

NVIDIA GPU Computing Theater

The theater will feature talks given by experts on a wide range of topics on high performance computing. Open to all attendees of SC10, the theater is located in the NVIDIA booth (#1343) and will feature scientists, developers, and industry luminaries giving technical talks on the latest research in the HPC field.

Tuesday, November 16th

Time	Topic / Presenter / Affiliation
10:00-10:30	High Performance Molecular Simulation, Visualization, and Analysis on GPUs John Stone, University of Illinois at Urbana-Champaign
12:00-12:30	Large-Scale GPU Computing of Multi-Phase Flow Wei Ge, Institute of Process Engineering, Chinese Academy of Sciences
1:00-1:30	Supercomputing for the Masses: Killer-Apps, Parallel Mappings, Scalability and Application Lifespan Robert Farber, Pacific Northwest National Laboratory
2:00-2:30	Keeneland – An NSF Heterogeneous Computing Resource for Open Computational Science Jeffrey Vetter, Oakridge National Laboratory/Georgia Tech
4:00-4:30	Large Scale Distributed GPU Isosurfacing Paul Navratil, Texas Advanced Computing Center

Wednesday, November 17th

Time	Topic / Presenter / Affiliation
10:00-10:30	GPU Cloud Computing: Examples in 3D City Planning using RealityServer® Paul Arden, mental images and Tamrat Belayneh, ESRI
12:00-12:30	TSUBAME 2.0 Satoshi Matsuoka, Global Scientific Information and Computing Center (GSIC) of Tokyo Institute of Technology (Tokyo Tech)

1:00-1:30	PGI CUDA Compiler for x86 Michael Wolfe, The Portland Group
2:00-2:30	Large-Scale CFD Applications and a Full GPU Implementation of a Weather Prediction Code on the TSUBAME Supercomputer Takayuki Aoki, Tokyo Institute of Technology
3:00-3:30	A Linear Algebra Library for Multicore/Accelerators: the PLASMA/MAGMA Collection Jack Dongarra, University of Tennessee

Thursday, November 18th

Time	Topic / Presenter / Affiliation
11:00-11:30	The Future of GPU Computing Bill Dally, NVIDIA
12:00-12:30	Using GPUs to Run Next-Generation Weather Models Mark Govett, National Oceanic and Atmospheric Administration
1:00-1:30	Data Analysis and Visualization on Emerging Architectures Patrick McCormick, Los Alamos National Laboratory

Speaker Biographies

Takayuki Aoki

Professor Aoki is a deputy director of Global Scientific Information and Computing center (GSIC), Tokyo Institute of Technology. His research area is Computational Fluid Dynamics (CFD) and he has developing high-accurate numerical schemes. He is a leader of GPU computing for large-scale CFD and organizes GPU Computing Study Group with 600 members and GPU sessions in many conferences. Last year, he published a CUDA programming textbook in Japanese.

Bill Dally

Bill Dally joined NVIDIA in January 2009 as chief scientist, after spending 12 years at Stanford University, where he was chairman of the computer science department. Dally and his Stanford team developed the system architecture, network architecture, signaling, routing and synchronization technology that is found in most large parallel computers today.

Dally was previously at the Massachusetts Institute of Technology from 1986 to 1997, where he and his team built the J-Machine and the M-Machine, experimental parallel computer systems that pioneered the separation of mechanism from programming models and demonstrated very low overhead synchronization and communication mechanisms. From 1983 to 1986, he was at California Institute of Technology (CalTech), where he designed the MOSSIM Simulation Engine and the Torus Routing chip, which pioneered “wormhole” routing and virtual-channel flow control.

Dally is a cofounder of Velio Communications and Stream Processors. He is a member of the National Academy of Engineering, a Fellow of the American Academy of Arts & Sciences, a Fellow of the IEEE and the ACM, and has received the IEEE Seymour Cray Award and the ACM Maurice Wilkes award. He has published over 200 papers, holds over 50 issued patents, and is an author of two textbooks.

He received a bachelor’s degree in Electrical Engineering from Virginia Tech, a master’s in Electrical Engineering from Stanford University and a PhD in Computer Science from CalTech.

Jack Dongarra

Jack Dongarra received a Bachelor of Science in Mathematics from Chicago State University in 1972 and a Master of Science in Computer Science from the Illinois Institute of Technology in 1973. He received his Ph.D. in Applied Mathematics from the University of New Mexico in 1980. He worked at the Argonne National Laboratory until 1989, becoming a senior scientist. He now holds an appointment as University Distinguished Professor of Computer Science in the Computer Science Department at the University of Tennessee and holds the title of Distinguished Research

Staff in the Computer Science and Mathematics Division at Oak Ridge National Laboratory (ORNL), Turing Fellow at Manchester University, and an Adjunct Professor in the Computer Science Department at Rice University. He is the director of the Innovative Computing Laboratory at the University of Tennessee. He is also the director of the Center for Information Technology Research at the University of Tennessee which coordinates and facilitates IT research efforts at the University.

He specializes in numerical algorithms in linear algebra, parallel computing, the use of advanced-computer architectures, programming methodology, and tools for parallel computers. His research includes the development, testing and documentation of high quality mathematical software. He has contributed to the design and implementation of the following open source software packages and systems: EISPACK, LINPACK, the BLAS, LAPACK, ScaLAPACK, Netlib, PVM, MPI, NetSolve, Top500, ATLAS, and PAPI. He has published approximately 200 articles, papers, reports and technical memoranda and he is coauthor of several books. He was awarded the IEEE Sid Fernbach Award in 2004 for his contributions in the application of high performance computers using innovative approaches and in 2008 he was the recipient of the first IEEE Medal of Excellence in Scalable Computing. He is a Fellow of the AAAS, ACM, IEEE, and SIAM and a member of the National Academy of Engineering.

Robert Farber

Rob Farber has worked with massively parallel computers and algorithms since the early 1980s as a scientist at prestigious institutions such as Los Alamos National Laboratory, NERSC and PNNL, as a consultant, and as a co-founder of two computation-based companies that achieved liquidity events. Recently, Rob has been teaching people how to think about and program in CUDA through his article series on the Doctor Dobbs Journal site as well as Scientific Computing and other venues.

Wu-chun Feng

Wu-chun Feng - or simply "Wu" - is an Associate Professor of Computer Science and Electrical & Computer Engineering at Virginia Tech (VT) as well as an Adjunct Associate Professor of Cancer Biology and the Translational Science Institute at Wake Forest University (WFU). He also directs the Synergy Laboratory at VT and several major projects including mpiBLAST (<http://www.mpiblast.org/>), Supercomputing in Small Spaces (<http://sss.cs.vt.edu/>), and The Green500 (<http://www.green500.org/>). Previous professional stints include The Ohio State University, Purdue University, University of Illinois at Urbana-Champaign, Orion Multisystems, Vosaic, IBM T.J. Watson Research Center, NASA Ames Research Center, and most recently, Los Alamos National Laboratory.

He received a B.S. in Electrical & Computer Engineering and in Music (Honors) in 1988 and an M.S. in Computer Engineering from Penn State University in 1990. He earned a Ph.D. in Computer Science from the University of Illinois at Urbana-Champaign in 1996.

Wei Ge

Professor of Chemical Engineering and Simulation at Institute of Process Engineering, Chinese Academy of Sciences. Born 1970, B.Sc and Ph.D of Harbin Institute of Technology, China.

Mark Govett

Mark manages NOAA Earth System Research Laboratory's Advanced Computing Section, a software group that supports weather model development, code parallelization, and exploring advanced computing technologies including GPUs. He has a background in high performance computing, code parallelization and compiler development. Recently, he wrote a Fortran to CUDA compiler to parallelize and run a next generation weather model on GPUs.

Satoshi Matsuoka

Satoshi Matsuoka is a full Professor at the Global Scientific Information and Computing Center (GSIC) of Tokyo Institute of Technology (Tokyo Tech). He is the technical leader in the construction of the TSUBAME supercomputer, which became the fastest supercomputer in Asia-Pacific in June, 2006, and continued to be Japan's #1 for 4 consecutive Top 500s during June 2006 - May 2008. He has also pioneered Grid computing research in Japan, and in particular co-lead the Japanese \$100 million NAREGI project, whose aim was to develop and foster research grid infrastructure in Japan. He has chaired many ACM/IEEE international conferences, and will be the technical papers chair for SC09. He has won many awards including the JSPS Prize from the Japan Society for Promotion of Science in 2006, from his Majesty Prince Akishinomiya, for the first time as a Computer Scientist.

Patrick McCormick

Patrick McCormick is a project and team leader in the Computer, Computational, and Statistical Sciences Division at Los Alamos National Laboratory. During his 16 year career at Los Alamos his research interests have spanned methods for parallel rendering and the visualization of extremely large scientific data sets. He is currently exploring the development of programming models and languages for data analysis and visualization on emerging processor architectures.

Paul Navratil

Paul is a Visualization Scientist in the Data and Information Analysis division of the Texas Advanced Computing Center (TACC) at the University of Texas at Austin. His research interests include efficient algorithms for large-scale parallel visualization and data analysis (VDA) and innovative design for large-scale VDA systems. His recent work includes algorithms for ultrascale distributed-memory ray tracing, work that enables photo-realistic rendering of the largest datasets produced on supercomputers today, such as cosmologic simulations of the Universe and computational fluid dynamics (CFD) simulations at unprecedented levels of detail. Paul also helps manage TACC's visualization systems, including Longhorn, the largest supercomputer in the world dedicated to VDA, and Stallion, the highest-resolution tiled display in the world.

John Stone

John Stone is a Senior Research Programmer in the Theoretical and Computational Biophysics Group at the Beckman Institute for Advanced Science and Technology, and Associate Director of the NVIDIA CUDA Center of Excellence at the University of Illinois. Mr. Stone is the lead developer of VMD, a high performance molecular visualization tool used by researchers all over the world.

His research interests include molecular visualization, GPU computing, parallel processing, ray tracing, haptics, and virtual environments. Mr. Stone also provides consulting services for projects involving computer graphics, GPU computing, and high performance computing in general. Prior to joining University of Illinois in 1998, Mr. Stone helped develop the award winning MPEG Power Professional line of video compression tools at Heuris.

Jeffrey Vetter

Jeffrey Vetter is a computer scientist in the Computer Science and Mathematics Division (CSM) of Oak Ridge National Laboratory (ORNL) where he leads the Future Technologies Group. His research interests are largely in the areas of experimental software systems and architectures for high-end computing. Vetter earned his Ph.D. in computer science from Georgia Tech and joined ORNL in 2003.