Layered Assurance Workshop
KEYNOTES and Invited Talks

Supposing we got a "Do over"?

Most of our computer hardware and systems designs are the legacy of a bygone era when resources were scarce and every transistor or line of code had the sole purpose of increasing performance. That has led us to today where our biggest problem is the security, integrity and robustness of our computer systems. In this talk, I'll describe two programs sponsored by DARPA, CRASH (Clean-slate Design of Resilient, Adaptive, Secure Hosts) and MRC (Mission-oriented Resilient Clouds), that are taking a clean-slate approach to design systems (from hardware, through system software, middleware and application code) that are both highly resistant to attack and capable of operating usefully even after attacks have succeeded.

Howie Shrobe is a long term member of the MIT Computer Science and Artificial Intelligence Laboratory. He is currently serving his second term at DARPA, having previously served from 1994 - 1997. During his first term, he initiated the Information Survivability effort. He is currently leading two Cyber Security programs: CRASH, Cleanslate Design of Resilient Adaptive Hosts and Mission-oriented Resilient Clouds (MRC). Both draw on biological and social inspirations. CRASH is trying to develop a clean-slate host architecture that uses novel hardware and operating system techniques to erect strong barriers to penetration. CRASH goes beyond this to build adaptive protections that detect, diagnose, recover and patch novel attacks. These two approaches mimic the "innate" and "adaptive" components of the immune system. Finally diversity is used to raise the attacker's work factor further. In MRC these same ideas are applied, but at a community scale, mirroring the elements of a public health system.

Using Eliminative Induction and Defeasible Reasoning to Assess Assurance Case Confidence

Given a claim supported by some argument or evidence, how much confidence should we have in the truth of the claim? What does “confidence” mean in this case? Should it be the likelihood that the claim is true, or should it be the degree of belief that the claim is true. We propose an approach that focuses on establishing the degree of belief one should have in a claim. The approach identifies reasons for doubting the truth of a claim and then showing how confidence grows as each of these sources of doubt are eliminated. The approach is illustrated with a simple example. Considerable additional work is needed to assess its practicality, but it does provide a way of thinking about assurance case confidence that has several benefits.

John Goodenough is a Fellow of the Software Engineering Institute (SEI) and most recently led a research project on system of systems software assurance. He and his team have worked with assurance cases for almost ten years, investigating their application to assuring the safety of medical devices, assuring the security of the software supply chain, and evaluating whether a software project is developing the right artifacts to be successful. He has served as the SEI’s Chief Technical Officer and is a Fellow of the Association for Computing Machinery (ACM).

A Layered Assurance Perspective: Lessons from the Formal Analysis of Fault-Tolerant Systems

A key component of any approach to layered assurance is reuse of both design artifacts and assurance arguments. Neither of these is a trivial undertaking. This talk will present some lessons learned from an attempt to develop both reusable design artifacts and assurance arguments for a family of fault-tolerant integrated modular avionics architectures.

Paul S. Miner is a senior research engineer in the Safety-Critical Avionics Systems Branch at NASA's Langley Research Center. He currently supports NASA's Aviation Safety Program as co-Technical lead for Validation and Verification of Flight-Critical Systems. His principal research interests are the development and application of formal methods for the analysis of safety-critical systems with a particular emphasis on the design and analysis of distributed systems. He was the principal architect for the Scalable Processor-Independent Design for Extended Reliability (SPIDER) family of fault-tolerant architectures developed at NASA Langley. He has been an active participant in the Networking and Information Technology Research and Development High Confidence Software and Systems Coordinating Group since 2003 (http://nitrd.gov). He was a member of the RTCA Committees that developed RTCA DO-254 "Design Assurance Guidance for Airborne Electronic Hardware" and RTCA DO-297 "Integrated Modular Avionics (IMA) Development Guidance and Certification Considerations." He holds a Ph.D. in computer science from Indiana University, an M.S. in computer science from the College of William and Mary, and a B.S. in computer science from Old Dominion University.

Bluespec Codesign Language: A Unified Language to Enable HW / SW Codesign

Embedded systems are almost always built with portions implemented in both hardware and software. Despite the increased demand for new embedded systems aimed at different points on the price-performance-power curve, all current industrial design methodologies make the decision regarding how the design should be split between hardware and software early in the design process. This leaves little room for design explorations without significantly delaying the time-to-market. This talk introduces Bluespec Codesign Language, a unified language for hardware-software design. We provide an easy and natural way of specifying which parts of the design should be implemented in hardware and which in software without obscuring important design decisions about the interfacing. We show how this can enable valuable design exploration of HW/SW systems.

Nirav Dave is currently a Computer Scientist at SRI International. He received his PhD from Massachusetts Institute of Technology in 2011. His research focuses on improving the abstraction, representation, and formal reasoning of hardware-software codesign via a unified language in an effort to enable extreme exploration of heterogeneous computing systems. He received his S.M. in Electrical Engineering and Computer Science in 2005 at Massachusetts Institute of Technology and his B.S. in Electrical and Computer Engineering in 2002 from Carnegie Mellon University.

Rigorous Component-based System Design Using the BIP Framework

Rigorous system design requires use of a single powerful component framework allowing a representation of the designed system at different levels of detail, from a high-level model of application software to its implementation. The use of a single framework allows one to maintain overall coherency and correctness by comparing different architectural solutions and their properties. In this talk, I will first present the BIP (Behavior, Interaction, Priority) component framework which encompasses an expressive notion of composition for heterogeneous components by combining interactions and priorities. This allows description at different levels of abstraction from high-level application software to mixed hardware/software systems. Second, I will introduce a rigorous design approach that uses BIP as a unifying semantic model to derive from an application software, a model of the target architecture and a mapping, a correct implementation. Correctness of implementation is ensured by application of source-to-source transformations in BIP which preserve correctness of essential design properties. The design is fully automated and supported by a toolset including a compiler, the D-Finder verification tool, and model transformers. We illustrate the use of BIP as a modeling formalism as well as crucial aspects of the design flow for ensuring correctness, through an autonomous robot case study.

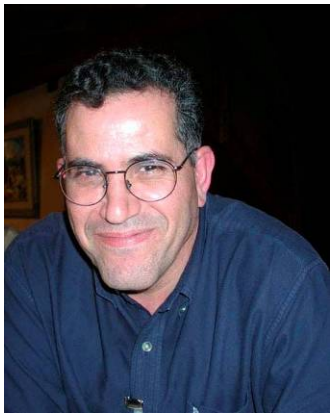

Saddek Bensalem is a professor at the University Joseph Fourier, Grenoble, France in the Distributed and Complex Systems Group, Verimag Laboratory. His research interests are in the study of formal specification and verification, algorithmic verification techniques, infinite-state systems, real-time systems, program verification, automata and logics. His work has been on the topic of automatic formal verification of systems that have unbounded numbers of states. He pioneered development of techniques that combine the complementary strengths of deductive and algorithmic methods. He developed the theory of property-preserving abstraction and relationship to compositional verification, automatic invariant discovery, predicate abstraction and a verification method for parametrized systems. All these techniques are employed in the tool InVeSt developed in collaboration with Yassine Lakhnech and Sam Owre from the SRI. Currently, he collaborates with Joseph Sifakis on the research project "Component based Construction - BIP". The aim is to develop the theory, methods and tools for building real-time systems consisting of heterogeneous components. The focus is on the following challenging problems : 1) Develop a framework for the incremental composition of heterogeneous components. Three different sources of heterogeneity are considered related to interaction, execution and abstraction. 2) Develop results ensuring correctness-by-construction for essential system properties such as mutual exclusion, deadlock-freedom and progress in order to minimize a posteriori validation. 3) Provide automated support for component integration and generation of glue code meeting given requirements.

Panel Discussion

The Future of Highly Trustworthy Systems, Networks, Apps, and Clouds

Peter G Neumann, SRI International – Moderator

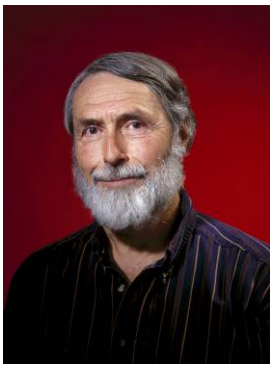

Peter G. Neumann has doctorates from Harvard and Darmstadt. After 10 years at Bell Labs in Murray Hill, New Jersey, in the 1960s, during which he was heavily involved in the Multics development jointly with MIT and Honeywell, he has been in SRI's Computer Science Lab since September 1971. He is concerned with computer systems and networks, trustworthiness/dependability, high assurance, security, reliability, survivability, safety, and many risks-related issues such as voting-system integrity, crypto applications and policies, health care, social implications, and human needs -- especially those including privacy. He moderates the Association for Computing Machinery (ACM) Risks Forum, has been responsible for Communications of the ACM's Inside Risks columns monthly from 1990 to 2007, tri-annually since then, chairs the ACM Committee on Computers and Public Policy, and chairs the National Committee for Voting Integrity. He created ACM SIGSOFT's Software Engineering Notes in 1976, was its editor for 19 years, and still contributes the RISKS section. He is on the editorial board of IEEE Security and Privacy. He has participated in four studies for the National Academies of Science: Multilevel Data Management Security (1982), Computers at Risk (1991), Cryptography's Role in Securing the Information Society (1996), and Improving Cybersecurity for the 21st Century: Rationalizing the Agenda (2007). His 1995 book, Computer-Related Risks, is still timely. He is a Fellow of the ACM, IEEE, and AAAS, and is also an SRI Fellow. He received the National Computer System Security Award in 2002 and the ACM SIGSAC Outstanding Contributions Award in 2005. He is a member of the U.S. Government Accountability Office Executive Council on Information Management and Technology, and the California Office of Privacy Protection advisory council. He co-founded People For Internet Responsibility (PFIR, http://www.PFIR.org). He has taught courses at Darmstadt, Stanford, U.C. Berkeley, and the University of Maryland. See his website (http://www.csl.sri.com/neumann) for testimonies for the U.S. Senate and House and California state Senate and Legislature, papers, bibliography, further background, etc.

Panel Participants:

*** Howard Shrobe, DARPA I20**
*** Rance DeLong, LynuxWorks**
*** W Mark Vanfleet, NSA/NCSC**
*** Greg Sullivan, BAE Systems**

LAW Chairman

Rance J. DeLong is a staff scientist at LynuxWorks, consultant, and adjunct professor at Santa Clara University. He has more than 35 years of software experience, with an emphasis on security and high-assurance systems. He is the consulting security expert for LynuxWorks' secure operating systems development. Mr. DeLong is a member of the Center for the Advanced Study and Practice of Information Assurance at Santa Clara University, where he teaches in the Computer Engineering graduate program. His research activities include methods and tools for principled approaches to the compositional construction and certification of high-assurance secure systems, and is a key contributor to ongoing MILS research at SRI International. Mr. DeLong is active in the leadership of The Open Group's Real-Time and Embedded Systems Forum. He is chairman of the MILS Application Programming Interface (API) Working Group and organizer of the Layered Assurance Workshop. He currently serves as a member of the External Oversight Group for the CRASH-worthy Trustworthy Systems Research and Development project, performed by SRI International and Cambridge University, a part of the DARPA CRASH program. Mr. DeLong holds a B.S. in Physics and a B.A. in Philosophy from Moravian College, and has done extensive post-graduate study in computer science at Lehigh University and Stanford University.