April 5, 2016

7:00-8:00 AM: Breakfast and Registration

8:00-8:15 AM : Welcome and Introductions - Neena Imam

8:15-9:15 AM : **Keynote 1:** Dr. Jeff Nichols, Associate Lab Director, Computing and Computational Sciences Directorate, Oak Ridge National Laboratory

Session 1: Nanomaterials for Future HPC

Since we are rapidly approaching the minimum effective size of CMOS circuitry, nano-electronics might hold the answers to increasing the device capabilities at reduced footprint and power. In this session we will discuss disruptive technologies such as Carbon nanotube based transistors, HPC applications of Carbon based nanomaterials such as Graphene, etc. We will also discuss the technical challenges of maintaining the advantageous properties of nanomaterials at very large scale.

Session Chair: Dr. Barney Maccabe, Computing and Computational Sciences Directorate,
Oak Ridge National Laboratory

9:15-9:50 AM: Dr. Victor Zhirnov, Semiconductor Research Corporation, *Future Information Processing*: *From Materials to Systems*

Abstract: It is becoming increasingly clear that in future information processing applications, synergistic innovations in materials, devices, and system architectures will be a key to achieving new levels of performance. We will examine the physics of extreme scaling of information processing devices and systems, with a focus on energy minimization and also discuss the materials and architectural implications on system scaling. The connection of the device physics in the Boltzmann-Heisenberg limits and the parameters of the digital circuits implemented from these devices will be explored. An abstraction of a Minimal Turing Machine built from the limiting devices will be used to provide insight on the "intelligence" that could be expected from a volume of matter.

Bio: Victor Zhirnov is Chief Scientist at the Semiconductor Research Corporation. His research interests include nanoelectronics devices and systems, properties of materials at the nanoscale, bio-inspired electronic systems etc. He has authored and co-authored over 100 technical papers and contributions to books. Victor Zhirnov served as the Chair of the Emerging Research Device (ERD) Working Group for the International Technology Roadmap for Semiconductors (ITRS). Victor Zhirnov also holds adjunct faculty position at North Carolina State University and has served as an advisor to a number of government, industrial, and academic institutions. Victor Zhirnov received the M.S. in applied physics from the Ural Polytechnic Institute, Ekaterinburg, Russia, and the Ph.D. in solid state electronics and microelectronics from the Institute of Physics and Technology, Moscow, in 1989 and 1992, respectively. From 1992 to 1998 he was a senior scientist at the Institute of Crystallography of Russian Academy of Science in Moscow. From

1998 to 2004 he was research professor at North Carolina State University. He joined SRC in 2004.

9:50-10:25 AM: Dr. Olga S. Ovchinnikova, Center for Nanophase Material Sciences, Oak Ridge National Laboratory, *Directing Matter: Towards Atomic Scale 3D Nanofabrication*

Abstract: To move beyond current technologies, design of both silicon-based and functional materials requires material control on the level of individual atomic configurations and single atom positions. This is also important for understanding and harnessing energy where disruptive improvements in efficiencies are required. Achieving this ultimate limit in science and technology of material and device performance–from quantum computing to nonvolatile memories, from thermoelectrics to superconductors, requires design and control of matter with atomic, molecular, and mesoscale fidelity by developing precision synthesis tools. Here we will discuss development of pathways to direct matter in a scalable fashion to fabricate/design three-dimensional structures of a variety of materials with atom-by-atom and defect control of their shape and composition.

Bio: Dr. Olga Ovchinnikova received her Ph.D. in 2011 from the University of Tennessee, Knoxville in Chemical Physics. Following her Ph.D. she worked as a postdoctoral associate and research staff scientist in the Organic and Biological Mass Spectrometry Group in the Chemical Sciences Division at Oak Ridge National Laboratory (ORNL) developing chemical imaging. Currently she is a research staff scientist in the Nanofabrication Research Laboratory at the Center for Nanophase Materials Sciences at ORNL investigating relationships between physical structure and chemical functionality at the nanoscale through the development of multimodal imaging platforms. Additionally, she works on engineering control in functional materials using helium and neon ion beams, and investigating the use of ion beams for atomically precise 3D manufacturing.

10:25-11:00 AM: Mr. Max Shulaker, Stanford University, *Transforming Nanodevices into Nanosystems: the N3XT 1,000X*

Abstract: Future computing demands far exceed the capabilities of today's electronics, and cannot be met by isolated improvements in transistor technologies, memories, or integrated circuit (IC) architectures alone. The N3XT (Nano-Engineered Computing Systems Technology) approach overcomes these challenges through recent advances across the computing stack: (a) transistors using nanomaterials such as one-dimensional carbon nanotubes (and two-dimensional semiconductors) for high performance and energy efficiency, (b) high-density non-volatile resistive and magnetic memories, (c) ultra-dense (e.g., monolithic) three-dimensional integration of logic and memory for fine-grained connectivity, (d) new architectures for computation immersed in memory, and (e) new materials technologies and their integration for efficient heat removal. N3XT hardware prototypes represent leading examples of transforming scientifically-interesting nanomaterials and nanodevices into actual nanosystems. Compared to conventional approaches, N3XT architectures promise to improve energy efficiency significantly, in the range of three orders of magnitude, thereby enabling new frontiers of high performance computing.

Bio: Max Shulaker is a Ph.D. candidate in Electrical Engineering at Stanford University, under the supervision of Professor Subhasish Mitra. He received his B.S. from Stanford University in Electrical Engineering. Max's current research interests are in the broad area of nanosystems. His research results include the demonstration of the first carbon nanotube computer (highlighted on the cover of Nature, Sept. 2013), the first digital sub-systems built entirely using carbon nanotube FETs (awarded the ISSCC Jack Raper Award for Outstanding Technology-Directions Paper, 2013), the first monolithically-integrated 3D integrated circuits combining arbitrary vertical stacking of logic and memory (IEDM 2014), the highest-performance CNFETs to-date (IEDM 2014), and the first highly-scaled CNFETs fabricated in a VLSI-compatible manner (IEDM 2015). Max also enjoys teaching and has been a guest lecturer in several classes at Stanford. He is a Fannie and John Hertz Fellow and a Stanford Graduate Fellow.

11:00-11:15 AM: Short Break

11:15-11:50 AM: Dr. Suman Datta, University of Notre Dame, *Band Engineered 3D Tri-gate Transistors for High Performance Computing*

Abstract: The 3D Tri-gate or FinFET technology market is expected to grow from $5 billion USD in 2015 to $35 billion by 2022, spanning across several technology nodes from 22nm, 14nm, 10nm to 7nm and a range of product offerings from CPUs, SoCs, FPGAs, GPUs, MCUs to Network Processors. 3D Tri-gate transistor enjoys fundamental advantages over its planar transistor counterpart with respect to superior performance at lower supply voltages, significantly higher drive current per layout width and improved intrinsic threshold voltage variation. In this talk, I will cover the basic operation principles of Tri-gate transistors from device physics and circuit design perspective. I will also present opportunities and challenges of advanced band engineered FinFET concepts such as Germanium and Compound Semiconductor based Quantum-Well FinFETs and stacked gate-all-around (GAA) nanowire FETs that are being actively pursued to extend technology scaling beyond 5nm node for HPC applications for the next decade.

Bio: Suman Datta is the Chang Family Chair Professor of Engineering Innovation in the department of Electrical Engineering at University of Notre Dame. He was previously a Professor of Electrical Engineering at Penn State University from 2007 to 2015. Before joining Penn State, from 1999 till 2007, he was in the Advanced Transistor Group at Intel Corporation, where he developed several generations of logic transistor technologies including high-k/metal gate, 3D Tri-gate and non-Silicon channel CMOS transistors. He continues to work with leading semiconductor companies on advanced CMOS transistor options such as band engineered FinFETs, inter-band Tunnel FETs, negative capacitance Ferro FETs and phase transistor FETs. His group also works on non-Boolean computing paradigm harnessing collective state of coupled dynamical systems. He was a recipient of the Intel Achievement Award (2003) for demonstration of high-performance high-k metal gate CMOS, the Intel Logic Technology Quality Award (2002) for invention of Tri-gate CMOS, the Penn State Engineering Alumni Association (PSEAS) Outstanding Research Award (2012), the SEMI Award for North America (2012), IEEE Device Research Conference Best Paper Award (2010, 2011) and the PSEAS Premier Research Award (2015). He has published over 210 papers and holds over 163 US patents. He is a Fellow of IEEE.

11:50 AM-12:25 PM: Dr. Ali Adibi, Department of Electrical and Computer Engineering, Georgia Tech, *Integrated Nanophotonic Structures for Optical Computing*

Abstract: The use of rich nonlinear dynamics in photonic resonators has emerged as a potential means for analog computing. In this talk, a review of the recent advance in optical computing will be presented, and the unique features of nanophotonic structures for forming analog computing modules in an integrated chip-scale platform will be presented. In addition, the recent achievements in the development of ultra-compact and high-performance resonance-based reconfigurable photonic devices and subsystems (as enabling tools for optical computing) will be demonstrated.

Bio: Ali Adibi is the director of Bio and Environmental Sensing Technologies (BEST) and a professor and Joseph M. Pettit chair in the School of Electrical and Computer Engineering, Georgia Institute of Technology. His research group has pioneered several structures in the field of integrated nanophotonics for both information processing and sensing. He is the author of more than 140 journal papers and 400 conference papers. He is the editor-in-chief of the Journal of Nanophotonics, and the nanophotonic program track chair of the Photonics West meeting. He is the recipient of several awards including Presidential Early Career Award for Scientists and Engineers, Packard Fellowship, NSF CAREER Award, and the SPIE Technology Achievement Award. He is also a fellow of OSA, SPIE, and AAAS.

12:25-1:30 PM: Lunch and Networking

Session 2: Quantum Computing

Quantum computing paradigms challenge our traditional understanding of nature by utilizing quantum phenomenon called superposition of states. In theory, quantum computing can provide parallel processing capabilities that cannot be matched by traditional computing logic. However, there are many technical challenges for quantum computing that must be addressed: such as the extraction of information from the system, materials for physically implementing a quantum computer, and lack of software stack for programming quantum processors. In this session we will discuss the programming environment, quantum compilers, benchmarks/metrics, and materials research for quantum computing.

Session Chair: Dr. Neena Imam, Computing and Computational Sciences Directorate,
Oak Ridge National Laboratory

1:30-2:10 PM: Dr. David Dean, Physical Sciences Directorate, Oak Ridge National Laboratory, *Quantum Computing Materials and Interfaces*

Abstract: This talk will describe ORNL's Laboratory Directed Research and Development Initiative on Quantum Computing Materials and Interfaces. This initiative, which began with the selection of four projects that started in October, 2015, will integrate core competencies in materials, modeling, and isotopes to establish a broad R&D effort in quantum computing. The initiative will create an S&T base at the Laboratory to drive computing beyond the exascale. Our

strategy is to develop tools necessary to characterize and design high-fidelity physical qubits, to explore methods to interface qubits to traditional computers, to develop a multi-qubilt research test bed, to research methods to program multi-qubit systems, and to foster multiagency ties to perform such research in the longer term.

Bio: Dr. David Dean is the Physics Division Director and Isotope Program Director at Oak Ridge National Laboratory. In this role, he oversees DOE Nuclear Physics and High Energy Physics programs at ORNL as well as the ORNL Isotope Program. From 2009 until 2011, Dean was Senior Advisor to the DOE Under Secretary of Science, and from 2007 until 2009 he was the Director of Institutional Planning at ORNL. He is a Fellow of the American Physical Society and the Institute of Physics (London) and is Vice Chair of the Division of Nuclear Physics of the American Physical Society. Dean performs research in the quantum many-body problem and computational physics, with a particular emphasis on coupled-cluster theory and quantum Monte Carlo algorithms. Dean obtained a Ph.D. from Vanderbilt University in 1991 and was a Post-Doctoral Fellow in Physics at Caltech from 1992 until 1995 when he joined ORNL.

2:10-2:50 PM: Mr. Steve Reinhardt, D-WAVE Systems, *Co-Evolving Quantum Architecture with Early Application Feedback*

Abstract: Moore's Law advances in classical computing have slowed recently, creating an opportunity for other types of computing to emerge. Quantum computing is one such candidate, with many implementations being explored. The D-Wave 2X™ quantum annealer, currently with a thousand qubits, is the only commercially available quantum computer; Google has reported differentiated performance results for important common problems. Starting from a description of today's D-Wave hardware and software architecture, we describe some possible evolutionary paths to greater generality. The actual path, and date of widespread usefulness, will depend heavily on the feedback D-Wave receives from early application developers and users over the next few years.

Bio: Steve Reinhardt has led and participated in teams driving significant innovations in high-end computing and analytics, from shared- and distributed-memory programming mechanisms to graph-analytic interfaces for subject-matter experts. He currently works at D-Wave Systems, mapping customer problems to the D-Wave quantum annealer and developing tools to make that easier.

2:50-3:20 PM: Afternoon Break and Refreshments

3:20-4:00 PM: Dr. Mark Novotny, Department Head, Physics and Astronomy Department, Mississippi State University, *Quantum Computers: Testing and Selected Applications*

Abstract: Adiabatic Quantum Computing (AQC) holds the potential to enable solutions to problems that are not in the complexity class P in computational complexity theory. Availability of D-Wave machines with more than 1000 qubits provides a new tool in computational studies. We have tested the D-Wave machines to better understand their limitations, applications, and how the next generation may be enhanced. We describe our tests of the D-Wave, which include random spanning tree studies, planted solutions, and an answer checking procedure. Our main

application is to Boltzmann machines for machine learning. Our enhancements for the next generation include an answer checking paradigm and small-world connections added to the D-Wave Chimera lattice.

Bio: Mark A. Novotny earned both of his degrees in physics, his B.S. degree from North Dakota State University and his Ph.D. degree from Stanford University. His research interests are in computational physics, including algorithm development and applications to materials, devices, and statistical mechanics. He is currently Professor and Head of the Department of Physics and Astronomy, and a William L. Giles Distinguished Professor, at Mississippi State University. He is a Fellow of both the American Physical Society and AAAS.

4:00-4:40 PM: Dr. Travis Humble, Computing and Computational Sciences Directorate, Oak Ridge National Laboratory, *Software Ecosystems for Quantum Computing*

Abstract: Quantum computing promises new opportunities for solving hard computational problems, but harnessing this novelty will require breakthrough concepts in the design, operation, and application of computing systems. In this talk, we define some of the challenges facing the development of quantum computing systems as well as software-based approaches that can be used to overcome these challenges. Following a brief overview of the state of the art, we present recent advances in the modeling and simulation of quantum computing systems, the development of architectures for hybrid high-performance computing systems, and the realization of software stacks for quantum computing. This leads to a discussion of the role that conventional computing plays in the quantum paradigm and how some of the current challenges for exascale computing overlap with those facing quantum computing.

Bio: Dr. Travis Humble is Director of the Quantum Computing Institute at Oak Ridge National Laboratory. His research focuses on the design, development, and benchmarking of new quantum computing platforms. He received his doctorate in theoretical chemistry from the University of Oregon before coming to ORNL in 2005. Dr. Humble is a member of the Center for Engineering Science Advanced Research, the Complex Systems Group, and the Quantum Information Science Group. He is also an associate professor with the Bredesen Center for Interdisciplinary Research and Graduate Education at the University of Tennessee.

4:40-5:20 PM: Dr. David C. McKay, IBM T.J. Watson Research Center, *Universal Quantum Computing at IBM: Successes and Challenges*

Abstract: A universal quantum computer – a computer that harnesses the power of quantum entanglement – promises to revolutionize our ability to solve certain problems, such as factoring, search and molecular design. However, quantum states are fragile, and so constructing a robust quantum computing device is a formidable challenge. At IBM we are pursuing a bottom-up approach by first demonstrating high fidelity gates between a small number of quantum bits ("qubits") as we move towards producing logical qubits that are protected against errors by continuous "quantum parity" measurement. Most recently we demonstrated the ability to detect arbitrary errors using a four qubit configuration [Corcoles et al. Nat. Comm. 6, 6979 (2015)]. In this talk, I will discuss our current progress towards a logical qubit and on novel designs for high fidelity qubit interactions. I will also discuss the challenges that we must address before we can

realize a fully universal quantum computer and the applications for first-generation systems with order of 100 qubits.

Bio: Dr. McKay is a research staff member at IBM in the experimental quantum computing group and primarily studies novel gate architectures for coupling qubits and characterization methods for quantum gates. Dr. McKay received his PhD in physics from the University of Illinois where he studied quantum simulation using ultracold atoms in optical lattices and was a postdoctoral scholar at the University of Chicago studying superconducting quantum computing. David joined IBM in 2015.

5:20 PM: Adjourn

April 6, 2016

7:30-8:30 AM : Breakfast and Registration

8:30-9:30 AM : **Keynote 2**: Dr. Eng Lim Goh, Chief Technology Officer and Senior Vice President, SGI, *Challenges Awaiting Solutions Beyond CMOS*

Bio: Dr. Eng Lim Goh joined SGI in 1989, becoming a chief engineer in 1998 and then chief technology officer in 2000. He oversees technical computing programs with the goal to develop the next generation computer architecture for the new many-core era. His current research interest is in the progression from data intensive computing to analytics, machine learning, artificial specific to general intelligence and autonomous systems. Since joining SGI, he has continued his studies in human perception for user interfaces and virtual and augmented reality.
In 2005, InfoWorld named Dr. Goh one of the World's 25 Most Influential CTOs. He was included twice in the HPCwire list of "People to Watch"; 2005 and 2015. In 2007, he was named "Champions 2.0" of the industry by BioIT World magazine, and received the HPC Community Recognition Award from HPCwire. Dr. Goh is a frequent industry speaker and he continues to discuss, in different forums, innovative technologies and their applications. He co-presented with NASA at the 1st plenary of the Supercomputing 2014 Conference to an audience of 2,500. Before joining SGI, Dr. Goh worked for Intergraph Systems, Schlumberger Wireline and Shell Research. A Shell Cambridge University Scholar, Dr. Goh completed his Ph.D. research and dissertation on parallel architectures and computer graphics, and holds a first-class honors degree in mechanical engineering from Birmingham University in the U.K. Dr. Goh has been granted four U.S. patents, two of which as the inventor and the others as co-inventor.

Abstract: In addition to classes of large Non-deterministic Polynomial problems, there are also challenges with current compute-, data-intensive and aspiring artificially-intelligent systems that may only have practical solutions Beyond CMOS. We discuss their examples.

Session 3: Superconducting Computing

Superconducting computing may be an attractive alternative to CMOS circuits with many potential advantages such as ultrafast switching time and very low energy consumption. This session will discuss the recent investments in superconducting computing by federal agencies,

the state of the art in superconducting circuits, the very difficult task of suitable cryogenic memory design, and fabrication challenges.

Session Chair: Mr. Adam Wurtz, United States Department of Defense

9:30-10:15 AM: Dr. Scott Holmes, IARPA, *Superconducting Computing and the IARPA C3 Program*

Abstract: Superconducting computing is a potential solution for energy-efficient, large-scale computing. The Intelligence Advanced Research Projects Activity (IARPA) Cryogenic Computing Complexity (C3) program was established to demonstrate a complete superconducting computer including processing units and cryogenic memory. IARPA expects that the C3 program will be a five-year two-phase program. Phase one, which is ongoing and encompasses the first three years, primarily serves to develop the technologies that are required to separately demonstrate a small superconducting processor and memory units. Phase two, which is expected to last an additional two years, will integrate those new technologies into a small-scale working model of a superconducting computer. Program goals and metrics are presented along with the main technical challenges and approaches to overcome them.

Bio: Dr. Scott Holmes is a technical consultant with Booz Allen Hamilton. His interests include superconducting electronics, future computing, and learning systems. He holds a joint Ph.D. in Materials Science and Nuclear Engineering & Engineering Physics from the University of Wisconsin, Madison. He is a Senior Member of the IEEE and a member of the IEEE Council on Superconductivity. Currently he supports the IARPA Cryogenic Computing Complexity (C3) program to develop energy-efficient superconducting computing technologies.

10:15-11:00 AM: Dr. Mark Gouker, MIT Lincoln Laboratory, *SFQ Circuits Fabrication and EDA Status*

Abstract: This talk will provide an overview of the status of single flux quantum (SFQ) superconducting circuit fabrication and electronic design automation (EDA) tools. It will include a summary of SFQ fabrication processes around the world, and it will compare and contrast SFQ fabrication with CMOS fabrication. We will describe the IARPA roadmap for SFQ fabrication that is being performed within the Cryogenic Computation Complexity (C3) Program being performed at MIT Lincoln Laboratory.

Bio: Dr. Mark Gouker leads the Quantum Information and Integrated Nanosystems Group at MIT Lincoln Laboratory. This group has robust activities in superconducting circuit fabrication, 90nm CMOS technology for specialized application, and integrated photonics. The group also performs work in quantum information processing with superconducting and trapped-ion qubit modalities and quantum sensing with NV-centers in diamond.

11:00-11:30 AM: Break

11:30 AM-12:15 PM: Dr. Kenneth Zick, Northrop Grumman Corporation, *Superconducting Supercomputing with Reciprocal Quantum Logic: A View from the C3 Program*

Abstract: We will review recent progress in creating digital superconducting circuits using Reciprocal Quantum Logic in the IARPA Cryogenic Computing Complexity (C3) program.

Bio: Kenneth Zick is a Systems Architect at Northrop Grumman Corp., focused on future computer architectures and harnessing emerging physical effects. His background includes 18 years in chip design and adaptive systems-on-a-chip research. He received a Ph.D. in Computer Science and Engineering from the University of Michigan.

12:15-1:15 PM: Lunch and Networking

Session 4: Emerging Processor and Memory Technology

The fourth focus area of this workshop will be emerging processor and memory technologies that can overcome the Von Neumann bottleneck for enhanced data analytics and storage support. The workshop will also focus on technologies such as Memristors, nanocrystal based HPC storage, and 3D processor and memory architectures.

Session Chair: Ms. Janice Elliott, United States Department of Defense

1:15-2:00 PM: Mr. David Mountain, United States Department of Defense, *Neuromorphic Computing*

Abstract: With the end of Moore's Law in sight, various technologies and models of computing are being actively pursued as potential replacements. The NSCI is explicitly including beyond CMOS technologies, and the recent OSTP grand challenge presumes neuromorphic computing is an important approach. This talk provides an overview of this emerging model of computing, highlighting existing initiatives, recent accomplishments, and current challenges.

Bio: David J. Mountain is the Senior Technical Director at the Laboratory for Physical Sciences at Research Park, a Department of Defense research lab in Catonsville, MD. The LPS-RP mission is to collaborate with industry, academia, and the government to drive innovative research that will improve advanced computing systems for a range of mission applications including cybersecurity, cryptanalysis, and complex data analytics. His responsibilities include research activities in neuromorphic computing, advanced modeling and simulation, energy efficiency, productivity, and resilience. His personal research projects have included radiation effects studies, hot carrier reliability characterization, and chip-on-flex process development utilizing ultra-thin circuits. He has been actively involved with 3D electronics research for over two decades. Mr. Mountain is the author of more than two dozen technical papers, has been awarded eight patents, and is a Senior Member of the IEEE.

2:00-2:45 PM: Dr. Rao Kosaraju, National Science Foundation, *An Overview of the Research Supported by the Division of Computing and Communication Foundations (CCF) at NSF*

Abstract: The talk first gives an overview of the research supported by the division of Computing and Communication Foundations. Then it will outline the existing and the new research initiatives to overcome the Moore's Law bottleneck.

Bio: Dr. Kosaraju is currently the Division Director of the Division of Computing and Communication Foundations (CCF) in the Computer and Information Science and Engineering (CISE) Directorate at the National Science Foundation. He is at NSF on assignment from the Johns Hopkins University where he holds the Edward J. Schaefer Chair in Engineering in the Department of Computer Science. He has broad research interests in Algorithms. He served on the editorial boards of many journals; he was the managing editor of the SIAM J. on Computing during 1980-1988. He has also served on the advisory and external committees of many computer science departments. He was recognized for teaching excellence several times. He is a fellow of ACM and IEEE.

2:45-3:15 PM: Afternoon Break and Refreshments

3:15-4:00 PM: Mr. Terry Leslie, Director, Research & Business Development-Automata Computing, Micron Technology, *Micron's Automata Processor*

Abstract: The Automata Processor (AP) is an upcoming co-processor which can be programmed to compute a large set of user-defined Nondeterministic Finite Automata (NFA) in parallel against a single query datastream. Defining an NFA to be programmed into the processor is simple, akin to creating classical state diagrams. Once the automata are compiled and programmed into the processor, a query stream can be streamed to the processor. Internally, each byte in the data-stream is broadcast to all the processing elements in the chip such that all paths in an NFA, and all NFAs in the chip can be examined in parallel. Any NFA matched during the streaming is reported along with the offset in the query stream where the match occurred. This provides a greatly simplified programming and execution interface without the need to handle communication delays, race conditions, etc. typical of other contemporary reconfigurable processors. The AP-board is designed to combine the processing capability of 32 AP-chips into a single PCIe-based accelerator board. Cumulatively, it provides the programmable resources to emulate a total of over 1.5 million edge transitions per board. Multiple chips on an AP-board can be combined to function as a single logical core, each processing a single data-stream at 1 Gbps. The AP-board can currently process up to 8 streams in parallel. This number is expected to double in the next generation of hardware. When the processor was publicly unveiled in 2013, it was six years into its making. Today, "engineering samples" have been showcased at the leading conferences running demo-applications based on our previous work. A limited set of "pre-production" AP-boards are being supplied to select research labs, universities, and industrial collaborators now with general engineering availability scheduled for June and full scale manufacturing in the second half of 2016. Also, researchers can obtain and use the AP-SDK for design validation and run-time simulation. A Center for Automata Processing (CAP) has also been set up at the University of Virginia to allow members to access a computing cluster containing AP-boards. Owing to the architecture and the working of the Automata Processor, identifying features in query strings (defined as finite state machines) follows naturally. However, researchers have now developed a variety of core algorithmic techniques which have expanded the scope of usage of the processor significantly. These include the use of regular-

expression engines, de-novo motif searches for biological sequences, numerical analysis, machine learning, association rule mining and the development of algorithmic techniques for large scale graph analytics. This talk will briefly discuss the basics of the AP architecture, what algorithms are well suited for the AP and explore some examples of AP algorithm research and implementations.

Bio: Terrence C. Leslie is the Director of Business Development for the Automata Processing team at Micron Technology, Inc. He is responsible for the academic and government segments where he drives the development of university research programs and Government business opportunities. Mr. Leslie is also a Distinguished Member of the Micron Technical Staff. He previously held process, test and quality engineering management positions at Micron in Manassas, VA. Before joining Micron, Mr. Leslie held various positions in process and quality engineering and management at IBM Microelectronics and Dominion Semiconductor, including process engineering manager and vice president of quality at Dominion. Mr. Leslie holds a BS in electrical engineering from the University of Illinois, an MS in electrical engineering from the University of Vermont, and an MS in management of technology from the Massachusetts Institute of Technology (MIT).

4:00-4:45 PM: Dr. Samira Khan, University of Virginia*, Rethinking Memory and Storage for Future Computing Systems*

Abstract: For future HPC systems, we need to rethink and redesign our computing model focusing on minimizing data movement and storage. In this talk I will present three high-level directions to solve the memory challenges in the future systems focusing on across the stack solutions. First, I will talk about how we can continue DRAM scaling, yet improve reliability, latency, and cost by rethinking the interface between circuits and the system. Second, I will discuss how we can unify the conventional memory and storage hierarchy leveraging the emerging non-volatile memory technologies. Third, I will present my vision on how we need to fundamentally rethink our computing model and redesign processing as near data computation in every level of the system

Bio: Dr. Samira Khan is an Assistant Professor at the University of Virginia (UVa). Prior to joining UVa, she spent three years as a Post-Doctoral Researcher at Carnegie Mellon University, funded by Intel Labs. Her research focuses on improving the performance, efficiency, and reliability of the memory system. She received her PhD from the University of Texas at San Antonio. During her graduate studies, she worked at Intel, AMD, and EPFL.

4:45 PM: Adjourn

Panel Discussion: 6:15 PM-8 PM

The workshop will conclude with a panel discussion to address questions such as (1) what technology is “most likely to succeed” as the next big thing in HPC, (2) how long before it could be in a system doing something useful in HPC, and (3) key issues that need to be resolved for that technology.

Moderator: Mr. Steve Pritchard, United States Department of Defense
Panelists: TBD