privacy is a big consideration for confidence building in new markets in the U.S. and sales overseas. Strong guarantees of privacy in our National Strategy to Secure Cyberspace will enable greater nationwide compliance, information sharing between at risk industrial players, and rapid responses to new threats. Protecting privacy should not be an optional afterthought. Personal and organizational freedoms are as much or more at risk than our physical infrastructures. New technology solutions exist for not compromising on security at the expense of privacy. Participants felt that more emphasis could be given to the building on of privacy into secure products.

5. **Can we develop a national R&D strategy to support the cybersecurity strategy?**
   There was consensus about the need to develop a model for outsourced funding of research and development. This should be in addition to a government-funded national laboratory model. There was a great deal of support for a HS-ARPA model inside the Department of Homeland Security, with a planning process to determine a funding model (6.1-6.5 basic to applied to EMD), the size of individual awards and their mix, the kinds of programs and program management functions and transition funding. Such a research program should be coordinated with other agencies that are funding research in cybersecurity, such as the DoD and DARPA, NSF, NIST, DoE, etc. The question of how such a research program is coordinated with the designated mission lead agencies needs to be addressed.

6. **What were some significant omissions in the plan?**

a. Some technology trends that are of concern to industry are not addressed explicitly in the report, such as large-scale web service protection, and **securing networked embedded systems**. While the vulnerability of SCADA/DCS systems is a current concern, the infrastructure is evolving towards networked embedded systems that need to be made secure. This includes untethered communication and computing devices.

b. **Liability and Insurance Considerations**. Industry groups felt that the government ought to engage industry in the development of liability standards and consequently in assessing levels of risk and insurance needs.

c. The National Strategy needs to take a **firm stance in its recommendations**, rather than being advisory to sectors.

d. **ISPs as a sector** need to be given specific recommendations for cybersecurity as a separate critical sector.

e. The security strategy for infrastructure must address commercial needs in order to present **compelling incentives** for the main IT industry to address major security concerns.

Appendix A Background: The National Strategy to Secure Cyberspace

The background for the workshop was the National Strategy to Secure Cyberspace which had been released for comment by the Presidents Critical Infrastructure Protection Board (PCIPB) the previous day (September 18th, 2002) before the first day of the workshop by its Chair Mr. Richard Clarke and Vice Chair Dr. Howard Schmidt. The period of comment for this document is 60 days. The document is available at http://cybersecurity.gov . The development of this draft strategy has been ongoing for sometime now with a succession of town hall meetings planned and fifty-three clusters of key questions sent out for comment. We review the key elements of the draft here:

1. Cyberspace Threats and Vulnerabilities: A Case for Action
    a. Cyber-incidents are increasing in numbers, sophistication, severity and cost
    b. Economy is increasingly dependent on cyberspace and this has introduced new vulnerabilities and interdependencies and single points of failure
    c. Infrastructure Disasters have cascading impacts
    d. Fix vulnerabilities before emerging threats
    e. Past levels of cyber damage are not good indicators of future risk
    f. Everyone must secure their own piece of cyberspace
    g. Common Defense depends on a public private partnership
2. National Policies and Guiding Principles
    a. Federal Government will perform homeland and national security missions
    b. Lead agencies for each sector:
        i. DHS Information and Telecommunications, transportation, Postal and Shipping, Emergency Services, Continuity of Government
        ii. Treasury: Banking and Finance
        iii. HHS: Public Health, Food
        iv. DoE: energy, electric power, gas and oil production and storage
        v. EPA: water chemical industry and hazardous materials

vi. Agriculture: agriculture, food
vii. DoD: defense industrial base

c. State and Local governments to maintain order and deliver public services
d. Private Sector to endure orderly functioning of the economy:
    i. Avoid regulation
    ii. Safeguard Civil Liberties and Privacy
    iii. Enhance Public private partnership

The goal of the overall national strategy goal is to empower all Americans to secure their portions of cyberspace through the following means:

1. **Awareness and Information**
    a. Home users and businesses have an important role by securing their own computer systems
    b. PCIPB's Awareness Committee should foster a public private partnership to develop and disseminate cybersecurity materials
    c. State and local governments should identify guidelines covering cyber awareness, literacy, training and education.
2. **Technology and Tools**
    a. A public private partnership should develop best practices and new technology to increase security of Digital Control Systems (DCS) and Supervisory Control and Data Acquisition (SCADA) systems.
    b. PCIPB should coordinated with Director of OSTP to develop a program of research and development including intrusion detection, internet infrastructure security, application security, denial of service and other areas which are listed on the InfoSec Research Council's hard problems list
    c. Public Private partnerships for identifying cross-sectoral cyber and physical interdependencies and reduce vulnerabilities. The National Infrastructure Simulation and Analysis Center could help with modeling efforts.
3. **Training and Education**
    a. States should consider expansion of the Federal Cyber Corps (scholarships for service) program for undergraduates and graduate students.
    b. The CIO council and Federal agencies should consider establishing a Cyberspace Academy linking cybersecurity and forensics training programs
    c. Explore feasibility of a nationally recognized certification program for cybersecurity personnel.
4. **Roles and Partnerships**
    a. CEOs should consider forming integrated security councils
    b. State and local government should consider establishing IT security programs for their departments and agencies including awareness, audits, and standards.
    c. Internet service providers should consider adopting a "code of good conduct".

d. Federal government should identify and remove barriers to public private information sharing to promote cyberspace security.
e. Colleges and Universities should consider establishing information sharing and analysis centers (ISACs) to deal with cyber attacks and vulnerabilities.

5. **Federal Leadership**
    a. To enhance the procurement of more secure IT products with a comprehensive program review of the National Information Assurance Program (NIAP) by 4QFY03
    b. Expand the use of automated enterprise wide security assessment and security policy enforcement tools, and deploy threat management tools to pre-empt attacks.
    c. Consider the cost effectiveness of scenario based security and contingency preparedness exercise. Resultant weaknesses are to be included in the Government Information Security Reform Act (GSRA).
    d. Study and respond to new vulnerabilities through wireless communications
    e. Produce annual IT security audits.
6. **Coordination and Crisis Management**
    a. ISPs, hardware and software vendors, IT security related companies, CERTs and ISACs should consider establishing a Cyberspace NOC. The NOC is to be private bur co-managed with coordination with the Federal government.
    b. Industry should in partnership with the Federal government complete and regularly update cybersecurity crisis contingency plans.
    c. The law enforcement and national security community should develop a system to detect a national cyber attack and to plan an immediate response.
    d. Owners and operators of information system networks and network data centers should develop remediation and contingency plans to reduce the consequence of large-scale physical damage to facilities supporting the networks.
    e. The US should work with individual nations and non-governmental organizations to promote the establishment of national and international watch and warning networks.

The strategy provides a roadmap for groups of the American people divided into five different audience levels:

- **Level 1: Home Users and Small Businesses**. The recommendations here include
    - Install firewall software for home DSL or cable modem usage
    - Home owners and small businesses are encouraged to use regularly updated anti-virus systems
    - Recommend programs for filtering spam
    - Operating system updates for the home and small business user

- ISPs and other software vendors should make it easy to obtain security software and updates

- **Level 2: Large Enterprises**
  - CEOs should consider forming enterprise wide corporate security councils
  - CEOs should consider independent security audits, remediation programs and best practice reviews
  - Diversity in IT service providers to mitigate risk
  - Develop IT security and best practices. Share information on IT security through an appropriate Information Sharing and Analysis Center (ISAC).
  - Public private Awards programs for progress in cybersecurity
  - Review mainframe security software and procedure to ensure that effective technology and procedures are being utilized

- **Level 3:Federal Government**. A key step is to understand the current state and effectiveness of security and privacy controls, and maintain this through a cycle of risk assessment as feature in the Government Information Security Reform Act, GISRA of 2000. Current gaps and weaknesses include
  - Lack of senior management attention
  - Lack of Performance measurement
  - Poor security education and awareness
  - Failure to fully fund and integrate security into capital planning and investment control
  - Ensuring that contractor services are adequately secure
  - Failure to detect, report and share information on vulnerabilities.
  - Inadequate authentication
  - Inconsistent contingency planning

  The recommendations include
    - Enhance procurement of secure IT products by the Federal government after conducting a review using the National Infrastructure Assurance Program (NIAP)
    - Consider having private sector security service providers be certified
    - Use E-government model to explore benefits of cross government acquisition, operation and maintenance of security tools and services
    - Use the ongoing E-authentication initiative to provide better physical and logical access .tools and authentication mechanisms
    - Expand the use of automated enterprise wide security policy assessment and enforcement tools for federal agencies
    - Assess the use of VPNs, private line networks, etc.
    - Federal government should lead in the adoption of secure network protocols
    - OMB will determine with the CIO council on a case-by-case basis whether to employ a lead agency concept for government wide security measures including GSA, NIST, Department of Homeland Security, and DoD.

- **Level 3: State and Local Governments**
    - State and local governments should consider establishing IT security programs for their departments and agencies including awareness, audits, and standards
    - States and local governments should consider participating in the established information sharing and analysis centers (ISACs).
    - State and local governments should expand training programs in computer crime for law enforcement officials, including judges, prosecutors, and police.
- **Level 3: Critical Sectors: Higher Education**
    - Each college and university should consider establishing a point of contact to ISPs and the law enforcement officials in the event that the school's IT systems are discovered to be launching cyber attacks.
    - Colleges and universities should consider establishing ISACs to deal with cyber attacks and vulnerabilities, model guidelines empowering CIOs to address cybersecurity.
- **Level 3: Private Sectors**
    - Each sector should consider establishing ISACs with cooperative agreements and analysis and warning centers
    - Each sector should consider a technology and R&D gap analysis
    - Each infrastructure sector group should consider developing best practices for cyber security and work on security awareness campaigns
    - Each sector should establish mutual assistance programs for cybersecurity emergencies.
- **Level 4: National issues and efforts**
    - *Securing Shared Systems*
        - Securing mechanisms of the internet
            - Public private partnership for S-BGP, DNS-SEC and others protocols to be implemented
            - Secure router technology
            - ISPs should consider a "code of good conduct"
            - Fundamental technology needs for the internet
        - SCADA/DCS research
            - New approaches, technology and practices for plugging vulnerabilities in DCS/SCADA
            - Prioritized plan for improving security of SCADA/DCS possibly starting with DoE's 21 Steps to Improve Cybersecurity of SCADA networks
        - Highly secure and trustworthy computing
            - R&D committee of PCIPB should conduct gap analysis
            - Develop near term (1-3 years), mid term (3-5 years) and long term (5 years out and longer) research plans.
            - Federally funded programs including Internet infrastructure security, application security, denial of

service, communication security, SCADA security and secure systems composition. (See Info Sec Research Council's Hard Problems List discussed below in this report).
        - Private sector should consider research funding.
    - Securing Emerging Systems
        - Vulnerabilities of wireless systems and networks
    - Vulnerability Remediation
        - Clearing house for patch implementation
        - More secure "out of box" implementation of products
        - Promulgate best practices and methodology promotion integrity, security and reliability.
- *Fostering a Reinforcing Economic and Social Framework*
    - Awareness
    - Training and Education
        - CyberCorps Scholarship programs
        - Cyberspace Academy for computer forensics
        - Cyberdefenders and other red team activities
        - PCIPB should consider multi-department corps of IT and cybersecurity specialists
        - State and local officials should develop programs for primary and secondary schools.
    - Certification
    - Information Sharing
    - Cybercrime
        - Encourage the reporting of cybercrime
        - Coordination by FBI and Secret Service
        - Improved Information Sharing between Federal, State and local authorities
        - Collect survey data on cyber crime to establish a base line
    - Market forces
        - Review Federal and State regulations which may impede market forces from contributing to cybersecurity
        - PCCIB working with insurance industry to develop risk assessment, modeling and loss economics
        - Corporations should consider disclosing the identity of their IT security audit firm
    - Privacy and Civil Liberties
        - Consult with privacy advocates in implementation of security solutions
        - Implementation of Gramm, Leach, Bliley Financial Modernization Act and Health Insurance Portability and Accountability Act
- *Developing National Plans and Policy*
    - Analysis and warning

        - Establishment of Cyberspace Network Operation Center (NOC) in cooperation with CERTs, ISACs to support health and reliability of operations.
        - Cyber Warning Information Network to key government and non-government cyber-security related network operation centers.
    - Continuity of operations, reconstitution and recovery
        - Voluntary partnership of Federal government with industry to develop recovery plan
        - Emergency plans with local and emergency authorities
    - National security
        - Establish program to counter cyber based intelligence gathering against US government, industry and universities.
        - Improve understanding of incidence response and coordination between law enforcement, national security and defense agencies
        - Capability to attribute threats, attacks and actions to suppress threats.
        - Capability to respond in appropriate fashion when nation states or terrorist groups threaten vital interests.
    - Interdependency and physical security
        - Validated models for simulation of interdependencies
        - Remediation steps for large utility, infrastructure owners to reduce damage after attack
- **Level 5: Global Issues**
    - Work with other nations, private sector, governmental, nongovernmental and international agencies to foster the development of watch and warning networks
    - Encourage nations to accede to the Council of Europe Convention on Cybercrime.
    - Secure North American critical infrastructures with Canada and Mexico.
    - Develop and foster global “security culture”.
    - Encourage appointment of national cyberspace coordinator within each country.
    - Draw on global Science and Technology base.

Appendix B: AGENDA

September 19, 2002

8:30- 8:45 ***Welcome*** - Peter Freeman, National Science Foundation

8:45-9:15 **Keynote Address** - Richard Russell, Associate Director for Technology, White House Office of Science and Technology Policy

9:15-9:45 ***National Strategy to Secure Cyberspace*** - Thomas Cabe, Homeland Security Office

10:00-10:30 ***CIP Hard Problems*** - Doug Maughan, DARPA ATO

10:30 - 10:40 ***Introduction to NSF-OSTP Workshop*** -
Mark LeBlanc, White House Office of Science and Technology Policy & Helen Gill, National Science Foundation

10:40 -11:00 ***Charge and Goals of the Workshop*** - Shankar Sastry, UC Berkeley

11:00-12:00 ***CIP and Networked Embedded Systems***
Panel chaired by Helen Gill, National Science Foundation
- ***Networked Embedded Systems and Software Control***
Janos Sztipanovits, Vanderbilt University
- ***Open Distributed Computing Platforms***
Doug Schmidt, DARPA
- ***Strategic View of Networking Research***
Mari Maeda, National Science Foundation

1:00-2:00 ***Models and analysis of Interdependences between IT and Critical Infrastructure***
Panel chaired by Sam Varnado, Sandia National Laboratories
Panelists:
- Steven Wicker, Cornell University
- Linda Nozick, Cornell University
- Miriam Heller, National Science Foundation (not able to attend)

2:00 - 3:00 ***Overview of Critical Infrastructure Systems: Power Grid and SCADA***
Panel chaired by Massoud Amin, EPRI
Panelists:
- Jose R. Gracia, Tennessee Valley Authority
- Robert Hutchinson, Sandia National Laboratories
- Marija Ilic, Carnegie Mellon University
- Robert Thomas, Cornell University

3:30 - 5:00 **Breakout sessions**
Power Grid Issues (Chaired by Robert Thomas, Cornell University)
SCADA Issues (Chaired by Roy Maxion, Carnegie Mellon University)
Interdependency Issues (Chaired by Sri Kumar, DARPA)

September 20, 2002

8:30-9:00 ***Welcome, Charge and Introduction to Day 2***
Helen Gill, National Science Foundation & Shankar Sastry, UC Berkeley

9:00- 9:20 Keynote: ***Protecting the Federal Aviation Administration from Cyber-Attack*** - Art Pyster, FAA

9:20- 9:40 Keynote: ***Aviation Safety and Efficiency*** – Marshall Potter, FAA

9:40-10:00 ***Security with Privacy*** -Doug Tygar, UC Berkeley

10:00-11:00 ***Protecting Critical Infrastructure from Cyber-terrorism***
Panel chaired by Jack Stankovic, University of Virginia
Panelists:
- French Caldwell, Gartner, Inc.
- Richard Hunter, Gartner, Inc.
- Jose R. Gracia, Tennessee Valley Authority

11:00 -12:00 ***Overview of Critical Infrastructure Systems: Airspace Management Panel Discussion***
Panel Discussion: CHAIR Feisal Keblawi, Office of Research and Acquisitions (ARA), FAA
Panelists:
- Tim Wallace, FAA
- David Sharp, Boeing Phantom Works
- Chip Meserole, Boeing Air Traffic Management
- Edward Lee, UC Berkeley

1:00-2:00 **Breakout Groups:** ***Airspace Management***

Air traffic control Ground Control Issues (Chair: Jack Stankovic, University of Virginia)
Flight Control Onboard Issues (Chair: Jon Ward, Rockwell Collins Advanced Technology Center)
Interdependency with other transportation modalities issues (Chair: Joseph Cross, Lockheed Martin Co.)

2:00-3:30 **Networked Embedded Systems**
Chaired by Steve Wicker, Cornell
- Anish Arora, Ohio State University
- Bhaskar Krishnamachari, University of Southern California
- William Merrill, Sensoria Corporation
- David Nicol, Dartmouth College & ISTS

3:30 - 4:30 **Wrap Up Discussion on Network Embedded Systems and CIP**
Moderators: Helen Gill, Shankar Sastry, Jack Stankovic, Janos Sztipanovits

4:30 - 5:00 Summary for US-EU meeting (Helen Gill / Shankar Sastry) and CRA meeting (John Stankovic)

Appendix C PARTICIPANT LIST

Participant List: US Technical Workshop on Information Technology for Critical Infrastructure Protection

S Kamal Abdali
National Science Foundation
kabdali@nsf.gov

Massoud Amin
Electric Power Research Institute (EPRI)
mamin@epri.com

William Arbaugh
University of Maryland, College
Parkwaa@cs.umd.edu

Anish Arora
Ohio State University
anish@cis.ohio-state.edu

John Bay
DARPA
jbay@darpa.mil

Terry V. Benzel
Network Associates
tbenzel@nai.com

Askold Boretsky
MITRE
aboretsk@mitre.org

Raymond Bortner
Air Force Research Laboratory, Air Vehicles Directorate
raymond.bortner@wpafb.af.mil

Lawrence E. Brandt
National Science Foundationlbrandt@nsf.gov

Kenneth Button
George Mason University
kbutton@gmu.edu

Tom Cabe
Homeland Security Office
tcabe@cybersecurity.gov

French Caldwell
Gartner, Inc.
Not Available

Jagdish Chandra
George Washington University
jchandra@seas.gwu.edu

Melvyn Ciment
Consultant to Hewlett-Packard
mel@ciment.com

Edwin R. Coover
MITRE Corporation
coover@mitre.org

Tracey Goyette Cote
Institute for Information Infrastructure Protection
tcote@ists.dartmouth.edu

Tim Courington
Federal Aviation Administration
tim.courington@auatac.com

Joseph Cross
Lockheed Martin
joseph.k.cross@lmco.com

Francis C. Deckelman
Office of Naval Research
deckelf@onr.navy.mil

Sartaj S. Dhami
Federal Aviation Administration
sartaj.dhami@auatac.com

Darleen Fisher
National Science Foundation
dlfisher@nsf.gov

Stephanie Forrest
University of New Mexico
forrest@cs.unm.edu

Peter Freeman
National Science Foundation
freeman@cc.gatech.edu

Cita Furlani
National Coordination Office for Information Technology Research & Development
furlani@itrd.gov

Helen Gill
National Science Foundation
hgill@nsf.gov

Sean Gorman
Geotel
seangorman@geo-tel.com

Jose R. Gracia
Tennessee Valley Authority
jrgracia@tva.gov

Valerie Gregg
National Science Foundation
vgregg@nsf.gov

Richard Han
University of Colorado
rhan@cs.colorado.edu

Carl Hauser
Washington State University
hauser@eecs.wsu.edu

Sharon Heise
DARPA/IXO
saheise@darpa.mil

Miriam Heller
National Science Foundation
mheller@nsf.gov

Ian Hiskins
University of Wisconsin – Madison
hiskens@engr.wisc.edu

Sally Howe
National Coordination Office for Information Technology Research and Development
howe@itrd.gov

Dimitrios Hristu
University of Maryland, College Park
hristu@glue.umd.edu

Richard Hunter
Gartner, Inc
Not Available

Robert Hutchinson
Sandia National Laboratories
rlhutch@sandia.gov

Marija Ilic
Carnegie Mellon University
ilic@mit.edu

Feisal Keblawi
Federal Aviation Administration
feisal.keblawi@faa.gov

Pradeep K. Khosla
Carnegie Mellon University
pkk@ece.cmu.edu

Steven King
DUST (S&T)
steven.king@osd.mil

Eva Kingsbury
National Science Foundation
ekingsbu@nsf.gov

Jana Kosecka
George Mason University
kosecka@cs.gmu.edu

Bhaskar Krishnamachari
University of Southern California
bkrishna@usc.edu

Sri Kumar
DARPA
skumar@darpa.mil

Nicholas Kyriakopoulos
George Washington University
kyriak@seas.gwu.edu

Carl Landwehr
National Science Foundation
clandweh@nsf.gov

Mark LeBlanc
White House Office of Science and Technology Policy
mleblanc@ostp.eop.gov

Edward Lee
University of California, Berkeley
eal@eecs.berkeley.edu

Ernie Lucier
Federal Aviation Administration
ernest.lucier@faa.gov

Mari Maeda
National Science Foundation (CISE)
mmaeda@nsf.gov

Stephen R. Mahaney
National Science Foundation
smahaney@nsf.gov

Daniel F. Massey
University of Southern California / Information Sciences Institute
masseyd@isi.edu

Doug Maughan
DARPA
dmaughan@darpa.mil

Roy A. Maxion
Carnegie Mellon University
maxion@cs.cmu.edu

William M. Merrill
Sensoria Corporation
williamm@sensoria.com

Jere S. Meserole
Boeing Air Traffic Management
chip.meserole@boeing.com

Lamine Mili
Virginia Polytechnic Institute & State University
lmili@vt.edu

Kevin Mills
National Institute of Standards & Technology (NIST)
kmills@nist.gov

Al Mok
University of Texas at Austin
mok@cs.utexas.edu

Erica Layne Morrison
University of California, Berkeley
layney@eecs.berkeley.edu

Dr. Priscilla P. Nelson
National Science Foundation
pnelson@nsf.gov

Kyle Nelson
Honeywell Labs
kyle.nelson@honeywell.com

Cliff Neuman
University of Southern California /
Information Sciences Institute
bcn@isi.edu

David Nicol
Dartmouth College & ISTS
nicol@cs.dartmouth.edu

Linda Nozick
Cornell University
lkn3@cornell.edu

Marshall Potter
Federal Aviation Administration
marshall.potter@faa.gov

Art Pyster
Federal Aviation Administration
arthur.pyster@faa.gov

Stanley Riveles
Office of the S&T Adviser to the
Secretary of State
rivelessa@t.state.Gov

Richard Russell
White House Office of Science and
Technology Policy
Not Available

Shankar Sastry
University of California Berkeley
sastry@eecs.berkeley.edu

Peter W. Sauer
University of Illinois at Urbana-
Champaign
sauer@ece.uiuc.edu

Sami Saydjari
Stanford Research Institute International
ssaydjari@csl.sri.com

Doug Schmidt
DARPA
dschmidt@darpa.mil

Roman Shaffer
US Nuclear Regulatory Commission
ras3@nrc.gov

David Sharp
Boeing Corporation
david.sharp@boeing.com

Nozer Singpurwalla
George Washington University
nozer@research.circ.gwu.edu

Jonathan N. Smith
University of Pennsylvania
jms@cis.upenn.edu

Jonathan M. Sprinkle
Vanderbilt University
Jonathan.Sprinkle@vanderbilt.edu

John A. Stankovic
University of Virginia
stankovic@cs.virginia.edu

Victoria Stavridou
Stanford Research Institute International
victoria@sdl.sri.com

Gary W. Strong
National Science Foundation
gstrong@nsf.gov

Mark Swinson
Sandia National Laboratories
mlswins@sandia.gov

Janos Sztipanovits
Vanderbilt University
Janos.Sztipanovits@vanderbilt.edu

Simon Szykman
National Institute of Standards and Technology (NIST)
sszykman@ostp.eop.gov

Robert J. Thomas
Cornell University
rjt1@cornell.edu

Doug Tygar
University of California, Berkeley
tygar@cs.berkeley.edu

Sam Varnado
Sandia National Laboratories
sgvarna@sandia.gov

Joe Veoni
MITRE CAASD
jveoni@mitre.org

Tim Wallace
Federal Aviation Administration
Timothy.S.Wallace@faa.gov

Jon Ward
Rockwell Collins Advanced Technology Center
jaward1@rockwellcollins.com

Carmen Whitson
National Science Foundation
cwhitson@nsf.gov

Steve Wicker
Cornell University
wicker@ece.cornell.edu

Felix Wu
University of California Davis
wu@cs.ucdavis.edu

Doug Xuan
Ohio State University
xuan@cis.ohio-state.edu

Wei Zhao
Texas A&M University
w-zhao@tamu.edu