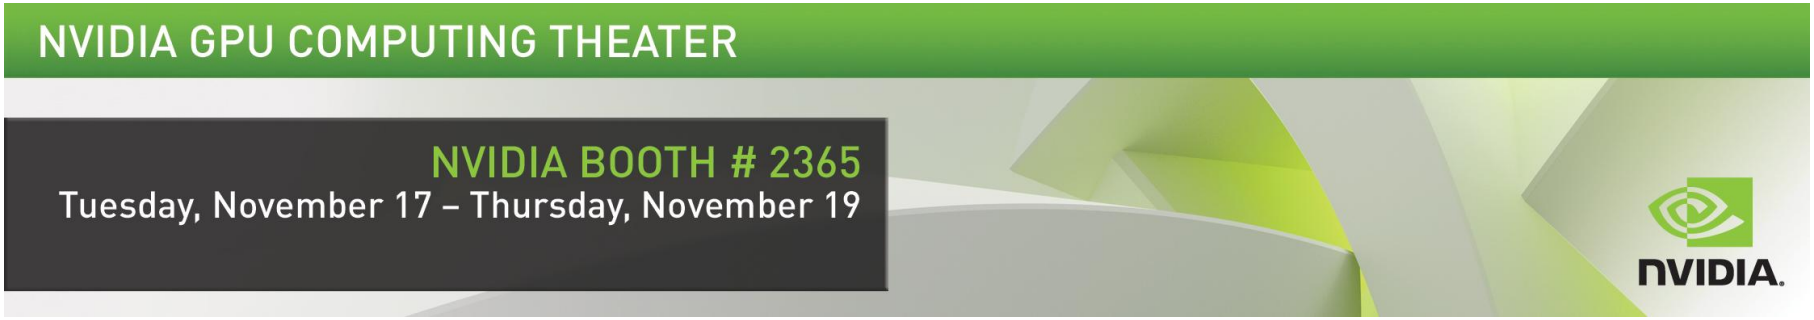

The theater will feature talks given by experts on a wide range of topics on high performance computing. Open to all attendees, the theater is located in the NVIDIA booth (#2365) and will feature scientists, developers, and industry luminaries.

Tuesday, November 17th

Time	Topic / Presenter / Affiliation
11:00-11:30	**Accelerating Market Value-at-Risk Estimation on GPUs** Matthew Dixon, University of California Davis
12:00-12:30	Tobias Brandvik, University of Cambridge
1:00-1:30	**GPU Acceleration: a Fad or the Yellow Brick Road onto Exascale?** Professor Satoshi Matsuoka, Global Scientific Information and Computing Center (GSIC) of Tokyo Institute of Technology (Tokyo Tech),
2:00-2:30	**A Powerful IDE for GPU Computing on Windows: Code Named Nexus**
3:00-3:30	**Languages, APIs and Tools for GPU Computing on the CUDA Architecture**
4:00-4:30	**Overview of NVIDIA CUDA Support in AMBER - Lessons Learned, Capabilities Gained** Ross Walker, Research Professor, University of California San Diego, San Diego Supercomputer Center
5:00-5:30	**The Next Generation CUDA Architecture, Code Named Fermi** **The Soul of a Supercomputer in the Body of a GPU**

NVIDIA GPU COMPUTING THEATER

NVIDIA BOOTH # 2365

Tuesday, November 17 – Thursday, November 19

Wednesday, November 18th

Time	Topic / Presenter / Affiliation
10:30-11:00	**A Linear Algebra Library for Multicore/Accelerators: the PLASMA/MAGMA Collection** Jack Dongarra, University of Tennessee
11:00-11:30	**Keeneland - An NSF Funded Partnership To Enable Large Scale Computational Science with Heterogeneous Computing** Jeffrey Vetter, Oakridge National Labs/Georgia Tech
12:00-12:30	**Harnessing GPU Speed to Accelerate LAMMPS Particle Simulations** Paul Crozier, Sandia National Lab
1:00-1:30	**Accelerating Molecular Modeling Applications with GPU Computing** John Stone, Beckman Institute for Advanced Science and Technology, and Associate Director of the NVIDIA CUDA Center of Excellence at the University of Illinois
1:30-2:00	**A Powerful IDE for GPU Computing on Windows: Codenamed Nexus**
2:00-2:30	**The Future of GPU Computing** Bill Dally, NVIDIA
3:00-3:30	**Application of GPUs to Solve the Systems of Linear Equations That Arise in Lattice Quantum Chromodynamics (lattice QCD)** Mike Clark, Harvard University
4:00-4:30	**Data Analysis and Visualization on Emerging Architectures** Patrick McCormick, Los Alamos National Lab
5:00-5:30	**Languages, APIs and Tools for GPU Computing on the CUDA Architecture**

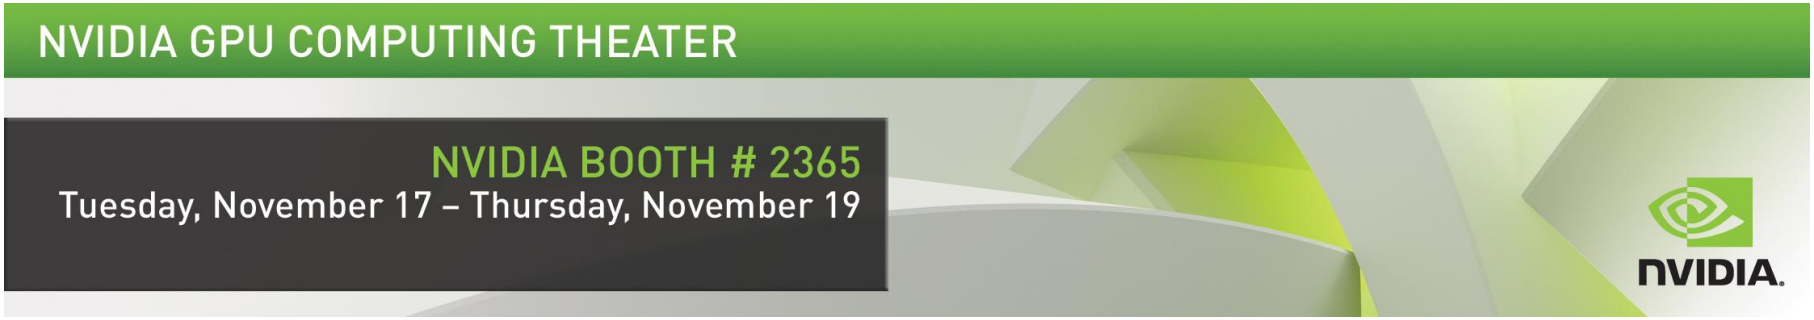

Thursday, November 19th

Time	Topic / Presenter / Affiliation
11:00-11:30	**Low Power Data-Intensive Computing Using the ION Platform** Alex Szalay, Johns Hopkins University
12:00-12:30	**A Powerful IDE for GPU Computing on Windows: Code Named Nexus**
1:00-1:30	**The Next Generation CUDA Architecture, Code Named Fermi** **The Soul of a Supercomputer in the Body of a GPU**
2:00-2:30	Wu-chun Feng, Virginia Tech **To GPU Synchronize or Not GPU Synchronize**
3:00-3:30	**Languages, APIs and Tools for GPU Computing on the CUDA Architecture**

Speaker Biographies

Michael Clark

Michael Clark received his PhD in theoretical particle physics in 2006 from The University of Edinburgh. During his PhD research, he spent 18 months working on the QCDOC supercomputer project, a precursor to BlueGene, at Columbia University. From 2005-2009 he performed postdoctoral research at Boston University, where his research focused upon multi-scale algorithms for computational particle physics. At this time he became interested in heterogeneous computation, this prompting his move to the Initiative in Innovative Computing at Harvard University where his current research is focused upon algorithms for various GPU computing applications.

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.

Matthew Dixon

Matthew Dixon is a postdoctoral fellow in the computer science department at UC Davis. In 2007 he graduated with a PhD in applied mathematics from Imperial College and in 2002, a MSc in parallel and scientific computation (with distinction) from Reading University, UK. Matthew has previously held a postdoctoral appointment with the Institute for Computational and Mathematical Engineering at Stanford University and worked and consulted for investment banks in Europe as a quantitative risk analyst. He has spoken at several computational finance conferences and workshops in Europe and is a member of the organizational committee for the first GPUs in finance workshop for quantitative analysts in London.

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.

Wu-chun Feng

Wu-chun Feng (a.k.a. "Wu") is an associate professor in the Department of Computer Science and Department of Electrical & Computer Engineering at Virginia Tech, where he also directs the Synergy Laboratory. Prior to joining Virginia Tech in 2006, Feng spent 7 years at Los Alamos National Laboratory, where he and his team created three artifacts for high-performance computing, each of which secured a R&D 100 Award: (1) Green Destiny: Super-Efficient Supercomputing, (2) Optimized 10-Gigabit Ethernet Adapter, and (3) mpiBLAST: A High-Speed Software Catalyst for Genetic Sequence Alignment. 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, and NASA Ames Research Center.

His research interests encompass synergistic aspects of high-performance computing from systems software and middleware to applications software, including green supercomputing, accelerator-based computing with GPUs, and bioinformatics. He has more than 150 peer-reviewed technical publications, and his work has been featured in media outlets such as The New York Times, CNN, and BBC News. He received a B.S. in Electrical & Computer Engineering and in Music (Honors) in 1988 and an M.S. in Computer Engineering from the Pennsylvania State University in 1990. He earned a Ph.D. in Computer Science from the University of Illinois at Urbana-Champaign in 1996. He is a senior member of the IEEE and was listed on HPCwire's Top People to Watch List in 2004.

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 forSC09. 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.

Alex Szalay

Alexander Szalay is the Alumni Centennial Professor of Astronomy at the Johns Hopkins University. He is also Professor in the Department of Computer Science. He is a cosmologist, working on the statistical measures of the spatial distribution of galaxies and galaxy formation. He was born and educated in Hungary. After graduation he spent postdoctoral periods at UC Berkeley and the University of Chicago, before accepting a faculty position at Johns Hopkins. He is the architect for the Science Archive of the Sloan Digital Sky Survey. He has been collaborating with Jim Gray of Microsoft to design an efficient system to perform data mining on the SDSS Terabyte sized archive, based on innovative spatial indexing techniques. He is leading a grass-roots standardization effort to bring the next generation Terabyte-sized databases in astronomy to a common basis, so that they will be interoperable – the Virtual Observatory. He is Project Director of the NSF-funded National Virtual Observatory. He is involved in creating testbed applications for the Computational Grid. He has written over 340 papers in various scientific journals, covering areas from theoretical cosmology to observational astronomy, spatial statistics and computer science. In 1990 he has been elected to the Hungarian Academy of Sciences as a Corresponding Member. In 2003 he was elected as a Fellow of the American Academy of Arts and Sciences. In 2004 he received one of the Alexander Von Humboldt Prizes in Physical Sciences.

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.

Ross Walker

Dr. Ross Walker is a research professor working at the University of California San Diego, San Diego Supercomputer Center interested in the fields of computational chemistry and molecular biology, especially in the development of efficient algorithms for parallel computation of combined QM/MM Molecular Dynamics. Ross currently works as part of the AMBER development team integrating new features into AMBER 10, including the GPU acceleration, as well as training and consultancy. He is a member of the National Science Foundation Resource Allocations Review panel, and regularly peer reviews a number of U.S. and International Journals.