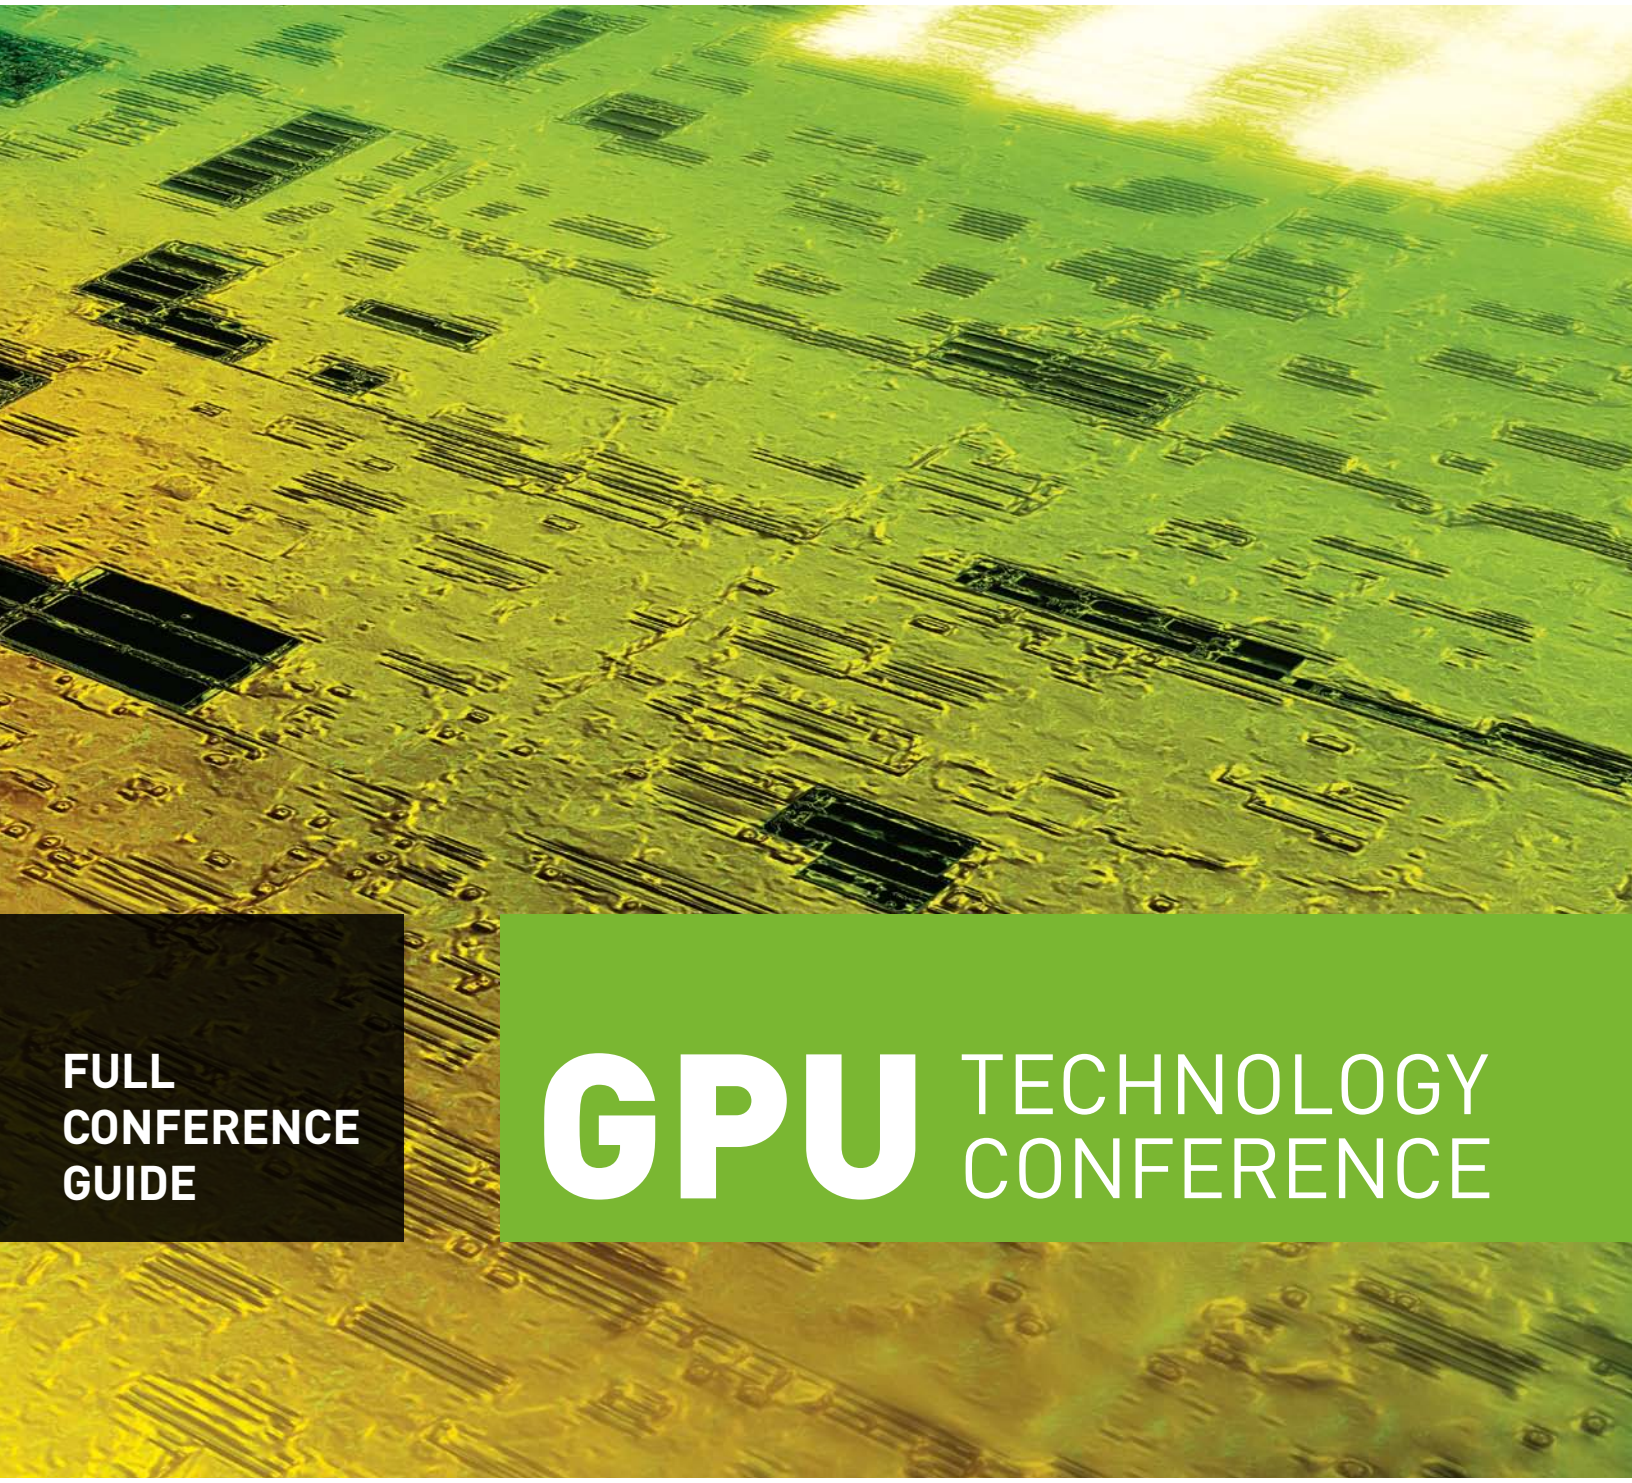

FULL CONFERENCE GUIDE

GPU TECHNOLOGY CONFERENCE

PRESENTED BY

NVIDIA

SPONSORED BY

Cooley GODWARD KRONISH LLP

SAMSUNG

ASUS

PNY

DELL

CHRISTIE

DELL PRECISION RACK –
A PERFECT FIT FOR YOUR DATA CENTER

Developed in close collaboration with hardware and software partners, the Dell Precision rack workstation delivers high performance in a flexible 2U chassis — an ideal solution for centralizing critical data and workstation assets in secure data center locations. To help maximize flexibility, the Dell Precision rack workstation is enhanced by a choice of industry-standard slot combinations. Designed for performance, reliability, and scalability in environments where space is at a premium, the Dell Precision rack workstation is ideal for high-performance clusters and rendering farms. For remote, crowded, heat-challenged, and acoustically sensitive environments such as financial trading or factory floors, Dell's optional remote access solution is the ideal solution for workstation users who need to interact with their system when elsewhere on the network.

"WITH THE EMERGING GPU OPPORTUNITIES, DELL PRECISION SHOULD BE THE PLATFORM OF CHOICE FOR ANYONE LOOKING TO ACCELERATE THEIR CGI WORKFLOW."

— Keith Jeffery, Head of Business development, Taylor James

Dell Precision workstations help meet your application demands for:

RELIABILITY:

- Dell Precision, the world's most preferred workstation family
- Compatibility for 90 applications based on certifications by 35 leading Independent Software Vendors (ISVs)

FLEXIBILITY:

- Broad range of form factors: Desktop, tower, mobile, and rack to meet a range of applications
- Over 1 billion ways of configuring select systems
- Highly portable 14" mobile workstation systems to full desktop replacements with 17" edge-to-edge display options
- Remote access solutions
- Flexible 2U rack chassis option

PERFORMANCE AND SCALABILITY:

- Intel®-based processing power options including Intel® Core™2 Extreme and Intel® Xeon® processors on select systems
- High performance and vastly scalable memory options
- Go-anywhere, workstation class performance
- Multiple processing cores

SERVICE OPTIONS:

- Dell Custom Factory Integration (CFI)
- Choose from a range of Dell service and support options

DELL PRECISION: our highest performance and most scalable business systems. the undisputed world leader in workstations for the last ten consecutive years.[1]

1. Source: IDC, Worldwide Quarterly Workstation Tracker, August 2009 (unit shipments)
2. Significant system memory may be used to support graphics, depending on system memory size and other factors

FOR MORE INFORMATION, VISIT DELL.COM/Precision

AD# G810003923 STOCK# 5076

Welcome to the inaugural GPU Technology Conference!

The era of GPU Computing is upon us. Seismic shifts are taking place in computing, triggered by the GPU and its rich possibilities.

The GPU Technology Conference is the single best place to learn about these changes, see the technologies and applications causing this disruption, and get the tools to solve the world's most complex computing challenges.

Among the attendees are officials from 150 of the world's premier universities, as well as 60 innovative start-ups, and top enterprises in the fields of energy exploration, finance, medical and life sciences, film and broadcast, and more. The caliber of attendees include chief scientists, chief technology officers, senior-level developers, engineers and researchers from 36 countries.

We invite you to take advantage of this unique opportunity to network with and learn from each other by attending the programs below:

Emerging Companies Summit

The second annual Emerging Companies Summit is a unique forum with presentations by 60 start-up companies from a dozen countries, discussing how they utilize GPUs to deliver innovative solutions in areas ranging from cloud computing and computer vision to gaming and energy exploration. Additional highlights include:

- Panel – Future Directions in GPU Computing
- Panel – Raising Capital for Emerging Companies in a Challenging Environment
- Fireside Chat with Jen-Hsun Huang, moderated by prominent industry analyst Jon Peddie

ACCESS: All GTC pass holders are invited to attend the Emerging Companies Summit and join the discussions that are reshaping the computing industry.

GPU Developers Summit

Designed to help developers of consumer, professional and HPC applications to harness the massively parallel processing power of the GPU, the GPU Developers Summit will feature experts from a broad range of industries, who will share insights and updates on state-of-the-art techniques in GPU Computing, media processing, advanced visualization and related areas.

ACCESS: GPU Developers Summit sessions are open to Full Conference, Research Summit and Press Pass Holders.

NVIDIA Research Summit

Designed for Academics, this cross-disciplinary forum targeting researchers interested in using GPUs in science and engineering. Attendees new to GPU computing will learn how GPU computing can drastically increase computational power and dramatically reduce time-to-discovery; attendees already using GPU computing can showcase their work, network with each other, learn advanced topics in GPU computing, and discuss their work with NVIDIA engineers and researchers.

ACCESS: GPU Developers Summit sessions are open to Full Conference, Research Summit and Press Pass Holders.

Quadro Professional Graphics

NVIDIA® Quadro® by PNY

Professional Solutions. Professional Support.

Visit the PNY booth and see how NVIDIA Quadro by PNY and Quadro Plex multi-GPU solutions are used today to enhance real-world professional graphics and HPC applications. Meet with our experienced product managers and partners to discuss your development and application needs. Learn how NVIDIA Quadro by PNY professional solutions enable new technical and business possibilities.

PNY Quadro® Advantage:

- 3 year warranty upon product registration
- Pre-sales support and system configuration assistance
- Dedicated Quadro Field Application engineers
- Direct tech support hot lines
- Certified software support and bug reporting
- Published product support and training materials
- Support for all workstation brands and complex installations
- Long product life cycles and availability

Ideal for professional applications:

- CAD/CAE
- Scientific
- Large Scale Visualization
- Oil and Gas Exploration
- Automotive Design
- Digital Content Creation
- HD Broadcast & Film

Find out more at: www.pny.com/quadro

NVIDIA QUADRO

PNY®

IMPORTANT INFORMATION

If there is anything else we can do to make your conference experience better, please stop by the info desk and let us know!

REGISTRATION / INFORMATION DESK HOURS

WEDNESDAY SEPTEMBER 30	8:00 AM to 6:30 PM
THURSDAY OCTOBER 1	8:00 AM to 6:30 PM
FRIDAY OCTOBER 2	8:00 AM to 1:30 PM

EXHIBIT HALL AND MEAL HOURS

WEDNESDAY SEPTEMBER 30	6:00 PM to 8:00 PM	Exhibits Open / Networking Happy Hour
THURSDAY OCTOBER 1	12:00 PM to 2:00 PM	Exhibits Open / Networking Lunch
	5:30 PM to 8:00 PM	Exhibits Open / Networking Happy Hour
FRIDAY OCTOBER 2	11:30 AM to 1:00 PM	Exhibits Open / Networking Lunch

BUSINESS CENTER

The Fairmont Hotel has a self-service business center that is open 24 hours, and staffed between 7:00 AM and 10:00 PM. It is located on the same floor as the GPU Technology Conference, near the Regency Ballroom.

WIRELESS INTERNET ACCESS

Wireless internet access should be available in most areas in the GPU Technology Conference meeting rooms under GPUTechConf.

FOLLOW US ON THE WEB

There is a lot going on at the conference, and we know how you can still miss the good stuff no matter how hard you try. No worries – log on to www.NVIDIA.com/gtc as we have converted the landing page of the GPU Technology Conference to give you play-by-play coverage of the event, including interviews, photos and videos. You can also follow us on Twitter at @nvidia_news or @nvidiadeveloper.

SPECIAL OPENING NIGHT CHARITY EVENT

Join us for a party and turn your drinks into donations at our Opening Night Charity Social. Hang out in the heart of downtown San Jose, let loose to some great music, and mingle with fellow GTC attendees. The local clubs we visit will donate a percentage of food and drink revenue to fund a computer lab at McKinley Elementary School, an underprivileged San Jose school adopted by NVIDIA for its annual employee community work day this December. The donation from drinks and a special raffle will be matched by the NVIDIA Foundation. We will depart from the Exhibit Hall at at 8:00PM on Wednesday, September 30 to kick off the festivities.

GO GREEN!

Take part in the shared goal of minimizing our collective impact on the environment. Please take only the conference materials you need. Recycle what you don't, including your badge at the conclusion of the conference. Also, if your pass includes a complimentary water bottle, and you have opted in to receive this item, we encourage you to use and reuse your bottle for your hot and cold beverages to avoid contributing to more waste to the environment.

BAG CHECK

Bag check is available on the first floor at the bell desk.

LOST AND FOUND

Please see the conference info desk should you lose or find an article.

FIRST AID / EMERGENCY

Should there be a medical emergency, please dial 911 and find the nearest hotel or conference personnel.

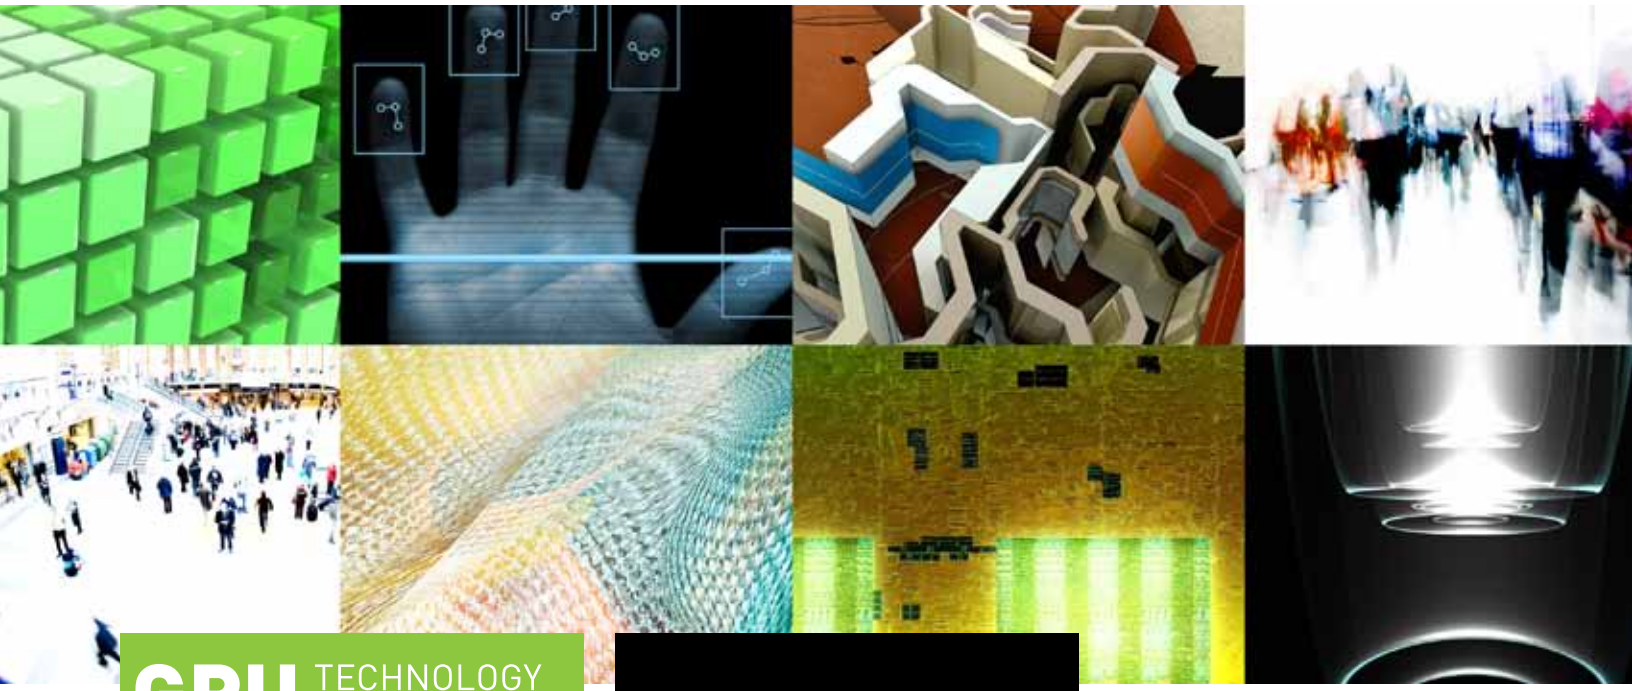

GPU TECHNOLOGY CONFERENCE

TABLE OF CONTENTS

The Republic of Gamers G51J and G71Gx notebooks deliver the high performance and unique gaming features that set it apart from the competition. Featuring Intel® Core™2 Quad processor technology and the sizzling NVIDIA® GeForce® GTX 260M GPU, these gaming power-houses let you experience the blinding explosions, reactive debris, realistic environments, and lifelike character motion that make games come to life. Plus, they're enhanced with illuminated gaming keyboards so you can dominate in the light or dark.

G51J

G71Gx

VISIT WWW.ASUSROG.COM FOR MORE INFORMATION

ASUS®
Inspiring Innovation • Persistent Perfection

REPUBLIC OF GAMERS

CHRISTIE®

3D/Advanced Visualization Solutions








World leaders in turnkey and customized integrated immersive display solutions

christiedigital.com

GPU TECHNOLOGY CONFERENCE

RECOMMENDED FOR NVIDIA RESEARCH SUMMIT ATTENDEES

WEDNESDAY / SEPTEMBER 30

SESSION ID	TITLE
1200	Opening Keynote with Jen-Hsun Huang, CEO and Co-Founder, NVIDIA
1201	Important Trends in Visual Computing (General Session)
1202	Breakthroughs in High Performance Computing (General Session)
1419	Exhibits Open / Networking Happy Hour / Research Posters Preview
1450	Opening Night Charity Social

THURSDAY / OCTOBER 30

SESSION ID	TITLE
1008	Memory Saving Discrete Fourier Transformation on CUDA
1009	Real Time Multi-Channel, Multi-Camera Image Processing using CUDA
1010	Volunteer Computing for GPUs: Petaflops for Free
1046	Research Summit Fast Forward with Bill Dally
1048	Interactive Ray Tracing with the OptiX ray tracing engine
1049	Supercomputing Super Session
1050	Quantum Chemistry on the GPU: Accelerating DFT Calculations
1051	GPU Accelerated Molecular Dynamics with AMBER
1053	GPU Accelerated Visualization and Analysis in VMD
1055	What Every CUDA Programmer Needs to Know About OpenGL
1056	Multiparticle Simulation
1421	Exhibits Open / Networking Lunch (Thursday)
1422	Day 2 Keynote with Hanspeter Pfister, Harvard University
1434	Unlocking Biologically-Inspired Computer Vision: a High-Throughput Approach
1437	Directing Experiments in the International Space Station With GPU-Assisted Image Analysis
1443	Thrust: A Parallel Template Library for CUDA
1449	Accelerating Quantum Chemistry Research Using GPUs-Two Electron Repulsion Integrals in GAMESS

GPU TECHNOLOGY CONFERENCE

FRIDAY / OCTOBER 2

SESSION ID	TITLE
1006	Implementing 3D Finite Difference Codes on the GPU
1047	Computer Vision with Horst Bischof
1058	Tridiagonal Solvers on the GPU and Applications to Fluid Simulation
1059	Fast Tridiagonal Solvers on GPU
1060	Graphcuts with CUDA and Applications in Image Processing
1062	Astrophysical Fluid Simulation Using Adaptive Meshes
1065	Diesel-Powered GPU Computing: Enabling a Real-Time Radio Telescope in the Australian Outback
1075	Reconstructing the Brain: Extracting Neural Circuitry with CUDA and MPI
1076	Optimizing Ion Channel Kinetics Using A Massively Parallel Genetic Algorithm on the GPU
1079	Scalable Multi Agent Simuluation on the GPU
1082	Advanced Numeric Computing
1083	CUDA Fortran Programming for NVIDIA GPUs
1090	Spatial Data Structures for Massively Parallel Computing
1106	Driving on Mars: Simulating Tracked Vehicle Operation on Granular Terrain
1112	Visualizing the Universe: Raycasting Astrophysical Simulation Data
1123	Domain Specific Languages for Programming GPUs
1134	Using OpenKODE to Create a Seamless Driving Experience
1142	Handheld Augmented Reality
1402	Applications of Graphics Processing Units to the Binary Black Hole Evolutions
1423	Day 3 Keynote with Richard Kerris, Lucasfilm
1424	Exhibits Open / Networking Lunch (Friday)
1441	Binary Black Holes using GPUs
1442	Numerical Cosmology Powered by GPUs

Adobe

Plan, create, and deliver anywhere with the intelligent post-production solution. **Adobe® Creative Suite® 4 Production Premium** is a must-have solution for creative professionals who need to craft world-class video, audio, and interactive media — on air, online, on device, and invariably on deadline.

Find your shortcut to brilliant at **adobe.com/go/creativesuiteproduction**

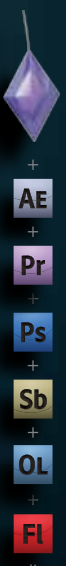

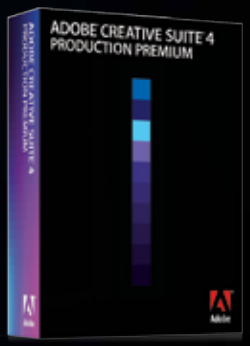

Nando Costa's Shortcut to Brilliant

DO NOT MISS THESE EVENTS!

WEDNESDAY

1:00 PM	Opening Keynote with Jen-Hsun Huang, CEO and Co-Founder, NVIDIA
3:00 PM	General Session: Important Trends in Visual Computing with David Luebke, NVIDIA Research, and special guests Horst Bischof (Graz University of Technology), Pat Hanrahan (Stanford University) and Blair McIntyre (Georgia Tech, GVU Center)
4:30 PM	General Session: Breakthroughs in High Performance Computing with Bill Dally, Chief Scientist, NVIDIA, and special guests Shawn Edwards (Bloomberg), Jeffrey Vetter (Oak Ridge National Laboratory and Georgie Tech), Steve Scott (Cray)
6:00 PM	Exhibits, Research Poster Preview and Networking Happy Hour
8:00 PM	Special Opening Night Charity Event


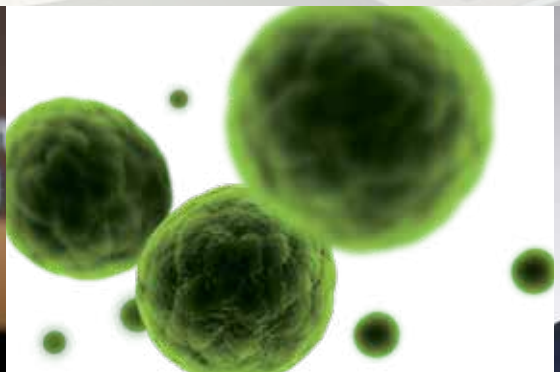

THURSDAY

9:00 AM	Day 2 Keynote with Hanspeter Pfister, Professor, Harvard University
11:00 AM	Emerging Companies Summit Opening Address
11:00 AM	NVIDIA Research Summit Fast Forward with Bill Dally
12:00 PM	Emerging Companies Summit Panel: Future Directions of GPU Computing with Steve Perlman, Founder and CEO of OnLive, Simon Hayhurst, Senior Director of Product Management of Adobe, Sean Varah, CEO of MotionDSP, and Bill Dally, Chief Scientist of NVIDIA
12:00 PM	Exhibits and Networking Lunch
3:00 PM	Supercomputing Super Session
6:00 PM	Exhibits and Networking Happy Hour
6:00 PM	NVIDIA Research Summit Poster Session and Social (open to Academics only)
All Day	Nexus Lab

FRIDAY

8:30 AM	Day 3 Keynote with Richard Kerris, CTO, Lucasfilm
10:00 AM	Emerging Companies Summit Fireside Chat with Jen-Hsun Huang and Jon Peddie
11:30 AM	Emerging Companies Summit Panel: Raising Capital in Difficult Environments
12:00 PM	Exhibits and Networking Lunch
3:30 PM	Advanced C for CUDA
3:30 PM	GPU Revolution in Film Production: Report from the Battlefield
All Day	Nexus Lab




GPU TECHNOLOGY CONFERENCE

SPECIAL OPENING NIGHT CHARITY EVENT

Join us for a party and turn your drinks into donations at our Opening Night Charity Social. Hang out in the heart of downtown San Jose, let loose to some great music, and mingle with fellow GTC attendees. The local clubs we visit will donate a percentage of food and drink revenue to fund a computer lab at McKinley Elementary School, an underprivileged San Jose school adopted by NVIDIA for its annual employee community work day this December. The donation from drinks and a special raffle will be matched by the NVIDIA Foundation. Meet at the Exhibit Hall on Wednesday night, Sept 30 at 8:00 PM to kick off the festivities.

SCHEDULE

8:00 PM	**Exhibit Hall (Imperial Ballroom)**
8:10 PM	**Tres Gringos** (83 S. Second Street, near San Fernando Street)
9:00 PM	**San Jose Bar & Grill** (85 S. Second Street, near San Fernando Street)
10:00 PM	**Voodoo Lounge** (14 S. Second Street, near Santa Clara Street)

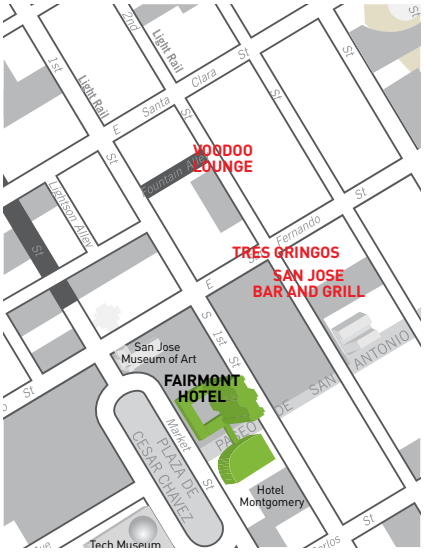

Credit: PQ Photography




SUPERMICRO®

GPU

Massively Parallel Processing

Performance Optimized for Scientific, Engineering and Graphics-Intensive Applications

SS7046GT-TRF-TC4

SS6016GT-TF-TM2

- 2 Teraflops in 1U, 4 Teraflops in 4U
- 2 non-blocking PCI-E x16 Gen II direct connections to GPU (1U)
- 4 non-blocking PCI-E x16 Gen II direct connections to GPU (4U)
- 1400W high-efficiency Gold Level 93% Power Supply with PMBus
- Dual Quad/Dual-Core Intel® Xeon® processors 5500 series
- Intelligent server IPMI management tools monitor GPU status information
- IPMI 2.0 with virtual media over LAN and KVM-over-LAN
- 8 hot-swap 3.5" SAS/SATA HDD bays (4U)

Ingram Micro	MA LABS	Synnex	SED International	Tech Data	AMAX	Arrow Electronics	ASI	Bell Micro
1-800-456-8000	1-408-941-0808	1-800-756-5974	800-444-8962	1-800-237-8931	1-800-800-6328	1-888-427-2250	1-800-2000-ASI	1-800-232-9920
www.ingrammicro.com	www.malabs.com	www.synnex.com	www.sedonline.com	www.techdata.com	www.amax.com	www.arrownacp.com	www.asipartner.com	www.bellmicro.com

www.supermicro.com

WEDNESDAY 9:00 AM – 10:30 AM
GOLD ROOM

1407 OpenGL: The Train Has Left the Station (Pre-Conference Tutorial)

The train has left the station! OpenGL innovation has greatly accelerated over the past year. Come learn about OpenGL 3 and other new features in NVIDIA's OpenGL drivers, including how to efficiently mix OpenGL and CUDA or OpenCL in your application.

Key Topic	Tools
Summit	GPU Developers Summit
Speakers	Michael Gold (NVIDIA) Mark Kilgard (NVIDIA) Barthold Lichtenbelt (NVIDIA)

WEDNESDAY 9:00 AM – 10:30 AM
CALIFORNIA ROOM

1408 DirectX 11 Overview (Pre-Conference Tutorial)

This presentation gives an overview of the DirectX 11 pipeline and how it extends previous DirectX versions to enable stunning visual effects in real-time graphics applications.

Key Topic	Tools
Summit	GPU Developers Summit
Speaker	Cem Cebenoyan (NVIDIA)

WEDNESDAY 9:00 AM – 10:30 AM
CRYSTAL ROOM

1412 Languages, APIs and Development Tools for GPU Computing (Pre-Conference Tutorial)

Get a head start on the conference with this first-day introduction to key technologies for GPU Computing. This tutorial session will cover the key features and differences between the major programming languages, APIs and development tools available today, including case studies of several applications and how they benefit from GPU Computing. Attendees will also learn several high level design patterns for consumer, professional and HPC applications, with practical programming considerations for each.

Key Topic	Tools
Summit	GPU Developers Summit
Speaker	Will Ramey (Product Manager, NVIDIA)

WEDNESDAY 10:45 AM – 12:15 PM
CRYSTAL ROOM

1409 OpenCL on the GPU (Pre-Conference Tutorial)

This presentation is an introduction to Khronos' OpenCL programming interface and how it is used to program GPUs. It describes OpenCL's platform, memory, programming, and execution models and illustrates these concepts with some simple code walkthrough.

Key Topic	Tools
Summit	GPU Developers Summit
Speakers	Neil Trevett (NVIDIA) Cyril Zeller (NVIDIA)

WEDNESDAY 10:45 AM – 12:15 PM
GOLD ROOM

1410 C on the GPU (Pre-Conference Tutorial)

This presentation teaches the basics of programming GPUs using the C language with CUDA extensions. No prior experience in GPU programming is required. The concepts – data transfers, kernel execution, memory model, synchronization – are introduced progressively and illustrated with step-by-step walkthroughs of code samples.

Key Topic	Tools
Summit	GPU Developers Summit
Speaker	Timo Stich (Developer Technology Engineer, NVIDIA)

WEDNESDAY 10:45 AM – 12:15 PM
CALIFORNIA ROOM

1411 DirectCompute (Pre-Conference Tutorial)

DirectCompute adds the new "Compute Shader" to its API, making possible GPU computing within the familiar programming environment. Compute shaders provide the programmer with greater flexibility and features such as thread cooperation, shared memory, and unordered data access. In this talk, we'll introduce the DirectCompute API by way of building simple examples from the ground up to illustrate the new processing possibilities of DirectCompute.

Key Topic	Tools
Summit	GPU Developers Summit
Speaker	James Fung (Developer Technology, NVIDIA)

WEDNESDAY 1:00 PM – 2:00 PM
REGENCY 1+2 ROOM

1200 Opening Keynote with Jen-Hsun Huang, CEO and Co-Founder, NVIDIA

Jen-Hsun Huang, NVIDIA's CEO and Co-Founder will open the GPU Technology Conference with this can't-miss keynote. Due to overwhelming response and limited capacity if you are interested in attending any of the keynotes or general sessions, be sure to schedule this session into your personal calendar.

Key Topic	General Interest
Summits	Emerging Companies Summit NVIDIA Research Summit GPU Developers Summit
Speaker	Jen-Hsun Huang (CEO and Co-Founder, NVIDIA)

WEDNESDAY 3:00 PM – 4:00 PM
REGENCY 1+2 ROOM

1201 Important Trends in Visual Computing (General Session)

Throughout the GPU Technology Conference, attendees will be exposed to some of the hottest trends and biggest ideas in the field of visual computing. In this session, moderated by David Luebke of NVIDIA Research, a collection of world-renowned experts will highlight and preview important visual computing trends ranging from computer vision, augmented reality and visual analytics.

Due to overwhelming response and limited capacity if you are interested in attending any of the keynotes or general sessions, be sure to schedule this session into your personal calendar.

Key Topics	Computer Vision, Visualization, General Interest
Summits	Emerging Companies Summit NVIDIA Research Summit GPU Developers Summit
Speakers	Horst Bischof (Professor, Graz University of Technology); Pat Hanrahan (Professor, Stanford University) Blair MacIntyre (Associate Professor, Georgia Tech, GVU Center)

WEDNESDAY 4:30 PM – 5:45 PM
REGENCY 1+2 ROOM

1202 Breakthroughs in High Performance Computing (General Session)

GPU computing is one of the biggest topics of discussion in the HPC industry today. This session, hosted by NVIDIA Chief Scientist Bill Dally, showcases perspectives from industry luminaries on the future of HPC and the role of the GPU.

Due to overwhelming response and limited capacity if you are interested in attending any of the keynotes or general sessions, be sure to schedule this session into your personal calendar.

Key Topic	General Interest
Summits	Emerging Companies Summit NVIDIA Research Summit GPU Developers Summit
Speaker	Bill Dally (Chief Scientist, NVIDIA)

WEDNESDAY 6:00 PM – 8:00 PM
IMPERIAL BALLROOM

1419 Exhibits Open / Networking Happy Hour / Research Posters Preview

Monday Exhibit / Networking / Posters Preview: Join your colleagues in the exhibit hall to preview emerging technologies and see some of the most innovative solutions available today. We also encourage you to take the opportunity to browse the NVIDIA Research Summit posters to see "what's next" out of the world of research and academia. Appetizers and drinks will be served.

Key Topic	General Interest
Summits	Emerging Companies Summit NVIDIA Research Summit GPU Developers Summit

WEDNESDAY 8:00 PM – 10:00 PM
SPECIAL EVENT

1450 Special Event & Charity Fundraiser

Join us for a party and turn your drinks into donations at the our Opening Night Charity Social. Hang out in the heart of downtown San Jose, let loose to some great music, and mingle with fellow GTC attendees. Donations will fund a computer lab at McKinley Elementary School, an underprivileged San Jose school adopted by NVIDIA for its annual employee community work day on December 5. Join the festivities as we kick off at 8 PM as we kick off from the Fairmont. See page 12 for details.

Key Topic	General Interest
Summits	Emerging Companies Summit NVIDIA Research Summit GPU Developers Summit

GPU TECHNOLOGY CONFERENCE

THURSDAY

THURSDAY 9:00 AM – 10:30 AM
REGENCY 1+2 ROOM

1422 Day 2 Keynote with Hanspeter Pfister, Harvard University

High-Throughput Science

How did the universe start? How is the brain wired? How does matter interact at the quantum level? These are some of the great scientific challenges of our times, and answering them requires bigger scientific instruments, increasingly precise imaging equipment and ever-more complex computer simulations. In his keynote address, Harvard professor, researcher and computing visionary Hanspeter Pfister will discuss the computational obstacles scientists face and how commodity high-throughput computing can enable high-throughput science, in which massive data streams are processed and analyzed rapidly -- from the instrument through to the desktop. Finally Professor Pfister will survey several groundbreaking projects at Harvard that leverage GPUs.

for high- throughput science, ranging from radio astronomy and neuroscience to quantum chemistry and physics.

Due to overwhelming response and limited capacity if you are interested in attending any of the keynotes or general sessions, be sure to schedule this session into your personal calendar.

Key Topic	General Interest
Summits	Emerging Companies Summit, NVIDIA Research Summit, GPU Developers Summit
Speaker	Hanspeter Pfister (Professor, Harvard University)

THURSDAY 11:00 AM – 11:30 AM
SACRAMENTO ROOM

1406 Embedded Tesla-Using CUDA and Tesla in a Medical Device

The TechniScan Whole Breast Ultrasound system images the female breast using a computationally intensive technique known as inverse scattering. Older prototypes used clusters of Pentium M computers to compute the WBU images in about four and one half hours. By porting the existing Fortran code to C and CUDA, the same algorithm could be run in approximately 16 minutes on 4 first generation Tesla GPUs. I will detail the process porting an MPI-based Fortran algorithm to multiple GPUs and CUDA, including issues and successes, cost, performance, and code design and implementation.

Key Topics	Computational Imaging, Medical Imaging & Visualization, Visualization
Summit	GPU Developers Summit
Speaker	Jim Hardwick, (Senior Software Engineer, TechniScan Medical Systems, Inc.)

THURSDAY 11:00 AM – 12:00 PM
BELVEDERE ROOM

1001 From Brook to CUDA C

While computing with the GPU may seem like a new technology, the foundations upon what is available today have a legacy which date back to the original parallel supercomputers. The research community first identified the GPU as a computing platform and shown the promise to solve many of the world's compute intensive problems, many orders of magnitude faster the conventional CPUs. Once the opportunity became obvious, the challenge was how to best evolve a general purpose programming model to preserve the GPU's architectural advantage. In this talk, I will provide both a perspective on how GPU computing evolved from early published results of using GPU as a computing platform, to the early programming models and tools, what the state of the art is today, and extrapolate where GPU computing is going.

Key Topic	General Interest
Summit	GPU Developers Summit
Speaker	Ian Buck (NVIDIA)

THURSDAY 11:00 AM – 12:00 PM
ATHERTON ROOM

1003 Overview of Computational Finance on the GPU

Financial institutions make extensive use of computers in solving models of the instruments they trade. In this talk we address the questions of how these instruments are modeled and what the results mean to the recipients. We explore applications of computational finance with a focus on the problems that are faced with existing systems today, why speed and accuracy are important, and how GPU technology can help.

Key Topics	Computational Finance, High Performance Computing
Summit	GPU Developers Summit
Speaker	Thomas Bradley (NVIDIA)

THURSDAY 11:00 AM – 12:00 PM
CALIFORNIA ROOM

1007 Digital 3D Entertainment

Digital 3D entertainment is suddenly catching fire. It is hard to find two letters that are changing the thinking and technology in so many industries in the digital entertainment value chain. This talk will cover digital 3D entertainment from the theater to the home. Along the way the audience will hear about some of the graphics intense challenges that people are facing.

Key Topic	3D
Summit	GPU Developers Summit
Speaker	Gary Sasaki (President, DIGDIA)

THURSDAY 11:00 AM – 12:00 PM
CRYSTAL ROOM

1013 Large Scale Visualization

Current GPUs can create incredibly realistic images for digital prototyping, however, sometimes realism alone is not sufficient to gain sufficient understanding of a proposed solution or design alternative. In many industries it is very important to further enhance realism in the image with large scale display technologies and immersive environments. This talk discusses many of the technologies and factors associated with large scale visualization, including synchronized multi-channel graphics, 3D stereo configurations, high bit-depth displays, as well as highlighting considerations when many such technologies are combined together.

Key Topic	Visualization
Summit	GPU Developers Summit
Speaker	Ian Williams (NVIDIA)

THURSDAY 11:00 AM – 12:00 PM
EMPIRE ROOM

1018 Sparse Linear Algebra – Iterative Solvers and Preconditioners on GPUs for Engineering Simulations

Sparse linear algebra solvers are used in many areas of scientific computing and are a key target for GPU acceleration. In this session, Acceleware will present an overview of its linear algebra libraries including support for direct and sparse iterative solvers. We will discuss the supported algorithms, performance results, and their application to engineering simulations. We will also look at some of the more commonly used preconditioners and some of the challenges faced when accelerating these algorithms on the GPU.

Key Topics	Algorithms & Numerical Techniques, Tools
Summit	GPU Developers Summit
Speaker	Michal Okoniewski (Interim President & CEO, Acceleware)

THURSDAY 11:00 AM – 12:00 PM
GARDEN ROOM

1023 NEXUS: A Powerful IDE for GPU Computing on Windows

Unleash your productivity with NEXUS, NVIDIA's new development environment for GPU Computing and graphics applications that use CUDA C, OpenCL, DirectCompute, Direct3D, or OpenGL. NEXUS introduces native GPU debugging and platform-wide performance analysis tools for both computing and graphics developers, fully integrated into Visual Studio 2008.

In this session, you will learn to use the powerful NEXUS debugger to set breakpoints and step through your CUDA C kernels or HLSL shader source code. Easily catch hard-to-find bugs using the NEXUS memory inspection tools and set data breakpoints on your running application. Visualize your application's workloads and performance characteristics across the CPU, GPU, and operating system, and then dig deeper using the NEXUS profilers. NEXUS also includes state of the art graphics debugging features and convenient API state inspection. This session will include information on how to register for the NEXUS Beta Program.

Key Topic	Tools
Summit	GPU Developers Summit
Speaker	Sebastien Domine (Sr. Dir. Dev Tech Tools, NVIDIA)

THURSDAY 11:00 AM – 12:00 PM
GOLD ROOM

1028 Performance Primitives for Video Codec and Image Processing

The NVIDIA Performance Primitives library (NVPP) is a library of highly optimized data-processing functions that execute on the CUDA architecture. It currently comprises 75 functions for image processing. The presentation will give an overview of the problem domains covered by those functions, what image types they support, and how they can be used to implement CUDA based image-processing solutions. The talk also gives data on how NVPP's performance compares to CPU based implementations as well as how to design for maximum performance using NVPP. Video encoding and decoding tasks are among the most computationally demanding features of consumer and professional applications today. NVIDIA's CUDA architecture for GPU Computing was designed to handle the large computational complexity of these tasks with relative ease. In addition to a hardware video decoder, we will present several software codec libraries written in CUDA and\or OpenCL.

Key Topics	Computational Imaging, Tools
Summit	GPU Developers Summit
Speakers	Frank Jargstorff (Developer Technology Engineer, NVIDIA); Anton Obukhov (Developer Technology Engineer, NVIDIA)

THURSDAY 11:00 AM – 12:00 PM
HILLSBOROUGH ROOM

1031 Designed for CUDA Program Overview

As the leader in Visual Computing, NVIDIA provides extensive development and marketing support for key ISV partners. ISVs can take advantage of NVIDIA's "Designed for CUDA" program, which provides a wide range of marketing and development tools available exclusively to developers of compelling new CUDA-enabled consumer applications. You'll learn how to leverage NVIDIA's worldwide marketing influence in PR, distribution, promotions, online presence and more. In addition, we will discuss details of OpenAutomate, a technical resource that provides comprehensive compatibility testing, performance measurement and tuning across the full spectrum of NVIDIA configurations. Technical prerequisites for certification of the program will be explained in detail, including automation, installation and licensing requirements.

Key Topic	General Interest
Summit	GPU Developers Summit
Speakers	John Spitzer (Sr Director of Content Infrastructure, NVIDIA); Michael Steele (General Manager, Visual Consumer Solutions, NVIDIA)

THURSDAY 11:00 AM – 12:00 PM
REGENCY 2 ROOM

1046 Research Summit Fast Forward with Bill Dally

NVIDIA embraces emerging technology and those who dedicate their research to growth in visual and parallel Computing. We believe that these minds lead the future in our industry and we are proud to support the 2009-2010 NVIDIA Fellows. It is our pleasure to showcase their research during the Fellow Fast Forward. Research topics include: Computer Vision, Ray Tracing, Quantum Chemistry, VLSI, Shadow Mapping and much more. For more information on the 2009-2010 NVIDIA Fellows, please visit http://www.nvidia.com/page/fellowship_programs.html.

Key Topic	General Interest
Summit	NVIDIA Research Summit
Speakers	2009-2010 NVIDIA Fellows: Anjul Patney (University of California, Davis), Bryan Catanzaro (University of California, Berkeley), Erik Sintorn (Chalmers University of Technology), Gregory Diamos (Georgia Institute of Technology), Huy T. Vo (University of Utah), Ivan Ufimtsev (Stanford University), Jiayuan Meng (University of Virginia), Nicolas Pinto (Massachusetts Institute of Technology), Rahul Garg (University of Washington) and Yen-Tzu Lin (Carnegie Mellon University)

THURSDAY 11:00 AM – 12:00 PM
VALLEY ROOM

1050 Quantum Chemistry on the GPU: Accelerating DFT Calculations

In this talk we describe a high performance exchange-correlation Fock matrix implementation in DFT accelerated using NVIDIA CUDA toolkit. A BLAS3 kernel based algorithm has been developed to achieve sustained performance on both GPU and CPUs. Benchmark calculation shows more than 30 times speedup using CPU (Quadcore Phenom X4 940) and GPU (Tesla C1060) together compared with original code on single Phenom core. Optimization techniques and heterogeneous computing strategies for acceleration computing will be discussed and illustrated with implementation examples.

Key Topic	Quantum Chemistry
Summits	GPU Developers Summit, NVIDIA Research Summit
Speaker	Zhengting Gan (Q-Chem, Inc.)

THURSDAY 11:00 AM – 12:00 PM
PIEDMONT ROOM

1098 NEXUS Lab

Experience the future of GPU Computing development today! NEXUS is NVIDIA's new GPU Computing development environment for Microsoft Visual Studio. In this hands-on lab, attendees will learn to use the NEXUS source debugger and profiler to efficiently solve common problems encountered when debugging and profiling CUDA C code, including

> Debugging logic errors using hardware breakpoints and variable inspection

> Reading and visualizing GPU memory directly in the Visual Studio IDE (no manual readbacks required!)

> Real-time expression evaluation on the GPU while your application is halted at a breakpoint

> Hardware-evaluated breakpoint conditionals for debugging a specific block or thread

> Detecting out-of-bounds memory errors

> Finding erroneous memory overwrites using data breakpoints

> Understanding the performance of your whole system using the Nexus Analysis tools

NVIDIA engineers will be available at each session to answer your NEXUS and general GPU Computing questions. Lab workstations are limited to the first 20 attendees per session. NVIDIA would like to thank HP and Parallels for their generous support of SLI Multi-OS workstations.

Key Topics	General Interest, Tools
Summit	GPU Developers Summit
Speakers	Brian Kretzler (NVIDIA); Greg Smith (Manager of Developer Tools, NVIDIA)

THURSDAY 11:00 AM – 12:00 PM
FAIRFIELD ROOM

1101 GPU Computing in Java: Generation and Scheduling of PTX Assembly

This talk will discuss an ongoing progress regarding the development of a Java-based library for rapid kernel prototyping in NVIDIA PTX and PTX instruction scheduling. It is aimed at developers seeking total control of emitted PTX, highly parametric emission of, and tunable instruction reordering. It is primarily used for code development at ICHEC but is also hoped that NVIDIA GPU community will also find it beneficial.

Key Topic	Tools
Summit	GPU Developers Summit
Speakers	Christos Kartsaklis (Computational Scientist, Irish Centre for High-End Computing (ICHEC))

THURSDAY 11:00 AM – 12:00 PM
REGENCY 2 ROOM

1446 Emerging Companies Summit Opening Address

The Emerging Companies Summit is a unique forum for startup companies to showcase innovative applications that leverage the GPU to solve visual and high-performance computing problems. The Opening Address includes selected demos from a few presenting companies. The ECS is a great opportunity to discover new players in the visual and high performance computing ecosystems, find great investments, explore partnership opportunities, network/build relationships, and discuss the future of an industry that is reshaping computing.

Key Topic	General Interest
Summit	Emerging Companies Summit
Speaker	Jeff Herbst (Vice President of Business Development, NVIDIA) Demonstrated by: Edge 3 Technologies, Luminova, MirriAd and Viewdle

In more than 100 countries around the world, Citi is helping companies, governments and institutions overcome business challenges, raise capital, mitigate risk and extend their reach.

When you partner with Citi, you gain access to an unparalleled global platform, capital markets, insightful advice and award-winning solutions — so you can realize your goals today and in the future.

Success requires both financial and intellectual capital, and that's why Citi never sleeps.

Please visit us at www.icg.citi.com.

Citi never sleeps℠

citi®

THURSDAY 11:30 AM – 12:00 PM
GLEN ELLEN ROOM

1456 Accelerating Finite Difference Time Domain (FDTD) Simulation with CUDA

Advances in a variety of fields ranging from cellular phone design to MRI coil design necessitate the use of high accuracy electromagnetic computational methods such as Finite Difference Time Domain (FDTD) to analyze complex new products. High fidelity techniques such as FDTD generally yield highly accurate results at the cost of increased computation time. Using the processing power of NVIDIA's CUDA architecture allows this time to be reduced by over two orders of magnitude. This session presents an implementation of the FDTD method using NVIDIA's CUDA technology. The core algorithm is presented in C code and translated to a CUDA kernel. Branching statements are eliminated in order to improve parallelism. Shared memory and an optimized looping method are employed to avoid memory communication bottlenecks. Timing comparisons against an optimized commercial FDTD code are given.

Key Topic	General Interest
Summit	GPU Developers Summit
Speaker	James Stack (Remcom Inc.)

THURSDAY 12:00 PM – 12:30 PM
HILLSBOROUGH ROOM

1032 Compute Application Testing

In applications ranging from entertainment to oil/gas research to financial market prediction, GPU Computing offers many-fold increases in performance over traditional CPU-based approaches. But with this also comes some complexity on the development side with multiple flavors of APIs and GPUs to take into account. In this moderated roundtable, NVIDIA engineers will lead a discussion on best practices for testing GPU Computing applications across the full spectrum of potential end-users. Attendees will share their approaches, successes and failures in testing such compute applications for compatibility, stability and performance.

Key Topic	General Interest
Summit	GPU Developers Summit
Speaker	John Spitzer (Sr Director of Content Infrastructure, NVIDIA)

THURSDAY 12:00 PM – 1:00 PM
REGENCY 2 ROOM

1425 Future Directions in GPU Computing

Join our panel of leading industry experts as they discuss some of the most exciting and disruptive current and future applications built to run on the GPU platform. Topics will cover both visual and GPU computing, including rendering, video processing, scientific computing, cloud computing, and more.

Key Topic	General Interest
Summit	Emerging Companies Summit
Speaker	Moderator: Rob Enderle, President and Principal Analyst – Enderle Group; Panelists: Bill Dally, Chief Scientist – NVIDIA; Simon Hayhurst, Senior Director of Product Management – Adobe; Steve Perlman, Founder and CEO – ONLIVE; and Sean Varah CEO – MotionDSP.

THURSDAY 12:00 PM – 2:00 PM
IMPERIAL BALLROOM

1421 Exhibits Open / Networking Lunch (Thursday)

Join your colleagues in the exhibit hall to preview emerging technologies and see some of the most innovative solutions available today. Lunch will be served to Full Conference, Research Summit and Emerging Companies Summit pass holders.

Key Topic	General Interest
Summit	Emerging Companies Summit, GPU Developers Summit, NVIDIA Research Summit

THURSDAY 1:00 PM – 1:30 PM
FAIRFIELD ROOM

2002 Viewdle Inc. Startup Presentation

By unleashing Viewdle technology, content producers and distributors can effectively leverage their video assets to open new revenue streams and attract, engage and retain customers/visitors. Initially focused on providing internal and private-label indexing and search capabilities to news and entertainment media, we will soon launch targeted advertising solutions, developer tools and a consumer-facing search site. Viewdle was founded in 2006 to revolutionize video search by solving the inefficiencies inherent in using text-based metadata and other existing tools for indexing moving media. The Viewdle engine fuses pioneering visual analysis technology, continuously developed over the past 35 years, along with other leading search techniques to go "inside the clip" and identify contextually relevant appearances based on actual video assets.

Key Topic	Computational Imaging, Visualization
Summit	Emerging Companies Summit
Speaker	Laurent Gil (CEO and President, Viewdle Inc.)

THURSDAY 1:00 PM – 1:30 PM
CUPERTINO ROOM

2014 Elemental Technologies, Inc. Startup Presentation

Elemental Technologies is the leading provider of massively parallel processing solutions for video applications. Using off-the-shelf, programmable graphics processing units (GPUs) for compute-intensive video processing and conversion tasks, Elemental's products are ideally suited for a variety of video applications including consumer media processing, professional video editing and Internet video processing. Founded in 2006, Elemental is headquartered in Portland, Oregon.

Key Topics	Computer Vision, General Interest
Summit	Emerging Companies Summit
Speaker	Sam Blackman (CEO, Elemental Technologies)

THURSDAY 1:00 PM – 1:30 PM
GLEN ELLEN ROOM

2018 C3 Technologies Startup Presentation

World's most realistic 3D city models

Swedish C3 Technologies fully automatically creates huge 3D city models from images only. The flagship product C3 Maps enables online mapping providers a much richer map display and the possibility for the user to fly around in 3D to find and explore locations. While current map solutions offer 2D maps and ortho photos C3 offers unique realistic high resolution 3D imagery where everyone can dive down to see their own back yard in 3D. The possibility of free navigation (fly overs) is today available for the web and desktop solutions and high-end mobile units and navigation devices. Using a highly efficient algorithm from the Swedish defence industry C3 computes a textured elevation model from thousands of images in a few hours time. A 3D model can be delivered within weeks from the photography, containing whole large cities, not only city centers. With tremendous amounts of 3D data streamed over the Internet, C3 carefully benefits from the full ability of 3D hardware, challenging old 3D city model concepts with more elaborated level of detail technology. The algorithms are currently used for aerial and satellite imagery and during fall of 2009 C3 plans a release of unique street level products using similar technology. C3 and its products has been noticed by slashdot and furthermore the models are claimed by the press to beat Google Earth and Microsoft Virtual Earth content.

Key Topics	3D, Embedded & Mobile
Summit	Emerging Companies Summit
Speaker	Mattias Astrom (Board member, C3 Technologies)

THURSDAY 1:00 PM – 1:30 PM
BELVEDERE ROOM

2041 Luminova Startup Presentation

Luminova utilises core NVIDIA-mental images 64-bit visualisation technologies in uniquely providing large scale construction/engineering project management systems through a web based, real-time interactive, highly detailed and photographic 3D visual information portal, utilising standard files from the existing project CAD system, where all of this data is communicated through our proprietary on-line interface application, Project Hub. Luminova's technology system uniquely handles data sets hundreds of times larger than any other existing technology. Our system revolutionises the speed, scale, accuracy and communication of the complete design and construction detail; resulting in significantly shortened design/engineering programs by providing a faster and a far more assured decision making process. The entire 'as built' construction detail can be readily assembled into a single environment that can all be viewed, measured, analysed, tested and shared in real time, by any number of concurrent users, anywhere in the world, through a standard web browser. This system has already been proven on many very large projects including Heathrow Airport – Terminal 5, Greenwich Peninsula, Wembly Precinct and London 2012 Olympic Games Facilities and many others.

Key Topics	3D, Visualization
Summit	Emerging Companies Summit
Speaker	Christopher Blewitt (CEO, Luminova)

THURSDAY 1:30 PM – 2:00 PM
CUPERTINO ROOM

2005 Stormfjord Startup Presentation

Stormfjord is a leader in interactive visualization of very large scale industrial data, specialising in Gpu-centric simulation. We provide next generation tools to drive novel industrial use-cases enhancing complex planning, work process modeling, robotic prototyping, training, control, and monitoring. Our tools are best in their class, derived from the bleeding edge of technology, and open to 3rd party developers.

Key Topic	Visualization
Summit	Emerging Companies Summit
Speaker	Thorolf Tonjum (Director of R&D, Stormfjord)

THURSDAY 1:30 PM – 2:00 PM
FAIRFIELD ROOM

2015 Edge 3 Technologies Startup Presentation

Edge 3 Technologies is an offshoot of the Embry-Riddle Machine Vision Lab, the premier research facility in 3D gesture recognition with GPU technology. We specialize in the development of GPU middleware with advanced learning, coupled with machine vision for sophisticated Fortune 500 companies. Our technology is up to 100x faster and more accurate than conventional CPU-based technology. The company's goal is to help redefine the relationship between humans and machines. Our methodology involves:

> CUDA/openCL/GPU to mitigate computationally challenging aspects and improve real-time object segmentation and tracking performance, and reduce error rates

> CUDA/GPU for machine learning/training on massive amounts of data to improve classification and identification of various user gestures and postures

> Machine learning of users' behavior and habits with hundreds of features

> A new AI capable of identifying facial gestures as well as hand gestures, within the same framework

> Differentiating users with various biometrics and depths

> Usability testing on thousands of subjects

> A portfolio of patent applications in the field

Key Topics	3D, Computer Vision
Summit	Emerging Companies Summit
Speaker	Tarek El Dokor (CEO and CTO, Edge 3 Technologies)

THURSDAY 1:30 PM – 2:00 PM
BELVEDERE ROOM

2029 MotionDSP Inc. Startup Presentation

MotionDSP makes revolutionary video software that dramatically improves video from a wide range of sources—from mobile phones to surveillance cameras. Its patented, multi-frame video enhancement technology is available in Ikena, a "CSI-style" video forensic application used by intelligence agencies and law enforcement, and in vReveal, a Windows application for consumers, launched in partnership with NVIDIA in March 2009, which offers a one-click fix for the issues that plague consumer-generated video: shake,

noise artifacts, pixelation, poor lighting, and low resolution. MotionDSP's investors include In-Q-Tel, the independent investment firm that invests to support the mission of the CIA and US Intelligence Agencies, and NVIDIA, the world leader in visual computing. MotionDSP is based in San Mateo with a development office in Nis, Serbia.

Key Topics	Embedded & Mobile, Visualization
Summit	Emerging Companies Summit
Speaker	Sean Varah (CEO, MotionDSP)

THURSDAY 1:30 PM – 2:00 PM
GLEN ELLEN ROOM

2042 Emergent Game Technologies Startup Presentation

Through our state-of-the-art software platforms, Gamebryo LightSpeed and Gamebryo, Emergent Game Technologies helps developers fulfill their maximum creative potential by giving them the power to endlessly – and efficiently – iterate development. From prototype to final product, Emergent fuels creativity while reducing risk.

Key Topic	General Interest
Summit	Emerging Companies Summit
Speaker	Geoff Selzer (CEO, Emergent Game Technologies)

THURSDAY 2:00 PM – 2:30 PM
VALLEY ROOM

1051 GPU Accelerated Molecular Dynamics with AMBER

Molecular Dynamics (MD) simulations represent a critical scientific tool in improving our understanding of enzymatic processes and are a key component in the development of advanced biological catalysts and next generation pharmaceuticals. AMBER represents one of the most widely used MD engines with worldwide reach and thus efficient porting of this software to NVIDIA GPUs promises to have a global impact on MD research. Molecular Dynamics simulations rely on two critical factors; the ability to simulate large regions of phase space quickly and to do this accurately such that both thermodynamic and kinetic data can be obtained. The main issue facing effective use of MD simulations for drug discovery and enzymatic optimization is the ability to simulate long enough timescales to observe rare, but important structural changes. These occur on time scales of a few nanoseconds to milliseconds. However, the inherent fast motions of atoms means that integration time steps are typically on the order of femtoseconds thus a single microsecond long protein MD trajectory requires on the order of 10^9 integrations all of which have to be carried out sequentially. The high arithmetic performance and intrinsic parallelism of GPUs makes then an enticing platform for accelerating MD simulations to the point where microseconds of MD can be carried out on a desktop instead of the extremely expensive and energy intensive supercomputers required at present. This talk will present details about the CUDA implementation of MD in AMBER. This will include an overview of the approaches used, example benchmarks, and lessons that were learned in ensuring that the simulation accuracy was preserved at all times. Approaches used to validate all aspects of the scientific information generated by the GPU code will be covered. Finally some example scientific projects that are benefiting from NVIDIA GPU support in AMBER will be presented along with details of future efforts.

Key Topic	Life Sciences, Molecular Dynamics
Summits	GPU Developers Summit, NVIDIA Research Summit
Speaker	Ross Walker (The Scripps Research Institute and San Diego Supercomputer Center)

THURSDAY 2:00 PM – 2:30 PM
SACRAMENTO ROOM

1414 State of the Art in GPU-Accelerated Techniques for Medical Imaging

This presentation will provide an overview of world-class medical imaging algorithms exploiting the immense computational power of modern GPUs. The range of techniques where NVIDIA GPUs have proven invaluable range from advanced visualization techniques such as volume rendering of large CT data sets and time-varying volumetric Ultrasound data to GPU-based reconstruction, non-linear filtering, non-rigid registration, segmentation and flow simulation. Siemens Corporate Research conducts research and development for all Siemens Healthcare business units and has a large world-wide network of clinical and academic collaborators. We will present examples of a broad spectrum of clinical applications, both for post-processing as well as real time interventional procedures, which have recently become feasible only thanks to advances in programmable GPUs.

Key Topic	Medical Imaging & Visualization
Summit	GPU Developers Summit
Speaker	Gianluca Paladini (Program Manager, Siemens Corporate Research)

THURSDAY 2:00 PM – 2:30 PM
ATHERTON ROOM

1431 GPU Acceleration of Mortgage CMO Option Adjusted Spread

Mortgage Backed Security Option Adjusted Spread (OAS) analysis requires computationally intensive Monte Carlo simulation of dynamic cash flows at different interest rate scenarios. This talk will describe a port of the Bloomberg L.P. MBS OAS legacy code to CUDA, and will compare the performance and cost of running the calculations on a CPU cluster versus a CUDA-enabled GPU cluster. Significant speedup and power savings are achieved using a CUDA enabled OAS server. Additionally, the development and rollout process of a CUDA-based application to production will be presented.

Key Topic	Computational Finance
Summit	GPU Developers Summit
Speakers	James Hook (Bloomberg); Xusheng Tian (Bloomberg)

THURSDAY 2:00 PM – 2:30 PM
CUPERTINO ROOM

2011 Metaio, Inc. Startup Presentation

Metaio designs, develops and markets solutions based on augmented reality. This innovative technology allows virtual 3D information to be superimposed seamlessly into the real environment at real-time. metaio was founded in 2003 and is fully owned by CEO Dr. Thomas Alt and CTO Peter Meier.

Key Topic	3D
Summit(	Emerging Companies Summit
Speaker	Thomas Alt (CEO, Metaio, Inc.)

THURSDAY 2:00 PM – 2:30 PM
FAIRFIELD ROOM

2045 NaturalMotion Limited Startup Presentation

NaturalMotion Ltd. is a leading entertainment software company with offices in Oxford (England) and San Francisco (California). The company produces the widely-adopted animation technologies euphoria,morpheme and endorphin, used across the game and movie industries by companies such as Rockstar Games, LucasArts, Disney and Bioware as well as in the development of Backbreaker, the company's first in-house game.

Key Topic	General Interest
Summit	Emerging Companies Summit
Speaker	Torsten Reil (CEO, NaturalMotion Ltd)

THURSDAY 2:00 PM – 2:30 PM
GLEN ELLEN ROOM

2047 IQ Engines Startup Presentation

IQ Engines was founded as a collaboration of computer neuroscientists at UC Berkeley and UC Davis. We are inspired by the brain's vision system – how the brain represents images and how the brain recognizes objects. Our team has mathematical expertise in sparse coding and neural networks. It is our goal to bring advances in biological vision models to practical image and video search, using algorithms that are hierarchical and massively parallel (fast and accurate). We deliver these advances in image and video search on a web server platform that can be used by our customers to solve image and video search applications. Our team is the first in the world to deliver any-image labelling with 100% success and near-100% accuracy. A demonstration of the image identification technology can be seen at http://www.iqengines.com/wb/update.php.

Key Topic	General Interest, Visualization
Summit	Emerging Companies Summit
Speaker	Gerry Pesavento (-)

THURSDAY 2:00 PM – 2:30 PM
BELVEDERE ROOM

2057 OnLive Startup Presentation

OnLive is launching the world's highest performance Games On Demand service, instantly delivering the latest high-end titles over home broadband Internet to the TV and entry-level PCs and Macs.

Founded by noted technology entrepreneur Steve Perlman (WebTV, QuickTime) and incubated within the Rearden media and technology incubator, OnLive spent seven years in stealth development before officially unveiling in March 2009.

OnLive, together with its Mova subsidiary, lies directly at the nexus of several key trends, all of which are reshaping the way we think about and use digital media:

> The shift to cloud computing, displacing the limitations, cost and complexity of local computing;

> An explosion of consumer broadband connectivity, bringing fast bandwidth to the home;

> Unprecedented innovation, creativity and expansion within the video game market.

Pioneering the delivery of rich interactive media to the home, OnLive will change the way that entertainment applications are created, delivered and consumed.

Key Topics	General Interest, Molecular Dynamics
Summit	Emerging Companies Summit
Speaker	Steve Perlman (Founder and CEO, OnLive)

THURSDAY 2:00 PM – 3:00 PM
CALIFORNIA ROOM

1011 Have You Heard What a GPU Can Do? – A Revolution in Audio

This talk is divided into two parts and details how GPUs are now being used as well as delivering significant benefits for audio processing and editing. The first section explains what features of GPUs and the CUDA language provide key benefits for audio processing. The second part will be given by Rudy Sarzo, a world renowned recording and performing professional rock bassist, and focuses on how the GPU can be integrated into a production audio setup as well as the benefits in doing so.

Key Topic	General Interest
Summit	GPU Developers Summit
Speakers	Rudy Sarzo (-); Ian Williams (NVIDIA)

THURSDAY 2:00 PM – 3:00 PM
EMPIRE ROOM

1019 CULA: Robust GPU Accelerated Linear Algebra Libraries

EM Photonics has been working in the field of accelerated linear algebra functions for several years. In a recent project funded by NASA, we have developed GPU versions of the most widely used LAPACK routines in partnership with NVIDIA. Available to the public in the summer of 2009, CULA will be a valuable asset to an extremely broad audience as it is easy to use and provides speedups to many applications. Our session at the GPU Developer Summit will show the broad range of functionality present in CULA We will demonstrate the ease of using CULA and then present some technical details and "lessons learned" while implementing our routines. CULA will also be on display at our booth, featuring live-running demos of some CULA-based applications.

Key Topics	Algorithms & Numerical Techniques, Tools
Summit	GPU Developers Summit
Speakers	John Humphrey (Senior Engineer, EM Photonics); Eric Kelmelis (CEO, EM Photonics); Daniel Price (EM Photonics)

THURSDAY 2:00 PM – 3:00 PM
HILLSBOROUGH ROOM

1033 Using CUDA within Mathematica

Mathematica comes with many extremely optimized numerical libraries integrated into the application, but they don't yet take advantage of the GPU. Thankfully, Mathematica provides an easy to use API for communicating between a large variety of

external resources, called MathLink. This tutorial will provide a hands-on introduction to start using CUDA within Mathematica, an introduction to the cuda mathematica plugin, as well as the different issues one has to keep in mind when writing MathLink applications using the CUDA Toolkit. Finally we will showcase a few real-world examples.

Key Topic	Tools
Summit	GPU Developers Summit
Speaker	Kashif Rasul (Nomad Labs)

THURSDAY 2:00 PM – 3:00 PM
REGENCY 2 ROOM

1048 Interactive Ray Tracing with the OptiX ray tracing engine

OptiX is the latest and most flexible API for interactive ray tracing, taking full advantage of the power and generality of the CUDA architecture. This new technology is valuable for anyone wanting to build a high-performance, ray tracing renderer (interactive or off-line), accelerate an existing ray tracing renderer, add raytrace capabilities to raster renderers, or even perform generic ray tracing functions for computation. By replacing fixed components of a traditional ray tracer with user-specified programs, OptiX is able to implement a rich array of rendering and non-rendering algorithms while still maintaining high performance. Examine these programmable components, and the flexibility they provide, through a guided tour of this exciting new engine.

Key Topic	General Interest, Visualization
Summit	NVIDIA Research Summit
Speaker	Steven Parker (Principal Research Scientist, NVIDIA)

THURSDAY 2:00 PM – 3:00 PM
GARDEN ROOM

1078 Adobe Photoshop CS4 and GPUs

Adobe Photoshop CS4 and Adobe Bridge CS4 take advantage of the processing power inherent in GPUs. When combined with GPU power, large Photoshop files can be manipulated faster and smoother than ever before. In this session, you'll see how Photoshop excels when used with NVIDIA GeForce and Quadro GPUs.

Key Topics	Computational Imaging, Film, Medical Imaging & Visualization, Visualization
Summit	GPU Developers Summit
Speakers	Jerry Harris (Sr Computer Scientist II, Adobe Systems); John Nack (Photoshop Product Mgr, Adobe Systems)

THURSDAY 2:00 PM – 3:00 PM
PIEDMONT ROOM

1098 NEXUS Lab

Experience the future of GPU Computing development today! NEXUS is NVIDIA's new GPU Computing development environment for Microsoft Visual Studio. In this hands-on lab, attendees will learn to use the NEXUS source debugger and profiler to efficiently solve common problems encountered when debugging and profiling CUDA C code, including

> Debugging logic errors using hardware breakpoints and variable inspection

> Reading and visualizing GPU memory directly in the Visual Studio IDE (no manual readbacks required!)

> Real-time expression evaluation on the GPU while your application is halted at a breakpoint

> Hardware-evaluated breakpoint conditionals for debugging a specific block or thread

> Detecting out-of-bounds memory errors

> Finding erroneous memory overwrites using data breakpoints

> Understanding the performance of your whole system using the Nexus Analysis tools

NVIDIA engineers will be available at each session to answer your NEXUS and general GPU Computing questions. Lab workstations are limited to the first 20 attendees per session. NVIDIA would like to thank HP and Parallels for their generous support of SLI Multi-OS workstations.

Key Topics	General Interest, Tools
Summit	GPU Developers Summit
Speakers	Brian Kretzler (NVIDIA); Greg Smith (Manager of Developer Tools, NVIDIA)

THURSDAY 2:00 PM – 3:00 PM
CRYSTAL ROOM

1443 Thrust: A Parallel Template Library for CUDA

Thrust is a parallel template library for developing CUDA applications. Modeled after the C++ Standard Template Library (STL), Thrust brings a familiar abstraction layer to the realm of GPU computing. Thrust provides host and device variants of the STL vector container to simplify memory management and facilitate data transfers. These containers are complemented with a large collection of generic data-parallel algorithms and a suite of useful iterator adaptors. Together, these features form a flexible high-level interface for GPU programming that greatly enhances developer productivity.

In this tutorial we'll walk through Thrust's features and demonstrate how Thrust is used to develop larger-scale libraries and applications. We'll also discuss how Thrust enables "kernel fusion" and circumvents the "array of structs" vs. "structure of arrays" problem. Lastly, we'll present evidence that Thrust implementations are fast, while remaining concise and readable.

Key Topic	Tools
Summits	GPU Developers Summit, NVIDIA Research Summit
Speaker	Nathan Bell (Research Scientist, NVIDIA)

THURSDAY 2:00 PM – 4:00 PM
GOLD ROOM

1029 C for CUDA – Advanced

This presentation covers the major CUDA optimizations. Topics will include: maximizing memory throughput, kernel launch configuration, using shared memory, and improving GPU/CPU interaction. We will also describe how the CUDA Visual Profiler and code instrumentation can be used to assess performance. While C for CUDA is used for illustration, the concepts covered

will apply equally to programs written with OpenCL and DirectCompute APIs.

Key Topic	Tools
Summit	GPU Developers Summit
Speaker	Paulius Micikevicius (NVIDIA)

THURSDAY 2:30 PM – 3:00 PM
ATHERTON ROOM

1203 Using GPUs to Estimate the Value-at-Risk of Portfolios

Financial institutions measure market risk by estimating potential losses in the market value of their traded assets. The estimation involves a simulation-based Value-at-Risk (VaR) computation over certain time horizon and provides a confidence interval to guide asset portfolio management decision making. This type of analysis remains one of the most ubiquitous techniques for measuring risk despite its high computational complexity. In this note we describe the implementation of a VaR Monte-Carlo simulation on GPUs. By grouping random experiments into blocks, we are able to exploit the computational efficiency of the CUDA BLAS3 SGEMM routine to significant effect. Using a portfolio of approximately 5000 instruments, we benchmark the performance of our CUDA implementation of the VaR simulation on a Nvidia GeForce GTX 280 against a 2.5 GHz quad-core Intel CPU. Simulation time is 6 times faster using the GPU than the CPU, reducing the simulation time to around 15 minutes. Further optimizations and performance diagnostics are presented.

Key Topics	Algorithms & Numerical Techniques, Computational Finance
Summit	GPU Developers Summit
Speakers	Jike Chong (Ph.D. Candidate, University of California, Berkeley); Matthew Dixon (Postdoctoral Fellow, UC Davis)

THURSDAY 2:30 PM – 3:00 PM
SACRAMENTO ROOM

1415 Efficient Integration of CUDA & OpenGL Proccessing in XIP

This case study presents the results of an effort to implement generic support for CUDA kernels to be executed using the eXtensible Imaging Platform, XIP – an Open Source project funded by the National Cancer Institute. Processing pipelines are constructed using a visual programming paradigm, through graphical representations of pipelines and scenegraphs based on Open Inventor. The developed framework enables efficient integration of CUDA and OpenGL processing by utilizing memory management exploiting PBO direct transfers between CUDA and OpenGL. The framework has been used to execute existing CUDA kernels from NVIDIA's CUDA SDK as well as a more complex image segmentation algorithm. Performance results and breakdown of the framework will be presented as well as examples using the visual programming interface.

Key Topic	Medical Imaging & Visualization
Summit	GPU Developers Summit
Speaker	Patric Ljung (Research Scientist, Siemens Corporate Research)

THURSDAY 2:30 PM – 3:00 PM
VALLEY ROOM

1433 GPU Accelerated Computational Biophysics in NAMD

This talk will present detailed algorithmic approaches to approximating long-range electrostatics on GPUs and optimizing performance for molecular dynamics, the simulation of molecules. GPU implementations of these algorithms differ significantly from CPU implementations both in terms of the relative costs of computation versus memory access and in maximizing data parallelization. To this end, we will describe the Generalized Born method as it was implemented both for Folding@Home and for AMBER, as well as the implementation of the Particle Mesh Ewald (PME) method as implemented for AMBER, focusing on their implementations, and describing the tricks and tradeoffs required to achieve both accuracy and performance. Finally, we will describe how to extend these algorithms to multiple GPUs.

Key Topics	Life Sciences, Molecular Dynamics
Summit	GPU Developers Summit
Speaker	Scott Le Grand (Senior Engineer, NVIDIA)

THURSDAY 2:30 PM – 3:00 PM
GLEN ELLEN ROOM

2004 TechniScan, Inc. Startup Presentation

The Company's Whole Breast Ultrasound (WBU) system produces three separate images; two images from sound transmission properties (one for the speed of sound and one for the attenuation of sound) and a more standard reflection image, is generated from data collected from three reflection transducers with different focal depths. All three images are produced as a series of planar slices of the breast that are reconstructed into a coherent three-dimensional CT image. Images produced by the WBU system possess important advantages over traditional ultrasound methods including: standardized, high-quality images that are generated independent of operator skill;

> images that are accurately registered in 3-D;

> images that are of high spatial resolution and low noise that allow for accurate localization of lesions and calculation of mass size; and

> images that show either coronal slice of the total breast or a complete 3-D representation of the section imaged, rather than a small flashlight view of the breast as shown by traditional ultrasound.

Importantly, the WBU system offers a painless and efficient method, without ionizing radiation, to screen and diagnose the breast and provide unique information about bulk tissue properties that may aid in the effective diagnosis of breast cancer.

Key Topic	Medical Imaging & Visualization
Summit	Emerging Companies Summit
Speaker	David Robinson (President & CEO, TechniScan, Inc.)

THURSDAY 2:30 PM – 3:00 PM
CUPERTINO ROOM

2025 Visualization Sciences Group, Inc

The Visualization Sciences Group ("VSG"), is the leading provider of a scalable application framework for 3D visualizing and understanding scientific and industrial complex and calculated datasets. VSG delivers 3D visualization solutions and services to developers through its Open Inventor toolkit product line and its Avizo end-user application to manage and visualize complex information, large quantities of data, and multi-modality datasets. Open Inventor provides cross-market solutions for the development of 3D interactive professional and commercial applications with state of the art data management and visualization. Avizo leverages all of the Open Inventor suite IP and delivers high-end visualization solutions to end-users in a number of vertical markets. VSG technology improves time to market, integrates and processes data from diverse sources, and presents visual analytics that enable users to understand their data. The technology continues to be applied to new applications across markets as the trend expands to acquire, integrate, visualize, process, analyze and present datasets in 3D.

Key Topic	3D
Summit	Emerging Companies Summit
Speaker	Jean Bernard CAZEAUX (CEO, Visualization Sciences Group, Inc)

THURSDAY 2:30 PM – 3:00 PM
FAIRFIELD ROOM

2032 Milabra Startup Presentation

Milabra is a visual recognition company that focuses on improving online advertising performance through the analysis and active utilization of commercially relevant visual data. Milabra uses advanced neural networks to identify the commercially valuable visual aspects of images, video and graphics, uses predictive analysis to synthesize the data collected, and provides partners with actionable data to advertising performance online. With over half of the Internet graphical in nature, and with over 8 billion photos and 500 million videos uploaded per month to social media in the United States alone, visual media is the most important and fastest growing component of the internet. The high cost of serving and storing this media cannot be offset with advertising dollars unless the contents, impact and meaning of that content is known. Milabra decodes that meaning and turns these visual liabilities into visible revenue. Milabra's Image Intelligence platform and commercial decisioning framework is fast, accurate and easy to use. Our sophisticated neural network technology has the ability to add value to every image and visual element we see.

Founded in 2008, the company was formed by a team of experienced entrepreneurs, industry veterans and scientists from the best visual recognition programs in the country. The company fuses neuroscience with proven advertising, psychology and grapic theory.

Key Topics	General Interest, Visualization
Summit	Emerging Companies Summit
Speaker	Samuel Cox (CEO, Milabra)

THURSDAY 2:30 PM – 3:00 PM
BELVEDERE ROOM

2043 Empulse GmbH Startup Presentation

Empulse GmbH is a Cologne/Germany based IT company focussing on innovative technology and Web applications. The Empulse Indexing Technology (EIT) is a framework for analyzing and searching large data sets on GPGPUs. This framework was developed for the implementation of a tourism web portal, accessing up to 5 billion travel offerings. Upcoming implementations of the EIT include OLAP and Data Warehouse style applications for telco and financial services industries.

Key Topic	General Interest
Summit	Emerging Companies Summit
Speaker	Joerg Bienert (CEO, Empulse GmbH)

THURSDAY 3:00 PM to 3:30 PM
VALLEY ROOM

1053 GPU Accelerated Visualization and Analysis in VMD

State-of-the-art graphics processing units (GPUs) contain hundreds of processing units and are able to perform trillions of floating point arithmetic operations per second. The newly available computational power brought by GPUs is enabling a new generation of scientific and engineering applications to perform calculations on "personal supercomputers" that previously required HPC clusters or that were otherwise impractical in everyday use. This talk will present recent successes in multi-GPU acceleration in VMD, a molecular dynamics visualization and analysis application in which GPU computing techniques have provided speedups ranging from 10 to over 100 times faster than commodity CPU cores. The talk will describe key challenges and algorithmic strategies involved in achieving high computational performance on GPUs, discuss methods for effectively using multiple GPUs in low-latency calculations that drive interactive visualizations, and will also include some examples of how these performance increases ultimately enable better science.

Key Topics	3D, Life Sciences, Visualization, Molecular Dynamics
Summits	GPU Developers Summit, NVIDIA Research Summit
Speaker	John Stone (Senior Research Programmer, University of Illinois at Urbana-Champaign)

THURSDAY 3:00 PM – 3:30 PM
SACRAMENTO ROOM

1445 GPU-Accelerated Imaging on Mobile X-ray C-arm

Digital X-ray on a mobile C-arm device is an imaging tool for assisting the surgeon during image-guided interventions. The current generation of C-arm systems comprises various imaging modes: traditional, real-time high-resolution fluoroscopy, and two new, near-time C-arm CT and tomosynthesis. To tune fluoroscopic IP chain, improve effectiveness of observer studies, and support research efforts, we discuss HW and SW implementation and review the early results associated with GE Healthcare-Surgery Interactive Imaging Chain Simulator. The Simulator has image processing functionality close to one that can be observed on industrial C-arms equipment. Using numerical X-ray CT Exerciser and experimental data from the C-arm prototype capable of rotational scan, we model cone-beam CT performance

on a mobile C-arm, benchmark GPU-accelerated image reconstruction, and demonstrate the imaging results. Our methodology is to create a friendly scalable framework utilizing a common workstation and running surgical applications on a cost-effective heterogeneous computing and visualization platform.

Key Topics	Medical Imaging & Visualization, Visualization
Summit	GPU Developers Summit
Speaker	Arvi Cheryauka (Sr Sys Eng, GE Healthcare – Surgery)

THURSDAY 3:00 PM – 3:30 PM
BELVEDERE ROOM

2003 VERTICE Startup Presentation

Founded in 2002, Vertice is located in Toulouse. Vertice is a Young Innovating Company (JEI) which was elected in 2003 1st Innovating Company. Vertice develops and distributes the Nova sotware range dedicated to 3D real time and accompanies its customers in the customization of applications. Vertice is a sister company of Bewise, one of the rare French software firms exclusively specialized on Microsoft development technologies.

Key Topic	3D
Summit	Emerging Companies Summit
Speaker	David Catuhe (CTO, VERTICE)

THURSDAY 3:00 PM – 3:30 PM
CUPERTINO ROOM

2012 Innovative Converged Devices Ltd Startup Presentation

Innovative Converged Devices Ltd (ICD) is a small, highly efficient company that designs consumer electronic products such as mobile phones, personal computers and new types of Internet devices. ICD achieves differentiation from the very large Far Eastern companies by creating products that are optimised for purpose with aesthetic design that appeals to Western tastes combined with the latest hardware & software technology available. ICD's business model is to conceive, design, prototype and prepare products for manufacture. ICD does not actually perform the manufacturing but arranges introductions between our customers and high-grade contract manufacturers who form their own commercial relationship. ICD facilitates this arrangement but does not participate in the financial transactions of manufacture. ICD describe this model as a Fabless Original Design Manufacturer or Fabless ODM.

Key Topics	Computer Vision, Embedded & Mobile
Summit	Emerging Companies Summit
Speaker	David Hayes (CEO, Innovative Converged Devices)

THURSDAY 3:00 PM – 3:30 PM
FAIRFIELD ROOM

2021 Immersive Media Startup Presentaiton

Immersive Media Corp. (IMC) is the world leader in the technology and production of full motion, high-resolution, 360-degree video experiences for websites. Well known as the developer of Google Street View visual mapping, IMC's fully immersive spherical videos allow viewers complete control to explore environments with a simple click and scroll of the mouse. Viewers can plunge themselves into live 360-degree streaming video on location or full motion prerecorded video/ webisodes. Among its many applications, IMC's immersive video is utilized by major consumer product brands to increase customer loyalty, website interest, and viral campaigns. Tourism and travel industries capitalize on immersive video's unique ability to put viewers directly into a scene. TV and documentary production companies enhance their storytelling and viewer experiences far beyond the traditional narrow viewpoint of a single camera. Noted IMC clients include Mercedes-Benz, Nike, National Geographic, NBC Sports, the NBA, Red Bull, Adidas, Armani and the Emmys telecast. IMC's immersive videos provide unparalleled visual intelligence in security applications for the military, fire and police departments and state and local government agencies. Whether it's on a school bus, in an office building, down an alley, or deep behind the lines in a military theater, immersive videos protect people, manage assets and ensure safety. Its pioneering work in developing Google Street View demonstrates how IMC's video.

Key Topics	General Interest, Visualization
Summit	Emerging Companies Summit
Speaker	Myles McGovern (CEO, Immersive Media)

THURSDAY 3:00 PM – 3:30 PM
GLEN ELLEN ROOM

2050 Total Immersion Startup Presentation

Founded in 1999 by Bruno Uzzan and Valentin Lefevre, Total Immersion is based in France and maintains a U.S. presence in Los Angeles. Its solutions are also available through a network of resellers worldwide. Total Immersion designs and produces innovative, interactive exhibits, serving diverse business segments on a global basis, including theme parks and attractions, special events, retail, and other digital marketing environments. A versatile "platform for discovery" with an undeniable wow factor, Total Immersion offers consumers a compelling way to interact with brands in their own environment. With augmented reality, the brand temporarily "resides" in the viewer's space. Imagine a favorite animated character sitting in the next chair, or a static product suddenly "come to life"– that's Total Immersion's augmented reality. Total Immersion brings to market sophisticated yet practical applications of industry protocols and standards, including the most advanced proprietary recognition, tracking and rendering available from any source. D'Fusion is the leading marketable technology that injects a brand or product representation into a live video stream in real time; the interactive experience is not confined by the physical space of the screen or the display. Operating on standard PC devices and compatible with HD, D'Fusion is highly interactive, practical, educational and entertaining.

Key TopicS	General Interest, Visualization
Summit	Emerging Companies Summit
Speaker	Bruno Uzzan (CEO, Total Immersion)

THURSDAY 3:00 PM – 4:00 PM
ATHERTON ROOM

1005 OPLib: A GPL Library of Elementary Pricing Functions in CUDA/OpenCL and OpenMP

The talk includes a crash course in Derivative Pricing Theory to explain how to reduce nearly all pricing tasks to this handful of routines, spanning the full spectrum of asset classes and model

types. High performance CUDA implementations for the key routines are given. As an application, we go through the code of a full-fledged scenario generation application for financial derivatives, also distributed as a C# project under GPL.

Key Topics	Computational Finance, Tools
Summit	GPU Developers Summit
Speaker	Claudio Albanese (Prof, Kings College)

THURSDAY 3:00 PM – 4:00 PM
CALIFORNIA ROOM

1010 Volunteer Computing for GPUs: Petaflops for Free

Volunteer computing is a paradigm for high-throughput computing in which computer owners donate computing resources to science projects. About 500,000 people currently participate, donating the use of about 800,000 computers and video game consoles. Together these devices provide over 6 PetaFLOPS of computing power to over 40 science projects, whose research areas include climate research, astronomy, high-energy physics, epidemiology, molecular biology, biomedicine, quantum computing, nanotechnology, and environmental studies. This session will cover how scientists with applications for NVIDIA GPUs can use BOINC to quickly and economically gain access to large numbers of GPUs.

Key Topics	Algorithms & Numerical Techniques, Astro GPU, General Interest
Summits	GPU Developers Summit, NVIDIA Research Summit
Speaker	David Anderson (Research Scientist, UC Berkeley)

THURSDAY 3:00 PM – 4:00 PM
CRYSTAL ROOM

1015 DirectCompute

This talk will describe the features and advantages of DirectCompute, give performance advice for optimizing DirectCompute Shaders on NVIDIA hardware, and discuss the various applications of DirectCompute.

Key Topic	Tools
Summit	GPU Developers Summit
Speaker	Tianyun Ni (3D Graphics Engineer, NVIDIA)

THURSDAY 3:00 PM – 4:00 PM
EMPIRE ROOM

1020 Advances in GPU-based Image Processing and Computer Vision

As a massively parallel processor, the GPU is well-suited for performing 'per-pixel' operations in image processing and computer vision. However, new changes in hardware, software, and algorithm mappings allow entire vision algorithms to be performed solely on GPU. In this session, we'll discuss how GPU programming now goes beyond per-pixel mappings, and is providing speedups in image feature processing and handling, frequency domain processing, graph cuts segmentation, and more.

Key Topics	Computational Imaging, Computer Vision
Summit	GPU Developers Summit
Speaker	James Fung (Developer Technology, NVIDIA)

THURSDAY 3:00 PM – 4:00 PM
GARDEN ROOM

1024 Par4All: Auto-Parallelizing C and Fortran for the CUDA Architecture

Par4All in CUDA is a new tool to translate C and Fortran code to CUDA to help programmers accelerate their codes using GPU computing. It is a tool based on the PIPS source-to-source framework that is developed by Mines ParisTech and others for 20 years and is one component of the new Par4All open-source initiative from HPC Project to promote parallelism. Par4All in CUDA uses an abstract interpretation framework based on linear algebra to compute many interesting properties in the program being analyzed. Region analysis is able to compute the array regions used and produced by any program statement. This information is used to parallelize nested loops, allocate the needed data on the GPU and generate communication primitives between the GPU and the host.

Key Topics	High Performance Computing, Tools
Summit	GPU Developers Summit
Speakers	Christopher Carothers (RPI); Beatrice Creusillet (HPC Project); Serge Guelton (Institut Telecom/ Telecom Bretagne); Francois Irigoin (Mines Paris Tech); Ronan Keryell (CSO, HPC Project)

THURSDAY 3:00 PM – 4:00 PM
HILLSBOROUGH ROOM

1034 You Might Also Like: A Multi-GPU Recommendation System

A recommendation system utilizes customers' known preferences to predict unknown preferences. We propose a multi-GPU algorithm for such recommendation systems. By an experiment, we have shown that our algorithm, when implemented on just two GPUs, runs more than 260 times faster than a single core implementation on a latest CPU.

Key Topic	Databases & Data Mining
Summit	GPU Developers Summit
Speaker	Kimikazu Kato (Researcher, Nihon Unisys, Ltd.)

THURSDAY 3:00 PM – 4:00 PM
PIEDMONT ROOM

1098 NEXUS Lab

Experience the future of GPU Computing development today! NEXUS is NVIDIA's new GPU Computing development environment for Microsoft Visual Studio. In this hands-on lab, attendees will learn to use the NEXUS source debugger and profiler to efficiently solve common problems encountered when debugging and profiling CUDA C code, including

> Debugging logic errors using hardware breakpoints and variable inspection

> Reading and visualizing GPU memory directly in the Visual Studio IDE (no manual readbacks required!)

> Real-time expression evaluation on the GPU while your application is halted at a breakpoint

> Hardware-evaluated breakpoint conditionals for debugging a specific block or thread

> Detecting out-of-bounds memory errors

> Finding erroneous memory overwrites using data breakpoints

> Understanding the performance of your whole system using the Nexus Analysis tools

NVIDIA engineers will be available at each session to answer your NEXUS and general GPU Computing questions. Lab workstations are limited to the first 20 attendees per session. NVIDIA would like to thank HP and Parallels for their generous support of SLI Multi-OS workstations.

Key Topic	General Interest, Tools
Summit	GPU Developers Summit
Speakers	Brian Kretzler (NVIDIA); Greg Smith (Manager of Developer Tools, NVIDIA)

THURSDAY 3:00 PM – 5:00 PM
REGENCY 2 ROOM

1049 Supercomputing Super Session

GPU computing is transforming the extreme high-end realms of supercomputing. NVIDIA Tesla GPUs already power several of the world's sixty fastest supercomputers, and this trend is accelerating. This three-hour "super session" will feature some of the world's premiere supercomputing experts, who will discuss their experience building and deploying GPU-based supercomputing clusters, and present case studies of designing and porting codes for "big iron" GPU supercomputers.

Key Topic	High Performance Computing
Summit	NVIDIA Research Summit
Speakers	Wen-mei Hwu (University of Illinois at Urbana-Champaign); Satoshi Matsuoka (Global Scientific Information and Computing Center (GSIC) of Tokyo Institute of Technology); James Phillips (Senior Research Programmer, University of Illinois at Urbana-Champaign); John Taylor (LEADER, CSIRO COMPUTATIONAL AND SIMULATION SCIENCES, CSIRO); Jeffrey Vetter (Oak Ridge National Laboratory and Georgia Tech)

THURSDAY 3:30 PM – 4:00 PM
SACRAMENTO ROOM

1444 GPU-based Decompression for Medical Imaging Applications

Medical imaging applications such as computed tomography (CT), ultrasound, and magnetic resonance imaging (MRI) capture an exponentially increasing amount of real-world sensor data. Real-time compression of this sensor data reduces bandwidth and storage costs of medical imaging equipment by 5% to 25%. This paper illustrates how real-time compression in field-programmable gate arrays (FPGAs) is being combined with GPU-based decompression and image reconstruction to enable new high-performance, low-cost imaging systems. We compare GPU-based decompression with CPU-based compression and identify ways that future GPUs could further reduce the decompression overhead during image reconstruction. More generally, transferring compressed packets across the GPU-CPU interface solves signal processing I/O problems in CUDA-based compute applications.

Key Topics	Medical Imaging & Visualization, Visualization
Summit	GPU Developers Summit
Speaker	Albert Wegener (CTO, Samplify Systems)

THURSDAY 3:30 PM – 4:00 PM
FAIRFIELD ROOM

2007 Realtime Technology AG Startup Presentation

RTTprovides 3D real-time visualization technology and services for innovative workflows in the auto, aeronautical and consumer-goods industries, as well as in the fields of interior design and architecture. It supports customers in coming up with new ways how to design and to market their products. The company supports its customers not only during the design and developing phases with regard to the virtual display of different product concepts (Virtual Prototyping), but also during the marketing and sales stages in presenting the products in all variants and with all additional options in the Showroom, on the Internet, at fairs or at dealerships (Virtual Marketing).

Key Topic	3D, Visualization
Summit	Emerging Companies Summit
Speaker	Ludwig A. Fuchs (Board Member, Realtime Technology AG)

THURSDAY 3:30 PM – 4:00 PM
CUPERTINO ROOM

2009 OptiTex Startup Presentation

OptiTex is the leading developer of 2D & 3D CAD solutions for virtually all sewn-products industries. OptiTex technologies allows designers to create, correct and adjust compelling designs before the first piece of fabric is cut, giving a new dimension to the motto, "Virtual is Real" and proving what OptiTex has long maintained – that corporate responsibility and an environmentally friendly approach are not mutually exclusive but in fact are intimately entwined. OptiTex presents a complete content creation solution, to create and visualize customized garments, to simulate fitting and draping of garments on fully parametric virtual models and to create movie clips in an immediate and direct manner. OptiTex system consists of three main components: cloth content creation system with our PDS software, 3D Runway Designer, a virtual try-on-system, which includes both cloth simulation and accurate 3D parametric mannequins; motion animation engine which enables the generation of motion sequences with interactive cloth. OptiTex brings a wealth of virtual textile experience to the gaming, feature animation and digital effects industries. OptiTex's products are second only to real life in depicting fabric movement and dynamics. The combination of OptiTex's rich graphics and data intensive computation and NVIDIA's CUDA-enabled GPUs provides a high-performance environment for animators and digital artists via the integration of OptiTex cloth engine technology with the Modeling and Anima

Key Topic	3D
Summit	Emerging Companies Summit
Speaker	Eri Rubin (Head of CUDA Project, OptiTex, LTD.)

THURSDAY 3:30 PM – 4:00 PM
BELVEDERE ROOM

2036 Anark Corporation Startup Presentation

Anark Corporation provides powerful enterprise software and solutions to manufacturing organizations that need to cost-effectively utilize their 3D CAD data and other graphics assets for visualization, simulation, CAE, and supply-chain data exchange. Anark helps market leaders such as Boeing,

Cessna, Pratt & Whitney, Lockheed Martin, and NASA unlock the potential of their 3D data assets to save money, accelerate product development, and to collaborate and communicate more effectively and securely.

Key Topic	3D, Visualization
Summit	Emerging Companies Summit
Speaker	Stephen Collins (CEO, Anark Corp)

THURSDAY 3:30 PM – 4:00 PM
GLEN ELLEN ROOM

2051 Unity Technologies Startup Presentation

Unity enables console-quality games for the desktop, the Web, the iPhone and the Wii console. The Unity tools have been hailed as revolutionary for enabling high quality results in an easily accessible development environment. Unity is the fastest growing platform with at least one new Unity game appearing for iPhone on a daily basis.

Key Topic	General Interest
Summit	Emerging Companies Summit
Speaker	David Helgason (CEO, Unity Technologies)

THURSDAY 4:00 PM – 4:30 PM
ATHERTON ROOM

1004 Derivative Price Modeling on the GPU

Abstract: To reproduce the implicit dynamics of stocks, increasingly complex models are used for pricing and hedging. Because of this complexity, the numerical methods mainly rely on Monte Carlo (MC) simulation. Because these methods are easy to parallelize, GPU based hardware is an excellent candidate for accelerating financial algorithms. This talk will summarize the overall performance that one expects from an effective utilization of GPUs using the CUDA or CG Toolkit on financial problems. This will include pricing American and Exotic European options and some Exotic Fixed-Income options studied with CALYON Bank. Also some additional computational challenges which GPUs could help solve will be discussed.

Key Topic	Computational Finance
Summit	GPU Developers Summit
Speakers	Lokman A. Abbas-Turki (PHD student in Applied Mathematics and Mathematical Finance, Paris-Est UMLV LAMA)

THURSDAY 4:00 PM – 4:30 PM
CALIFORNIA ROOM

1008 Memory Saving Discrete Fourier Transformation on CUDA

The Fourier transformation of large data sets is a computation intensive and memory consuming operation. It can be speeded up by transferring the computation to a graphics card where many operations can be done in parallel. As today even the low class consumer products are CUDA enabled, the transformation can be computed on the graphics device and offload the main processor. Our approach (called FlexFT) intends to make the benefit of parallel computing available on low end hardware running demanding image processing applications. In the talk we will show our algorithm and present the evaluated results in comparison to existing approaches. We will show, that our algorithm performs equal to CUFFT but does not suffer from the memory limitation at all.

Key Topic	Computational Imaging
Summit	GPU Developers Summit, NVIDIA Research Summit
Speakers	Steffen Frey (University of Stuttgart); Harald Sanftmann (University of Stuttgart)

THURSDAY 4:00 PM – 4:30 PM
SACRAMENTO ROOM

1088 Medical Image Registration with CUDA

This session will discuss the exceptionally fast 3D Medical Image Registrations achieved using CUDA. Speedups of up to 750 times were obtained as compared to code in daily use at Addenbrookes Hospital and Bio-Medical Campus. Some very recent results are shown in the figures. This work is of direct application in both research and clinical practice. A particular application is voxel based MRI morphometry in humans and in animal brains. Here small changes in particular regions of the brain are studied to monitor disease progression, response to treatment and effects of genetic variation etc.

Key Topics	Computational Imaging, Medical Imaging & Visualization
Summit	GPU Developers Summit
Speaker	Richard Ansorge (Senior Lecturer, University of Cambridge)

THURSDAY 4:00 PM – 4:30 PM
HILLSBOROUGH ROOM

1417 Using Flash SVG Animations in Enhanced & Virtual Reality

This presentation will cover the use of Adobe Flash (SVG) animations in virtual "worlds". Today's consumer products have dynamic and interactive multimedia user interfaces. Many businesses would benefit by employing these same sophisticated user interfaces in their traditional graphics-based applications. The integration of Flash animations into these primarily process driven applications increases their appeal, usability and functionality.

During the presentation, we will demonstrate the use of Animated Media's VGK product platform with an NVIDIA GPU in an industrial control solution that was developed for a client in the fire and safety business. The client's fire and safety control system was being upgraded to use current technologies and required 2D/3D visualization of real-time activities. In the solution, we virtualized and used multiple VGK engines to render the Flash animations to texture-mapped 3D accelerated OpenGL objects in the floor plans of multi-story buildings. The animations represent various fire system objects (sensor components) in the building such as fire control valves, sprinklers, smoke detectors, speakers, elevators, escape routes, and fire detectors. The building's fire and safety system, accessed through the building's fire panel, communicates in real-time using XML over a network. The application monitors the status of every fire system object in the building and commands the ActionScript to display the proper animation to match the object's status. Integrating standard Adobe Flash animations into a 2D or 3D software application adds attention getting visuals, such as a pulsating speakers and smoke detectors, enabling the building staff to readily identify troubled

areas within the building and direct firemen or safety resources to that exact location.

Key Topic	Computational Imaging
Summit	GPU Developers Summit
Speaker	Chris Brady (CEO, Animated Media); Bob Starr (Animated Media)

THURSDAY 4:00 PM – 4:30 PM
BELVEDERE ROOM

2010 Nurien Software Startup Presentation

Nurien Software, established in 2004, is a leading developer of next-generation 3D social networking platform and gaming services geared to deliver unprecedented experience to the end user. The Nurien platform includes user-centric social networking features a fully customizable 3D avatar system, 3D objects and user-created content support as well as game and other social applications including dance, fashion and quiz game applications. Users can create one-of-a-kind avatars to their exact liking, and use avatars to participate in games, form communities, and interact with other players. With stunning visuals and innovative game and entertainment functions, the Company aims to revolutionize next-generation 3D social network platform market.

Key Topic	3D
Summit	Emerging Companies Summit
Speaker	Taehoon Kim (Co-founder, President, Nurien Software)

THURSDAY 4:00 PM – 4:30 PM
CUPERTINO ROOM

2040 Wild Pockets Startup Presentation

Wild Pockets is the next generation solution for creating and monetizing interactive 3D content embedded right within your web-browser. Whether you are an independent developer working on games in your garage or a large company looking for the fastest, most powerful development environment to create stunning immersive 3D web content – Wild Pockets is your solution. The Wild Pockets platform is disrupting the notion that you have to license and pay big dollars for development software before your project has even begun. Instead, Wild Pockets provides FREE and equal access to all developers, professional, hobbyist and student alike, because you never know where the most innovative new ideas will emerge from! Wild Pockets is designed as a full end to end solution and surrounding eco-system because it is takes more than just creating content to be successful. All Wild Pockets content can be embedded and distributed instantly on any website anywhere online similar to a YouTube video, and the built-in transaction system and marketplace allows developers to monetize their content quickly and easily. Check out the demos and sign up for your free account today at www.wildpockets.com.

Key Topic	3D, Embedded & Mobile
Summit	Emerging Companies Summit
Speaker	Shanna Tellerman (CEO, Wild Pockets)

THURSDAY 4:00 PM – 4:30 PM
GLEN ELLEN ROOM

2044 Enodo Startup Presentation

Abstract: By adapting state of the art video-game technology to industrial project's needs, ENODO became a leading 3D visualization company, providing cutting-edge solutions for real time & interactive manipulation of industrial data.

Key Topics	3D, General Interest
Summit	Emerging Companies Summit
Speaker	laleh-chloe sahrai (International project manager, ENODO)

THURSDAY 4:00 PM – 5:00 PM
CRYSTAL ROOM

1012 MetaSL and the Evolution of Shading Languages

The performance of the GPU as a 3D geometry engine has revolutionized our expectations of rendered image complexity. However, attempts to tap the power of new hardware in 3D rendering have encouraged the design of a variety of complex and incompatible programming languages. MetaSL is a "meta shading language" that defines an abstraction layer above the specifics of any given hardware implementation. Rendering plugins (called "shaders) written in MetaSL are translated to the target shading language of a renderer through "back- end" compilers. The development of these compilers is supported by a software library and SDK that provide the intermediate data structures typical of software compilers, including the construction of an abstract syntax tree. This talk will provide an overview of MetaSL: its design philosophy, practical examples of shader implementation, and the custom development of new back-end compilers for current and future rendering pipelines.

Key Topics	Film, Visualization
Summit	GPU Developers Summit
Speaker	Andy Kopra (Mental Images)

THURSDAY 4:00 PM – 5:00 PM
GARDEN ROOM

1025 Large-Scale Text Mining on the GPU

We are quickly reaching an age in which a capability is needed for knowledge discovery from terabyte-scale unstructured text collection for prompt decision-making. Text mining (TM) is an important field for knowledge discovery in large scale text collection. In terms of speedup, utilizing the GPU for TM processing offers a much more cost-effective choice in comparison with using a high-performance supercomputer. Our research mainly focused on exploring GPU-enhanced solutions of three TM functions, text document encoding, text vector dimension reduction, and document clustering functions. In our experiments, the average speedup ranging from six to eight times for all three TM functions. Our successful results proofed the GPU-enhanced computer's ability in text mining the large-scale text dataset.

Key Topic	Databases & Data Mining
Summit	GPU Developers Summit
Speaker	Xiaohui Cui (Oak Ridge National Laboratory)

THURSDAY 4:00 PM – 5:00 PM
VALLEY ROOM

1055 What Every CUDA Programmer Needs to Know About OpenGL

CUDA has attracted thousands of new developers to GPU computing. Many of these new programmers have pursued CUDA development without any prior exposure to OpenGL or experience with graphics programming, but nonetheless could benefit from the powerful visualization capabilities GPUs also have to offer. This talk is intended to provide an introduction to OpenGL graphics programming specifically oriented towards the non-graphics CUDA programmer. We'll discuss basic OpenGL operations, simple drawing, image/video display and CUDA / OpenGL interoperability concepts for rendering data generated from CUDA programs.

Key Topics	General Interest, Tools
Summits	GPU Developers Summit, NVIDIA Research Summit
Speaker	Joe Stam (NVIDIA)

THURSDAY 4:00 PM – 5:00 PM
PIEDMONT ROOM

1098 NEXUS Lab

Experience the future of GPU Computing development today! NEXUS is NVIDIA's new GPU Computing development environment for Microsoft Visual Studio. In this hands-on lab, attendees will learn to use the NEXUS source debugger and profiler to efficiently solve common problems encountered when debugging and profiling CUDA C code, including

> Debugging logic errors using hardware breakpoints and variable inspection

> Reading and visualizing GPU memory directly in the Visual Studio IDE (no manual readbacks required!)

> Real-time expression evaluation on the GPU while your application is halted at a breakpoint

> Hardware-evaluated breakpoint conditionals for debugging a specific block or thread

> Detecting out-of-bounds memory errors

> Finding erroneous memory overwrites using data breakpoints

> Understanding the performance of your whole system using the Nexus Analysis tools

NVIDIA engineers will be available at each session to answer your NEXUS and general GPU Computing questions. Lab workstations are limited to the first 20 attendees per session. NVIDIA would like to thank HP and Parallels for their generous support of SLI Multi-OS workstations.

Key Topics	General Interest, Tools
Summit	GPU Developers Summit
Speakers	Brian Kretzler (NVIDIA); Greg Smith (Manager of Developer Tools, NVIDIA)

THURSDAY 4:00 PM – 5:00 PM
EMPIRE ROOM

1454 GPU Ventures, for the Entrepreneur

During this panel, we will discuss how Viewdle and NVIDIA have been working together and bring a product to market. Viewdle first came to NVIDIA with a technology that required massive amounts of raw compute. Working with NVIDIA's technical team, Viewdle was able to harness the power of the GPU, and bring advanced photo and video analysis to the consumer desktop. NVIDIA not only provided the technical expertise, but also the business and marketing support. Working together with NVIDIA, Viewdle was able to formulate a consumer strategy and secure distribution partners. This discussion will be from the point of view of the entrepreneur who would like to work with NVIDIA and will include real examples and demos of Viewdle's soon to launch software.

Key Topic	General Interest
Summit	GPU Developers Summit
Speakers	Laurent Gil (CEO and President, Viewdle Inc.); Jeff Herbst (Vice President of Business Development, NVIDIA)

THURSDAY 4:00 PM – 6:00 PM
GOLD ROOM

1030 GPU-based Libraries: Accelerating the Next Generation of Applications

Through continued technological advances, GPUs have become massively-parallel vector processors streamlined for high-performance computations. The increasing flexibility of the GPU has allowed it to enter new markets, and the vendors behind this hardware have increased GPU usability and opened up the underlying hardware constructs to general computing uses. GPUs are now allowing software normally targeted at high-performance computing platforms to be run on desktops and in GPU-enhanced clusters. This creates an opportunity for GPU-accelerated libraries that can be used to increase the performance of numerous applications. The panel will discuss some of the libraries that are currently being developed and answer questions concerning their future and integration into other applications.

Key Topics	High Performance Computing, Tools
Summit	GPU Developers Summit
Speakers	Eric Kelmelis (CEO, EM Photonics); Peter Messmer (VP Space Applications, Tech-X Corporation); Gallagher Pryor (CTO, AccelerEyes); Ryan Schneider (Chief Technology Officer, Acceleware)

THURSDAY 4:30 PM – 5:00 PM
HILLSBOROUGH ROOM

1009 Real Time Multi-Channel, Multi-Camera Image Processing using CUDA

At The Johns Hopkins University Applied Physics Laboratory we are developing and deploying advanced image processing techniques implemented in CUDA and running on off-the-shelf NVIDIA GPUs. Our experience and confidence with CUDA has allowed us to migrate existing image processing techniques from CPU to GPU implementations. The GPU implementations execute with such speed that we are now able to deploy custom-built CUDA implementations that take full advantage of the speed and robustness of the latest CUDA releases. GPUs and the CUDA architecture have allowed us to deploy more sophisticated processing techniques and run multiple techniques simultaneously on additional sensors, multiplying

by many times our capability. We will discuss in presentation format the evolution of our product from CPU to GPU and beyond.

Key Topic	Computational Imaging
Summits	GPU Developers Summit, NVIDIA Research Summit
Speakers	Sean Happel (Research and Development Engineer, The Johns Hopkins Applied Physics Laboratory); Todd Neighoff (System Architect, The Johns Hopkins Applied Physics Laboratory)

THURSDAY 4:30 PM – 5:00 PM
ATHERTON ROOM

1026 Particle Filters and Sequential Monte Carlo Techniques Using GPUs

Sequential Monte Carlo (SMC) techniques, which include "particle filters", have become important tools in a variety of signal processing and estimation applications over the past decade. The basic approach involves using a collection of point masses (particles) to create a non-parametric representation of a probability density of interest. SMC allows non-linear state dynamics and arbitrary (including multi-modal) densities, thus providing an important generalisation of previous approaches such as Kalman filters. Implementation of SMC algorithms on GPUs is attractive due to the fact that particle computations are highly data-parallel. Some operations however, such as weight re-normalisation and some resampling algorithms, are more difficult to parallelise. The proposed session will discuss implementation issues surrounding SMC algorithms, and facilitate sharing and evaluating implementation ideas.

Key Topics	Algorithms & Numerical Techniques, Computational Finance
Summit	GPU Developers Summit
Speaker	W. James MacLean (Senior Engineer, MDA Space Missions)

THURSDAY 4:30 PM – 5:00 PM
SACRAMENTO ROOM

1035 Harnessing the GPU for Surgical Training and Preoperative Planning

Surgical simulation and preoperative planning systems require high performance solutions to complex spatial and physically based problems often over large datasets. We detail the use of NVIDIA CUDA and GPGPU techniques in surgical training and pre-operative planning systems currently under development.

Key Topic	Life Sciences, Visualization, Physics Simulation, Medical Imaging & Visualization
Summit	GPU Developers Summit
Speaker	Josh Passenger (Project Leader, CSIRO)

THURSDAY 4:30 PM – 5:00 PM
CALIFORNIA ROOM

1455 Advanced CUDA Development in MATLAB with Jacket

In this tutorial we present a new approach to integrating CUDA code into MATLAB leveraging Jacket's Developer SDK. Jacket's Developer SDK makes integration of custom CUDA code into Jacket's runtime very easy. With a few simple jkt functions (which mimic standard MEX API functions), you can integrate custom CUDA kernels into Jacket. This enables your CUDA code to inherit Jacket's optimized memory management and kernel execution runtime system. In this tutorial, we share examples of using the Developer SDK to write a median filter in CUDA and integrate it into Jacket. We will start with the naïve approach, then show how to optimize by using shared memory, and finally show the impact of using texture memory on this problem. In each case, we will integrate the code into Jacket & MATLAB via the Developer SDK and do benchmarking in MATLAB. A Q/A session will enable audience members to ask specific questions about the Developer SDK, CUDA programming, and MATLAB-based GPU computing.

Key Topic	Tools
Summit	GPU Developers Summit
Speaker	James Malcolm (Chief Architect, AccelerEyes)

THURSDAY 4:30 PM – 5:00 PM
BELVEDERE ROOM

2006 SOFTKINETIC S.A. Startup Presentation

Softkinetic develops 3D gesture recognition middleware to build natural interfaces for the interactive digital entertainment and consumer markets.

Key Topic	3D
Summit	Emerging Companies Summit
Speaker	Michel TOMBROFF (CEO, SOFTKINETIC S.A.)

THURSDAY 4:30 PM – 5:00 PM
FAIRFIELD ROOM

2030 MirriAd Startup Presentation

MirriAd is the only end to end digital marketing solution that creates a new way to reach audiences by embedding brands into video content. MirriAd offers content owners and advertisers a complete marketplace with MirriAd's online services platform, a full on-line management system. Content owners can upload their content securely and automatically to the highly secure, scalable platform and offer the inventory for sale. Advertisers can view a catalog of video content files and search for the video or set of videos that best meets their campaign needs. Using MirriAd an advertiser may select their preferred zones, insert their brand images and preview a low quality stream of the embedded advertising video before they commit and execute a buy.Videos can be ingested individually or in volume with various degrees of automation. By managing engagements and campaigns online, content owners can exert as much editorial control over the process as they need or as required by advertising standards or regulatory guidelines. Brand images and video can be embedded in professionally produced content such as DVD, TV and online and even user generated content. MirriAd's unique campaign management platform manages multiple stakeholders securely, allowing each stakeholder the appropriate level of visibility of the process: before the placement, and at final sign-off stage. The platform also delivers sophisticated metrics detailing the effectiveness of any particular campaign on any platform.

Key Topics	Computational Imaging, General Interest
Summit	Emerging Companies Summit
Speaker	Mark Popkiewicz (CEO, MirriAd)

THURSDAY 4:30 PM – 5:00 PM
CUPERTINO ROOM

(2033) ICIDO GmbH Startup Presentation

ICIDO, pioneer and global leader of Virtual Reality and Virtual Engineering Solutions support companies to make their product development process sustainable, competitive and environmentally friendly. ICIDO's virtual engineering solutions minimize risks, reduce failure costs and decrease the use of resources. They also allow for faster and easier global collaboration and communication, ensure better visibility and visualization of early stage products. At the end they integrate customers, partners and other departments in the product lifecycle early. ICIDO's Virtual Reality solutions are used throughout the entire product development process from design to manufacturing. ICIDO's solutions also generates value in sales, marketing and maintenance.

Key Topics	General Interest, Visualization
Summit	Emerging Companies Summit
Speaker	Ralf Heimberg (CEO, ICIDO GmbH)

THURSDAY 4:30 PM – 5:00 PM
GLEN ELLEN ROOM

2053 Xulu Entertainment Startup Presentation

Xulu is developing a new type of interactive experience based on massively distributed physics simulation that delivers unprecedented realism in a virtual world. Capable of supporting up close and personal social interactions, alongside full high-end gaming and sporting experiences, the system provides open API's and a suite of tools that allow both professional and novice users alike to contribute a broad range of highest quality content. The Xulu platform scales gracefully for a large number of simultaneous users with minimal reliance on centralized servers, and harnesses the power of leading-edge GPUs and CPUs to deliver ultimate online end-user experiences. Compelling content is a key element of a successful launch, and is needed to attract a large initial audience that will drive demand for Xulu's platform and generate early revenue. Xulu's content will showcase the beginnings of an expanding universe with both real-world and fantasy-world elements. It will include a range of seed activities from extreme sports and gaming to gentle self-paced sightseeing excursions, episodic adventures, face-to-face socializing and play, and support for vibrant and viable communities. Third party and user-created entertainment, businesses, goods, and services will follow. Xulu is leveraging more than 100 man-years effort in a prior related project that included the development of the first generation of the technology, virtual destinations, and multi-player games.

Key Topics	General Interest, Visualization
Summit	Emerging Companies Summit
Speaker	Nanci Solomon (CEO, Xulu Entertainment)

THURSDAY 5:00 PM – 5:30 PM
ATHERTON ROOM

1027 Banking on Monte Carlo

Monte Carlo simulation is a common technique used to find solutions to a wide range of problems where it is impossible or impractical to compute a result with a deterministic method. Monte Carlo relies on repeated computation of random numbers and this can be accomplished extremely efficiently on GPUs. In this presentation we give an example from the Finance community of extreme performance gains using the random number generators (RNGs) in the NAG GPU Library. We go on to discuss why techniques such as Monte Carlo might play an increasing role in advancing science during the current step change in hardware technology.

Key Topics	Algorithms & Numerical Techniques, Computational Finance
Summit	GPU Developers Summit
Speaker	Ian Reid (Numerical Algorithms Group (NAG))

THURSDAY 5:00 PM – 5:30 PM
SACRAMENTO ROOM

1036 GPU Accelerated Solvers for ODEs Describing Cardiac Membrane Equations

Mathematical models describing cellular membranes form the basis of whole tissue models to describe the electrical activity of entire organs, such as the heart. Numerical simulations based on these models are useful for both basic science and increasingly for clinical diagnostic and therapeutic applications such as targeting ablation therapy for atrial arrhythmias, defibrillator design and cardiac resynchronization therapy. A common bottleneck in such simulations arises from solving large stiff systems of ordinary differential equations (ODEs) thousands of times for numerous integration points (representing cells) throughout a three-dimensional tissue or organ model. For some electrophysiology simulations, over 80% of the time is spent solving these systems of ODEs. While a cluster provides the required interactive response time to solve the ODEs, a desktop sized platform would enhance usability of the software in a laboratory setting. The audience will benefit by learning how a real-world, complex, HPC application can directly benefit by the use of CUDA technology. Participants will learn which optimization techniques yielded the best performance results on an actual application. We will also explore the benefits and limits of the use of single precision in certain scientific applications.

Key Topics	Algorithms & Numerical Techniques, Life Sciences, Medical Imaging & Visualization, Visualization
Summit	GPU Developers Summit
Speaker	Fred Lionetti (Researcher, University of California, San Diego)

THURSDAY 5:00 PM – 5:30 PM
HILLSBOROUGH ROOM

1037 Mapping Satellite Imagery on the GPU: Fast Orthorectification and Pan-Sharpening

In the mapping and remote sensing industries, the ever-increasing availability of high-resolution satellite and aerial imagery in recent years has had two main effects. First, the data volume required for a given project size (for example, a county or state) has increased due to improved spatial resolution of the data. Second, the improved availability of coverage is leading geospatial data producers to handle larger and larger projects. To better serve our customers in these industries, we have improved the throughput in our core image processing modules by taking advantage of the high performance offered by NVIDIA GPUs. Numerous software modules in our new ProLines GeoImaging Server product are being reengineered

to support this technology, but the first two completed modules (orthorectification and pansharpening) have already demonstrated a significant performance benefit (65x and 14x over conventional methods).

Key Topic	Computational Imaging
Summit	GPU Developers Summit
Speakers	Teodor C. Hanchevici (Technical Project Leader, PCI Geomatics); James Lutes (Software Developer (Imaging Geometry), PCI Geomatics)

THURSDAY 5:00 PM – 5:30 PM
GARDEN ROOM

1434 Unlocking Biologically-Inspired Computer Vision: a High-Throughput Approach

The study of biological vision and the creation of artificial vision systems are naturally intertwined – exploration of the neuronal substrates of visual processing provides clues and inspiration for artificial systems, and artificial systems, in turn, serve as important generators of new ideas and working hypotheses. However, while systems neuroscience has provided inspiration for some of the "broad-stroke" properties of the visual system, much is still unknown. Even for those qualitative properties that most biologically-inspired models share, experimental data currently provide little constraint on their key parameters. Consequently, it is difficult to truly evaluate a set of computational ideas, since the performance of a model depends strongly on its particular instantiation – the size of the pooling kernels, the number of units per layer, exponents in normalization operations, etc. To pave a way forward, we have developed a high-throughput approach to more expansively explore the possible range of biologically-inspired models, including models of larger, more realistic scale, leveraging recent advances in commodity stream processing hardware – particularly, high-end NVIDIA GPUs. In analogy to high-throughput screening approaches in molecular biology and genetics, we generated and trained thousands of potential network architectures and parameter instantiations, and "screened" the visual representations produced by these models using an object recognition task. From these candidate models, the most promising were selected for further analysis. We have shown that this approach can yield significant, reproducible gains in performance across an array of basic object recognition tasks, consistently outperforming a variety of state-of-the-art purpose-built vision systems from the literature, and that it can offer insight into which computational ideas are most important for achieving this performance. In this talk, I'll also highlight how the application of flexible programming tools, such as high-level scripting and template metaprogramming, can enable large performance gains, while managing complexity for the developer. As the scale of available computational power continues to expand, our approach holds great potential both for accelerating progress in artificial vision, and for generating new, experimentally-testable hypotheses for the study of biological vision.

Key Topic	Computer Vision
Summits	GPU Developers Summit, NVIDIA Research Summit
Speaker	Nicolas Pinto (MIT)

THURSDAY 5:00 PM – 5:30 PM
GOLD ROOM

1457 Examples of CUDA-based Realizations in Artificial Vision and Soft Computing

We present CUDA-based realizations of computationally expensive algorithms in artificial vision as well as soft computing, running at ultra-fast rates. One example addresses the problem of agglomerative clustering, exploiting parallel realizations to achieve very high performance gains. Another example highlights a new, biologically-inspired approach to the implementation of deep artificial neural networks in CUDA. To illustrate this approach, visualizations of deeper layer receptive fields will be presented, as well as a discussion regarding future implementations to improve both network accuracy, as well as its overall information capacity. Network training in CUDA is also presented. Benchmarking of these realizations against those with OpenCV/Intel platforms will also be done, as well as highlighting both GPU occupancy and performance with the Nexus profiler.

Key Topic	Computer Vision
Summit	GPU Developers Summit
Speaker	Tarek El Dokor (CEO and CTO, Edge 3 Technologies)

THURSDAY 5:00 PM – 5:30 PM
BELVEDERE ROOM

2020 Acceleware Startup Presentation

Acceleware provides software solutions to harness the parallel processing capabilities of multi-core GPUs/CPUs for the Electronic Design and Oil & Gas industries. Our Acceleration platform seamlessly integrates with applications from industry leading vendors in the Electromagnetic Simulation, Seismic Data Processing, Reservoir Modeling industries and Linear Algebra solvers. Acceleware solutions enable users of single-threaded applications to access multi-core processing hardware and achieve dramatic compute speed-ups. With Acceleware solutions installed, run times for data processing and simulation applications are reduced by more than 50 times. Our customers share a common and urgent need – the need for powerful and timely computer modeling and testing. They are pushing the boundaries of innovation and demand faster product-development cycles for more complex products. They want more effective tools to interpret vast amounts of data.

Key Topic	General Interest
Summit	Emerging Companies Summit
Speaker	Michal Okoniewski (Interim President & CEO, Acceleware)

THURSDAY 5:00 PM – 5:30 PM
FAIRFIELD ROOM

2024 GASS Company for Advanced Supercomputing Solutions LTD Startup Presentation

> Founded in 2008, after years of experience in Supercomputing and visualizations

> Privately owned

> Developing Hoopoe – a cloud infrastructure and service for GPU computing

> Providing libraries to assist developers with using GPU

technologies from .NET and Java applications (CUDA.NET, jCUDA, OpenCL.NET)

> Develops GPU based solutions for the Israeli industry targeting general computing aspects and visualization

Key Topics	General Interest, Visualization
Summit	Emerging Companies Summit
Speaker	Mordechai Butrashvily (CEO, CTO, GASS / Hoopoe Cloud)

THURSDAY 5:00 PM – 5:30 PM
GLEN ELLEN ROOM

2031 Mirics Semiconductor Inc. Startup Presentation

Mirics provides total system solutions for converged wireless connectivity on portable consumer electronic devices such as Notebook PCs, Media Players and Cellphones. The initial focus of the company is the delivery of free-to-air broadcast services and content to portable computing and communications devices. The Mirics FlexiTV platform leverages re-configurable RF tuner hardware combined with a host processor based software demodulator. The solution is capable of covering all broadcast standards in all terrestrial broadcast bands from simple AM radio at LW to DVB-T in UHF Band V. Currently focusing primarily on providing global TV and radio reception on PC platforms, FlexiTV offers substantial benefits in terms of system flexibility and solution cost and footprint when compared to conventional hardware based solutions. Current solution offerings with the FlexiTV platform include: AM/FM Radio, DAB Radio along with ISDB-T, DVB-T and T-DMB for TV applications. Mirics has been collaborating closely with NVIDIA to migrate increasing amounts of the FlexiTV software demodulator from the x86 CPU platform to the CUDA GPU platform. In doing so, Mirics and NVIDIA are demonstrating the long term viability of software based modems whilst at the same time illustrating the efficiency of the CUDA GPU as a signal processing platform for non-graphics applications.

Key Topics	Computer Vision, General Interest, Embedded & Mobile
Summit	Emerging Companies Summit
Speaker	Simon Atkinson (Chief Executive Officer, Mirics Semiconductor Inc)

THURSDAY 5:00 PM – 5:30 PM
CUPERTINO ROOM

2052 Visuvi, Inc. Startup Presentation

Visuvi develops targeted visual search engine solutions for a wide range of vertical applications in medicine, ecommerce and general-purpose visual search and maintains an index of images on the Internet. Visuvi has patent pending technology fort it's search engine that examines the content and patterns within an image, categorizes that information via mathematical indexing and delivers search results based on the image itself – no text or meta-tags required. Visuvi Inc. is a privately held company based in Redwood City, CA and is managed by a seasoned executive team consisting of Christopher Boone, President and CEO, Alexander Valenica, Chief Scientist and Co-Founder, Florian Brody, VP Marketing and Yuri Drozd, VP Product Management.

Key Topics	General Interest, Visualization
Summit	Emerging Companies Summit
Speaker	Christopher Boone (President & CEO, Visuvi Inc)

THURSDAY 5:00 PM – 6:00 PM
CALIFORNIA ROOM

1016 MATLAB GPU Computing Essentials Tutorial

In this tutorial, we will discuss AccelerEyes' Jacket software which connects MATLAB to the graphics processing unit (GPU). With the GPU as a backend computation engine, Jacket brings together the best of three important computational worlds: computational speed, visualization, and the user-friendliness of MATLAB programming. Jacket enables developers to write and run code on the GPU in the native M-Language used in MATLAB. Jacket accomplishes this by automatically wrapping the M-Language into a GPU compatible form. By simply casting input data to Jacket's GPU data structure, MATLAB functions are transformed into GPU functions. Jacket also preserves the interpretive nature of the M-Language by providing real-time, transparent access to the GPU compiler. The tutorial will provide examples of running MATLAB code on the GPU for image and signal processing, life science, finance, and other applications. Also, a Q/A session will enable audience members to ask specific questions about the Jacket project and MATLAB GPU computing.

Key Topic	Tools
Summit	GPU Developers Summit
Speaker	John Melonakos (CEO, AccelerEyes)

THURSDAY 5:00 PM – 6:00 PM
CRYSTAL ROOM

1017 The Art of Debugging for the CUDA Architecture

This session is an open discussion on CUDA debugging tips and discussion of existing debugger features in the CUDA architecture. Members of this panel have several years of experience developing various debuggers and will share their experience building the debugger for the CUDA architecture. This is excellent opportunity for existing and future customers and developers to come together to share their questions and comments on how to make CUDA debugger better and how to take CUDA debugging into every developer tool.

Key Topic	General Interest, Tools
Summit	GPU Developers Summit
Speakers	Satish Salian (CUDA Tools Manager, NVIDIA), David Lecomber (CTO Allinea Software, DDT debugger); Ed Hinkel (Senior Engineer, TotalView debugger, TotalView Technologies); Avinash Baliga (Senior Engineer, Windows Nexus Debugger, NVIDIA)

THURSDAY 5:00 PM – 6:00 PM
EMPIRE ROOM

1022 High Performance Remote Graphics

GPU computing is rapidly becoming a key technology for many industries. When associated to the latest virtualization technology, graphics intensive applications (CAD, DCC) can be centralized and published to offer users a solution that can

be rapidly deployed and easily maintained. Also, multi-GPU architectures can be leveraged for greater scalability and cost effective solutions. This talk will provide an overview of the hardware and software stack that will allow near desktop experience of a remote enterprise class graphics application.

Key Topic	Visualization
Summit	GPU Developers Summit
Speaker	Phillippe Rollin (Applied Engineer, NVIDIA)

THURSDAY 5:00 PM – 6:00 PM
VALLEY ROOM

1056 Multiparticle Simulation

A diverse array of science, engineering, and computer graphics applications involve simulations of large numbers of particles. These involve computation of interactions between many particles, potentially mediated by a spatial data structure such as a grid. Improvements in computation efficiency can be achieved by sorting particles to determine which particles are involved in interactions or undergo close approaches. Nearest neighbor or collision pair groupings can be used to reduce the total number of computation steps by reducing the number of queries for collisions or can speed up and improve accuracy of simulations via a multiple timestep integrator. Identification of nearest neighbor and collision partner groupings is a task that can be efficiently implemented in parallel on the GPU reducing the number of interactions that must be computed. A broad class of problems known as Particle-In-Cell (PIC) code advect particles through cells of a surrounding grid. During this roundtable we will discuss strategies for increasing the efficiency of multiparticle simulations as a general problem as well as challenges for multiparticle simulation in specific settings such as astrophysics, SPH, PIC, and granular flows.

Key Topic	Astro GPU
Summits	GPU Developers Summit, NVIDIA Research Summit
Speaker	Alice Quillen (Associate Professor, University of Rochester)

THURSDAY 5:00 PM – 6:00 PM
PIEDMONT ROOM

1098 NEXUS Lab

Experience the future of GPU Computing development today! NEXUS is NVIDIA's new GPU Computing development environment for Microsoft Visual Studio. In this hands-on lab, attendees will learn to use the NEXUS source debugger and profiler to efficiently solve common problems encountered when debugging and profiling CUDA C code, including

> Debugging logic errors using hardware breakpoints and variable inspection

> Reading and visualizing GPU memory directly in the Visual Studio IDE (no manual readbacks required!)

> Real-time expression evaluation on the GPU while your application is halted at a breakpoint

> Hardware-evaluated breakpoint conditionals for debugging a specific block or thread

> Detecting out-of-bounds memory errors

> Finding erroneous memory overwrites using data breakpoints

> Understanding the performance of your whole system using the Nexus Analysis tools

NVIDIA engineers will be available at each session to answer your NEXUS and general GPU Computing questions. Lab workstations are limited to the first 20 attendees per session. . NVIDIA would like to thank HP and Parallels for their generous support of SLI Multi-OS workstations.

Key Topics	General Interest, Tools
Summit	GPU Developers Summit
Speakers	Brian Kretzler (NVIDIA); Greg Smith (Manager of Developer Tools, NVIDIA)

THURSDAY 5:30 PM – 6:00 PM
GARDEN ROOM

1418 Heterogeneous Multicore Parallel Programming

Modifying legacy codes to take advantage of GPU acceleration can be a time consuming task, and portability concerns must be carefully considered at the same time. In this presentation we give an overview of HMPP (Heterogeneous Multicore Parallel Programming) workbench that addresses this challenge by providing:

> An incremental use of NVIDIA GPUs through the insertion of directives into C & Fortan source code, from which the CUDA code is automatically generated

> Performance tuning as well as CPU-GPU data transfers optimizations directives

We will also share some experiences from HMPP customers porting their key applications to NVIDIA Tesla systems.

Key Topics	High Performance Computing, Tools
Summit	GPU Developers Summit
Speaker	François Bodin (CTO, CAPS entreprise)

THURSDAY 5:30 PM – 6:00 PM
HILLSBOROUGH ROOM

1437 Directing Experiments in the International Space Station With GPU-Assisted Image Analysis

We implement image correlation, a fundamental component of many real-time imaging and tracking systems, on a graphics processing unit (GPU) using NVIDIAs CUDA. We use our code to analyze images of liquid-gas phase separation in a model colloid-polymer system, photographed in the absence of gravity aboard the International Space Station (ISS). Our GPU code is 4000 times faster than simple MATLAB code performing the same calculation on a central processing unit (CPU), 130 times faster than simple C code, and 30 times faster than optimized C++ code using single-instruction, multiple data (SIMD) extensions. The speed increases from these parallel algorithms enable us to analyze images downlinked from the ISS in a rapid fashion and send feedback to astronauts on orbit while the experiments are still being run.

Key Topic	Astro GPU
Summits	GPU Developers Summit, NVIDIA Research Summit
Speaker	Peter Lu (Post-Doctoral Research Fellow, Harvard University)

THURSDAY 5:30 PM – 6:00 PM
SACRAMENTO ROOM

1449 Uncontracted Rys Quadrature Implementation of Up to g Functions on GPUs

We present an implementation of Rys quadrature algorithms for Electron Repulsion Integrals (ERI) of up to g functions on Graphical Processing Units (GPUs). We outline the general GPU programming model, challenges associated with implementing the Rys quadrature on highly parallel emerging architectures, and our approach to implementing the quadrature. The performance of the implementation is evaluated for single and double precision on two GPU devices. The performance obtained is on par with the matrix-vector routine from the CUDA Basic Linear Algebra Subroutines (CUBLAS) library.

Key Topic	Algorithms & Numerical Techniques
Summits	GPU Developers Summit, NVIDIA Research Summit
Speaker	Veerendra Allada (Graduate Research Assistant, Ames Laboratory / Iowa State University)

THURSDAY 5:30 PM – 6:00 PM
ATHERTON ROOM

1452 Visual Innovation in the Classroom

3D models are transforming primary education. Research has proven that only 35% of students can think in 3D, so visualization has become an important tool in the classroom for improving engagement, comprehension and boosting test scores. Sean O'Brien, president of NavTech, will discuss the innovative technologies being leveraged in K – 12 classrooms, and provide a development overview of the technology used to drive these applications. Sean will share some of the models that secondary and college educators have built to improve learning in science, technology, engineering and math (STEM). Attendees will don 3D glasses and be given a demonstration of scientific concepts being taught in today's classrooms.

Key Topic	General Interest
Summit	GPU Developers Summit
Speaker	Sean OBrien (President, NavTech)

THURSDAY 5:30 PM – 6:00 PM
GOLD ROOM

1461 A Large Scale Simulation of Lattice QCD with a GPU Cluster

Quantum Chromodynamics (QCD) is the quantum field theory of the strong interaction, describing the interactions of the quarks and gluons making up hadrons (e.g., proton, neutron, and pion). Most importantly, it accounts for the nuclear energy inside an atom, as well as plays an important role in the evolution of the early universe. To solve QCD is a grand challenge among all sciences. Now the most promising approach to solve QCD nonperturbatively is to discretize the continuum space-time into a 4 dimensional lattice (i.e., lattice QCD), and to compute physical observables with Monte Carlo simulation. For lattice QCD with exact chiral symmetry, it often requires supercomputers (e.g., 10 racks of IBM BlueGene) to perform the simulations. The TWQCD Collaboration in Taiwan is the first lattice QCD group around the world to use a GPU cluster (with 120 GPUs) to perform large-scale unquenched simulations of lattice QCD with the optimal domain-wall fermions, attaining 14 Teraflops (sustained) at a price of $200,000. This has significant impacts to the lattice QCD, as well as the physics of the strong interaction.

Key Topics	General Interest, High Performance Computing
Summit	GPU Developers Summit
Speaker	Ting-Wai Chiu (Professor, National Taiwan University)

THURSDAY 6:00 PM – 8:00 PM
IMPERIAL BALLROOM

1453 Exhibits Open / Networking Happy Hour (Thursday)

Join your colleagues in the exhibit hall to preview emerging technologies and see some of the most innovative solutions available today. Appetizers and drinks will be served.

Key Topic	General Interest
Summit	GPU Developers Summit

GPU TECHNOLOGY CONFERENCE

FRIDAY

FRIDAY 8:30 AM – 10:00 AM
REGENCY 1+2 ROOM

1423 Day 3 Keynote with Richard Kerris, Lucasfilm

Games and interactive media have long been the beneficiaries of cutting edge GPU technology and it has not gone unnoticed in the world of feature film production. To date the visual effects industry had been a sideline observer of these advances while awaiting technology to reach maturity. At Lucasfilm, research and development has been on-going for some time and this past summer Industrial Light & Magic employed this technology in two of its summer blockbuster films. Lucasfilm CTO, Richard Kerris, will show a brief history of their computer graphics for film, and will then pull back the curtain on how they are now using GPU technology to advance the state of the art in visual effects and provide a glimpse of what's on the horizon for GPU's in future and how it will impact filmmaking. Due to overwhelming response and limited capacity if you are interested in attending any of the keynotes or general sessions, be sure to schedule this session into your personal calendar and arrive early.

Key Topic	General Interest
Summits	Emerging Companies Summit, NVIDIA Research Summit, GPU Developers Summit
Speaker	Richard Kerris (CTO, Lucasfilm Ltd)

FRIDAY 10:00 AM – 10:30 AM
CALIFORNIA ROOM

1442 Numerical Cosmology Powered by GPUs

By definition, cosmology cannot rely on lab experiments to reproduce the phenomenons observed in the sky and test its theories. For this very reason, the use of numerical simulations is widely spread within this community in order to understand the formation of the astrophysical objects and to put constrains on the physical ingredients that lead to the Universe as it is currently observed. Since 2001, I have been personnaly involved in trying to understand these questions through the intensive use of numerical simulations that reproduce the evolution of the Universe from the Big-Bang to our epoch. During the last two years, I have been investigating the new possibilities offered by GPUs to boost these numerical calculations, mostly using CUDA. At the current stage, three applications benefited from these studies and using 8800 GTX and Tesla C1060 devices, we found that accelerations range from factors 20 to 80 compared to CPU versions : – a cosmological N-Body integrator CUDAPM. It follows the evolution of millions of particles that interact through gravitation in an expanding Universe, modelling the rise of large scale structures. – a non-linear full multigrid solver, for the Poisson equation of modified Newtonian gravity (CUDAMOND). – a cosmological radiative transfer code CUDATON. It models the propagation of ionising radiation and its effect on the gas that filled the Early Universe. This application is multi-gpu and currently runs on 192 devices on the CCRT supercomputing centre.

Most of the techniques used in these applications are fairly standard and are not specific to astrophysics and cosmology. Therefore describing my own experience of porting these applications to GPUs as a physicist is likely to benefit to a large public of numerical scientists.

Key Topic	Astro GPU
Summits	GPU Developers Summit, NVIDIA Research Summit
Speaker	Dominique Aubert (Dr, Université de Strasbourg/ Observatoire Astronomique)

FRIDAY 10:00 AM – 11:00 AM
REGENCY 1+2 ROOM

1426 Fireside Chat with Jen-Hsun Huang and Jon Peddie

For the second year in a row, NVIDIA CEO Jen-Hsun Huang will participate in a freewheeling fireside chat with leading industry analyst Jon Peddie, of Jon Peddie Research. Topics will include the state of the industry, NVIDIA's strategy and the future of GPU computing, as well as other issues raised by the audience.

Key Topic	General Interest
Summit	Emerging Companies Summit
Speaker	Jen-Hsun Huang (CEO and Co-Founder, NVIDIA)

FRIDAY 10:00 AM – 11:30 AM
GOLD ROOM

1082 Advanced Numeric Computing

Modern GPUs offer tremendous potential for performance and efficiency gains in large-scale arithmetic-heavy scientific computing applications. To realize the very high potential throughput capacity of the latest GPUs, it is important to focus on the *basic formulation* of the numerical algorithm as well as the implementation in CUDA or OpenCL. We will discuss how the massive parallelism and relative abundance of floating point units on NVIDIA's latest GPUs can impact algorithmic choice.

Dr. Bell will discuss the design of sparse matrix-vector multiplication (SpMV) algorithms, a broadly important problem in numeric computing made difficult by the varied and sometimes highly irregular matrices encountered in sparse operations. Dr. Bell will explore SpMV methods that expose substantial fine-grained parallelism and impose sufficient regularity on execution paths and memory access patterns, achieving large speedups over modern quad-core CPU implementations. Prof. Warburton will discuss a new type of discontinuous Galerkin based solver for time-domain electromagnetics and acoustics that he has developed following a "flops-for-free" philosophy. The reformulation increases the number of floating point operations but massively decreases the memory footprint to achieve a much higher computational throughput.

Key Topic	Physics Simulation
Summit	NVIDIA Research Summit
Speakers	Nathan Bell (Research Scientist, NVIDIA); Timothy Warburton (Associate Professor, Rice University)

FRIDAY 10:30 AM – 11:00 AM
GARDEN ROOM

1108 Exploiting the GPU in Ultra High-End 4K Video Servers

Zaxtar is the highest performance video server that will feed up to 4GB/sec of video data to 4K projectors or 4K displays. This performance is accomplished by utilizing CPU, GPU and synchronization of multiple computers. The most important feature is mathematically lossless compression, which can compress video or graphics data at the ratio of 3 to 1 without losing any information. Mathematically lossless compression has been achieved by CPU up until now, but we have ported the algorithm to Quadro FX 5800. We will talk about the algorithm.

Key Topics	Algorithms & Numerical Techniques, Visualization, Computational Imaging
Summit	GPU Developers Summit
Speaker	Mark Marrin (Chief Engineer, Zaxel Systems)

FRIDAY 10:30 AM – 11:00 AM
FAIRFIELD ROOM

1130 SCVF: A Framework for Visualization of Grid and Particle Based CUDA Simulations

The Stone Ridge CUDA Visualization Framework (SCVF) provides a user API that allows developers to easily visualize their simulations on 2D grid, 3D grid and free particle systems. It allows users to rapidly create visual demos of their CUDA enhanced applications and make direct visual comparison between GPU and CPU implementations.

Key Topic	Visualization
Summit	GPU Developers Summit
Speakers	William Brouwer (Senior Physicist, Stone Ridge Technology); Ben Karel (Stone Ridge Technology)

FRIDAY 10:30 AM – 11:00 AM
HILLSBOROUGH ROOM

1139 CUDA Kernel Performance Analysis Using NVPerfMon

Performance optimization is one of the more difficult aspects of kernel development for NVIDIA's CUDA-enabled GPUs. It's often not easy to determine where compute cycles are spent, particularly complex kernels. In this talk I will describe a performance analysis methodology using NVPerfMon, a program annotation-based timing analysis framework that can be used to generate timelines of important events, such as thread start and stop times. NVPerfMon can provide deep visibility into the performance pathologies of a CUDA kernel and can expedite the time to attaining performance goals. MultiCoreWare has extended Nvidia's NVPerfMon tool for High-Performance Computing applications and environments, with full 64-bit capability, and full functionality in a multiple GPU environment.

Key Topics	High Performance Computing, Tools
Summit	GPU Developers Summit
Speaker	Sanjay Patel (CTO / Professor, MultiCoreWare / University of Illinois)

FRIDAY 10:30 AM – 11:00 AM
CUPERTINO ROOM

1144 Use of High Performance GPUs in Military/Aerospace and Defense Applications

GPUs have long been used in Military/Aerospace and Defense applications, but with the advent of GPGPU computing the range of problems that will be solved by GPUs with the CUDA architecture is set to expand dramatically. This session will look at some of the traditional and new applications for GPUs in the Military/Aerospace arena, together with some of the challenges involved in taking GPU technology into the hostile real-world environments of this market.

Key Topic	General Interest
Summit	GPU Developers Summit
Speaker	Jim Berlin (CTO, GE Fanuc)

FRIDAY 10:30 AM – 11:00 AM
BELVEDERE ROOM

1147 The High Life for High Performance

Samsung Semiconductor and Lenovo have teamed up to push the envelope on notebook performance and data reliability. In this presentation, the two leading-edge companies will take a close look at market-ready, state-of-the-art notebook technology for high performance computing applications with case study information that focuses on real-world applications. They will also discuss the PC industry's decisive movement to solid state drives, and future technology for product designers.

Key Topic	General Interest
Summit	GPU Developers Summit
Speakers	Henry Ramos (Senior Field Technical Specialist Lenovo); Brian Beard (SSD Marketing Manager Samsung Semiconductor, Inc.)

FRIDAY 10:30 AM – 11:00 AM
GLEN ELLEN ROOM

1404 Real Time Cloth: 3D Triangular Mesh Algorithms on the GPU

Rendering has been the classic challenge of 3D graphics. But today, with complex scenes having millions of polygons and wanting to add physical effects – just moving and deforming those polygons presents a growing challenge especially when trying to achieve real time performance. In some dynamic scenes, tens of thousands of vertexes need to have their

new position recalculated for each frame and in some cases (like physical simulation) every sub frame. For this reason, using the GPU to perform these calculations is very tempting. But one of the problems that arise from moving these calculations to the GPU is, triangular mesh topologies don't have a predefined vertex locality, it depends on the arbitrary indexing mechanism used when the topology was created. And even in optimal circumstances, there are always vertexes which destroy the locality. At OptiTex, we have been using CUDA since the initial 0.8 Beta. OptiTex specialty is cloth. Using CUDA we are on the way to deliver an implicit cloth engine, which will work in real time for cloth meshes up to 10,000 polygons. In our process we have ported (and are still working) all of our geometry calculations to the GPU. These include normals, smooth skin, collision detection, finite elements and more.

During the development of these parts, different problems occurred relating to vertex index locality. These were solved using a variety of methods and approaches.

Key Topic	3D
Summit	GPU Developers Summit
Speaker	Eri Rubin (Head of CUDA Project, OptiTex, LTD.)

FRIDAY 10:30 AM – 11:30 AM
EMPIRE ROOM

1006 Implementing 3D Finite Difference Codes on the GPU

This presentation reviews GPU parallelization of 3D finite difference computation over regular grids. 3DFD is a fundamental computation in many applications, including Reverse Time Migration in seismic computing. A single-GPU implementation is described first, followed up by a scalability study on a cluster of up to 8 GPUs. Performance results are compared to the theoretical limits of the hardware.

Key Topics	3D, Tools, Physics Simulation, Energy Exploration, Algorithms & Numerical Techniques
Summits	GPU Developers Summit, NVIDIA Research Summit
Speaker	Paulius Micikevicius (NVIDIA)

FRIDAY 10:30 AM – 11:30 AM
SACRAMENTO ROOM

1061 Real-time Visual Effects in Film and Video: the GPU Changes Everything

Powerful fully programmable GPUs and stream processing technologies such as CUDA are revolutionizing the field of professional film and video visual effects. Using stream-computing techniques, high quality effects can now be produced in real time, even at high definition resolutions and above. Traditional OpenGL techniques did not have enough image quality for professional use, and CPUs are too slow, even in parallel. Massively parallel computation changes the game for artists, who will be able to focus on telling the story rather than waiting for frames to render; studios, whose return on investment will increase measurably; and for viewers, who will reap the benefits of increased realism and fewer distractions from the story line. This will have future implications not only for film and video, but for online pro/consumer videos and games.

Key Topics	3D, Visualization, Film
Summit	GPU Developers Summit
Speaker	Gary Oberbrunner (Chief Scientist, GenArts, Inc.)

FRIDAY 10:30 AM – 11:30 AM
PIEDMONT ROOM

1098 NEXUS Lab

Experience the future of GPU Computing development today! NEXUS is NVIDIA's new GPU Computing development environment for Microsoft Visual Studio. In this hands-on lab, attendees will learn to use the NEXUS source debugger and profiler to efficiently solve common problems encountered when debugging and profiling CUDA C code, including

> Debugging logic errors using hardware breakpoints and variable inspection

> Reading and visualizing GPU memory directly in the Visual Studio IDE (no manual readbacks required!)

> Real-time expression evaluation on the GPU while your application is halted at a breakpoint

> Hardware-evaluated breakpoint conditionals for debugging a specific block or thread

> Detecting out-of-bounds memory errors

> Finding erroneous memory overwrites using data breakpoints

> Understanding the performance of your whole system using the Nexus Analysis tools

NVIDIA engineers will be available at each session to answer your NEXUS and general GPU Computing questions. Lab workstations are limited to the first 20 attendees per session. NVIDIA would like to thank HP and Parallels for their generous support of SLI Multi-OS workstations.

Key Topics	General Interest, Tools
Summit	GPU Developers Summit
Speakers	Brian Kretzler (NVIDIA); Greg Smith (Manager of Developer Tools, NVIDIA)

FRIDAY 10:30 AM – 11:30 AM
VALLEY ROOM

1106 Driving on Mars: Simulating Tracked Vehicle Operation on Granular Terrain

This work demonstrates the use of high performance GPU computing for the numerical simulation of tracked ground vehicles. The work closes a gap in physics based simulation related to the inability to accurately characterize the 3D mobility of tracked vehicles on granular terrains (sand and/or gravel). The problem of tracked vehicle mobility on granular material is approached using a discrete element method that accounts for the interaction between the track and each discrete particle in the terrain. Presently the approach captures the dynamics of systems with more than 300,000 bodies interacting simultaneously. The method is implemented on the GPU through CUDA.

Key Topics	Physics Simulation, Visualization
Summits	GPU Developers Summit, NVIDIA Research Summit
Speakers	Toby Heyn (Research Assistant, University of Wisconsin – Madison); Justin Madsen (Research Assistant, University of Wisconsin – Madison); Hammad Mazhar (Research Assistant, University of Wisconsin Madison); Dan Negrut (Professor, University of Wisconsin – Madison)

FRIDAY 10:30 AM – 11:30 AM
CALIFORNIA ROOM

1401 Convolution Soup: A Case Study in CUDA Optimization

This talk explores various optimization techniques using the simple example of image convolution. We begin with simple naive approaches and explore progressively more advanced optimizations and alternative memory access approaches, eventually arriving at an optimal implementation.

Key Topics	Algorithms & Numerical Techniques, Computational Imaging
Summit	GPU Developers Summit
Speaker	Joe Stam (NVIDIA)

FRIDAY 11:00 AM – 11:30 AM
FAIRFIELD ROOM

1110 Using the GPU for Gradient Reconstruction of Unstructured Meshes

Impressive performance gain has been obtained on field calculations for large volume of seismic survey data which is hierarchically represented [1]. Gradient reconstruction, for both scalar and vector unstructured fields, is yet another performance critical task in engineering simulations such as computational fluid dynamics (CFD) and finite element analysis (FEA). The latest GPU hardware has been improved significantly in terms of memory capacity and memory random access efficiency, which makes GPU computing attractive to engineering simulation[2]. Based on requirements drawn from cross disciplinary fields including geophysical modelling, material analysis and manufacturing, this study continues to investigate the double precision performance and its scalability across multiple GPUs. A software framework is designed where algorithms can be conveniently implemented in a heterogeneous computing environment with mixed CPU and GPU configurations. Attention has also been steered toward integration of GPU algorithms in end-user application software Avizo, which will enable the application of the algorithm in industrial aerodynamic simulations where mixed element unstructured mesh dominates.

Key Topics	Algorithms & Numerical Techniques, Visualization, Physics Simulation, Computational Fluid Dynamics
Summit	GPU Developers Summit
Speaker	Michael Heck (Technology Advisor, Visualization Sciences Group (VSG))

FRIDAY 11:00 AM – 11:30 AM
GLEN ELLEN ROOM

1113 Programming NVIDIA GPUs using PGI Accelerator Fortran and C Compilers

This talk provides an introduction to programming NVIDIA GPUs using the PGI Accelerator Programming Model in C and Fortran. It is suitable for application programmers, in particular those who are not expert GPU programmers. This talk introduces the compute-specific details of the NVIDIA GPU, and through examples, illustrates how to program common computational algorithms on NVIDIA GPUs using portable directive- based C and Fortran 95/03. The material covers programming language features, interpreting compiler feedback, performance analysis, and performance tuning. This talk includes a live component with a Linux workstation containing a Tesla card, and the latest PGI Accelerator compilers and tools.

Key Topic	Algorithms & Numerical Techniques
Summit	GPU Developers Summit
Speaker	Michael Wolfe (The Portland Group)

FRIDAY 11:00 AM – 11:30 AM
HILLSBOROUGH ROOM

1134 Using OpenKODE to Create a Seamless Driving Experience

The software for an in-car infotainment system requires many different graphical technologies, including 2D widget toolkits, 3D OpenGL ES graphics, navigation engines, and web browsers, all using different APIs and implemented with different coding techniques. Somehow, these technologies must all be brought together to operate as a single, purpose-built system. By leveraging the open standard OpenKODE API, development teams that use Nvidia GPUs can merge the output of these separate technologies into a single GUI, saving development effort while improving the user experience. In this session, attendees learn how the OpenKODE API can help development teams manipulate graphics to create the unified experience that customers demand.

Key Topics	Embedded & Mobile, Visualization
Summits	GPU Developers Summit, NVIDIA Research Summit
Speakers	Andy Gryc, Senior Product Marketing Manager, Graphics and Multimedia Technologies QNX Software Systems)

FRIDAY 11:00 AM – 11:30 AM
CUPERTINO ROOM

1145 Integrated Design that Shatters 1U Server Performance Records

Supermicro and NVIDIA recently announced a new class of servers that combines massively parallel GPUs with multi-core CPUs in a single server system. This unique configuration delivers performance at least an order of magnitude better than traditional quad-core CPU-based servers. This breakthrough technology immediately provides users with the ability to implement tasks that were traditionally addressed only with massive supercomputers or that were simply unsolvable. With GPUs and CPUs on a co-processing 1U server, data centers can achieve greater system manageability and a uniform building block for deployment.

Supermicro will provide an overview of thid groundbreaking product in this session.

Key Topic	General Interest
Summit	GPU Developers Summit
Speaker	Don Clegg (Vice President, Marketing and Business Development, Super Micro Computer, Inc.)

FRIDAY 11:00 AM – 11:30 AM
(Belvedere Room)

1146 Next-Generation Graphics DRAM: Challenges and Opportunities

GDDR5 is the most up-to-date graphics memory technology today, offering an impressive bandwidth of between 4 and 7Gbps. However, demand continues for the development of finer and finer graphics to be used in games as well as professional design applications. So, the memory Industry is considering how it is going to move beyond GDDR5 levels in a couple of years to reach bandwidths of more than 10Gbps. Considerable obstacles await. Samsung research indicates that it will be difficult to overcome a few major barriers to attaining and exceeding data transfer speeds in excess of 10Gbps with the single-ended technology that is used today.

The industry will need to uncover a cost-efficient way to move beyond single-ended memory chip designs. Memory vendors will have to closely examine a variety of factors including differential technology, and wide I/O schemes. This presentation will review GDDR5 technology and present various ideas on how best to exceed the bellwether threshold of 10Gbps.

Key Topic	General Interest
Summit	GPU Developers Summit
Speakers	Jimmy Chung (Technical Marketing Manager, Graphics Memory, Samsung Semiconductor, Inc.)

FRIDAY 11:00 AM – 11:30 AM
ATHERTON ROOM

1459 Multi-Frame Analysis: The Future of Video Processing

MotionDSP makes software that significantly improves video from a wide range of sources – from mobile phones to surveillance cameras. Its patented, multi-frame video enhancement technology is available in Ikena, a video forensic solution, and in vReveal, a Windows application for consumers.

High-fidelity motion estimation plays a critical role in almost any video processing technology. The open-loop nature of MotionDSP's multi-frame video reconstruction framework requires a departure from quantized block-based motion estimation concept, commonly used in video encoding, as super-resolution temporal filtering cannot rely on a prediction error to compensate for motion estimation inaccuracies. The need to use floating-point accurate motion fields, coupled with a multi-frame nature of the algorithm, makes GPGPU a natural choice when addressing this massive computational problem.

This talk will provide an overview of MotionDSP's video technology and products, as well as details of the design and performance challenges of MotionDSP's multi-frame video pipeline, one of the first commercial video processing solutions fully ported to GPU architecture.

Key Topic	Computational Imaging
Summit	GPU Developers Summit
Speaker	Nikola Bozinovic (VP Engineering, MotionDSP)

FRIDAY 11:30 AM – 12:30 PM
IMPERIAL BALLROOM

1424 Exhibits Open / Networking Lunch (Friday)

Join your colleagues in the exhibit hall to preview emerging technologies and see some of the most innovative solutions available today. Lunch will be served to Full Conference, Research Summit and Emerging Companies Summit pass holders.

Key Topic	General Interest
Summits	Emerging Companies Summit, NVIDIA Research Summit, GPU Developers Summit

FRIDAY 11:30 PM – 12:30 PM
REGENCY 1+2 ROOM

1428 Raising Capital in Difficult Environments

The current economic situation poses unique challenges for start-ups seeking to raise critical funds to grow their businesses. This panel will attempt to assess the current state of both the public and private markets, and will explore various strategies and options for raising capital. Topics will include equity and debt, as well as other creative forms of financing such as NRE arrangements, etc. Jeff Herbst, NVIDIA's Vice President of Business Development will lead this panel discussion, which will include leading industry experts from Sutter Hill Ventures, Citi, Silicon Valley Bank and Deloitte.

Key Topic	General Interest
Summit	Emerging Companies Summit
Speaker	Moderator: Jeff Herbst (Vice President of Business Development, NVIDIA); Panelists: Bill Frauenhofer, (Managing Director, Citi Technology West Coast Investment Banking Group); Jon Gavenman (Partner, Cooley Godward Kronish); Heather Gates-Massoudi (Director, Venture Capital Services Group, Deloitte Services LP); Peter Kidder (Division Risk Manager, Silicon Valley Bank); Andrew T. Sheehan (Managing Director, Sutter Hill Ventures)

FRIDAY 1:00 PM – 1:30 PM
CALIFORNIA ROOM

1062 Astrophysical Fluid Simulation Using Adaptive Meshes

Adaptive mesh fluid simulations play a crucial role in many areas of astrophysical research including the formation and explosion of stars, jets from black holes, etc. A parallel adaptive mesh multi-physics fluid code, Enzo, has been widely used in astrophysical community in recent years. In this talk I will describe a CUDA implementation of the finite volume fluid solver used in Enzo. The GPU version shows significant speed-up compared to the CPU version.

Key Topic	Astro GPU, High Performance Computing
Summits	GPU Developers Summit, NVIDIA Research Summit
Speaker	Peng Wang (NVIDIA)

FRIDAY 1:00 PM – 1:30 PM
GARDEN ROOM

1077 Bullet: A Case Study in Optimizing Physics Middleware for the GPU

We have been porting and implementing all stages of the collision detection and rigid body dynamics pipeline onto GPU. We describe the changes in data structures and algorithms to parallelize for CUDA and OpenCL. The implementation is available as open source in the Bullet Physics SDK.

Key Topic	General Interest
Summit	GPU Developers Summit
Speaker	Erwin Coumans (Physics Simulation Team Lead, Sony Computer Entertainment US R&D)

FRIDAY 1:00 PM – 1:30 PM
BELVEDERE ROOM

2008 Right Hemisphere Startup Presentation

Right Hemisphere is the leading provider of enterprise visual product communication, collaboration and publishing solutions that optimize mission critical business processes such as sourcing, manufacturing, marketing and customer support for leading Global 1000 manufacturers. Right Hemisphere integrates disparate product data readily available in CAD, PLM, and ERP systems and automates the delivery of intelligent product information as common documents and files or directly into business applications to downstream stakeholders on a global basis. Founded in 1997, Right Hemisphere is a privately held, venture-funded corporation based in Silicon Valley and Auckland, New Zealand. For more information please visit www.righthemisphere.com.

Key Topics	General Interest, Visualization
Summit	Emerging Companies Summit
Speaker	Michael Lynch (CEO, Right Hemisphere)

FRIDAY 12:00 PM to 1:00 PM
CALIFORNIA ROOM

1021 CUDA Research Roundtable: Mixed Precision GPU Computing

Many algorithms used in computational physics can be greatly accelerated by the use of GPUs. However, the full double-precision floating point operations to which scientists are accustomed can prove costly, especially in compute-intensive applications where floating-point computations rather than memory bandwidth limit performance. In fact, many scientific problems actually require double precision in only a small subset of the code. In these cases, the development of mixed-precision algorithms can bring substantial improvements without sacrificing overall accuracy. Double precision is used only where it is required; the remaining calculations are carried out in single precision. Although the task of identifying "necessary" double-precision code may be non-trivial, the performance payoff can be considerable. This approach also facilitates software emulation of double precision in key portions of the code, which can be effective in accelerating GPU double-precision performance. In practice, the emulation is neither complete nor IEEE-compliant, and is cumbersome to code without support at the compiler or processor levels, but impressive improvements in speed have been obtained. This roundtable will discuss progress made so far in mixed-precision calculations and emulation techniques for GPUs, and will consider the prospects for future development of these approaches.

Key Topic	General Interest
Summits	GPU Developers Summit, NVIDIA Research Summit
Speaker	Steve McMillan (Professor, Drexel University)

FRIDAY 1:00 PM – 1:30 PM
FAIRFIELD ROOM

2028 Teradici Corporation

Teradici Corporation has developed PCoIP (PC-over-IP) technology, a unique remote display protocol, which makes network delivered computing a viable corporate computing reality. Founded in 2004 and headquartered in Burnaby, BC, Canada, Teradici for the first time enables an exceptional end user experience for datacenter-based computing. Through a combination of unique graphics algorithms, high-performance silicon processing, and workstation/server add-in cards, the company is changing how personal computers are used, deployed and managed. More information is available at www.teradici.com.

Key Topic	3D
Summit	Emerging Companies Summit
Speaker	Dan Cordingley (President & CEO, Teradici Corporation)

FRIDAY 1:00 PM – 1:30 PM
CUPERTINO ROOM

2034 Hanweck Associates, LLC Startup Presentation

As an established leader in high-performance quantitative financial computing, Hanweck Associates pioneered commercial GPU-based computing solutions for the finance industry, focusing tomorrow's technology on today's risk management, research and trading problems. Our customers include top-tier banks, broker/dealers, asset managers and hedge funds around the world. Hanweck Associates' team brings over 40 years of hands-on expertise in financial markets and technology services. Our experience as traders, strategists, researchers and technologists at major financial institutions spans all major asset classes and products.

Key Topic	General Interest
Summit	Emerging Companies Summit
Speaker	Gerald Hanweck (Principal, Hanweck Associates, LLC)

FRIDAY 1:00 PM – 1:30 PM
GLEN ELLEN ROOM

2038 Affine Systems Startup Presentation

Affine Systems is an online video advertising company that can target advertisements against specific videos across all the top online video web sites. Affine uses its proprietary computer vision technology to scan all the videos on these web sites to find the best advertising opportunities for its clients. Affine's core vision technology is able to recognize people, brands and products in the visual stream of a video, as well as detect inappropriate content and extract important information about the quality and content of the videos that it scans. By using this metadata to target video advertising campaigns, Affine is able to achieve better ROI for advertisers and more revenue for publishers.

Key Topic	General Interest
Summit	Emerging Companies Summit
Speaker	Michael Sullivan (CEO, Affine Systems)

FRIDAY 1:00 PM – 2:00 PM
REGENCY 2 ROOM

1047 Computer Vision with Horst Bischof

The area of computer vision is entering an exciting period because it is for the first time in its history that we are not limited by the hardware, we have now enough camera resolution, enough memory and computing power and access to virtually an unlimited mount of images via the internet. The challenging question we have to answer now, is how to make use of that. In this talk I will highlight some of these recent technological trends and how they influence the area of computer vision and more important speculate about their future impact on the field. Among the topics I touch in the talk are computational photography, interactive vision, internet vision and of course GPU-based visual computing.

Key Topics	Computational Imaging, Computer Vision
Summit	NVIDIA Research Summit
Speaker	Horst Bischof (Professor, Graz University of Technology)

FRIDAY 1:00 PM – 2:00 PM
ATHERTON ROOM

1057 NEXUS: A Powerful IDE for GPU Computing on Windows

Unleash your productivity with NEXUS, NVIDIA's new development environment for GPU Computing and graphics applications that use CUDA C, OpenCL, DirectCompute, Direct3D, or OpenGL. NEXUS introduces native GPU debugging and platform-wide performance analysis tools for both computing and graphics developers, fully integrated into Visual Studio 2008. In this session, you will learn to use the powerful NEXUS debugger to set breakpoints and step through your CUDA C kernels or HLSL shader source code. Easily catch hard-to-find bugs using the NEXUS memory inspection tools and set data breakpoints on your running application. Visualize your application's workloads and performance characteristics across the CPU, GPU, and operating system, and then dig deeper using the NEXUS profilers. NEXUS also includes state of the art graphics debugging features and convenient API state inspection. This session will include information on how to register for the NEXUS Beta Program.

Key Topics	General Interest, Tools
Summit	GPU Developers Summit
Speaker	Sebastien Domine (Sr. Dir. Dev Tech Tools, NVIDIA)

FRIDAY 1:00 PM – 2:00 PM
GOLD ROOM

1083 CUDA Fortran Programming for NVIDIA GPUs

This talk provides an introduction to programming NVIDIA GPUs using CUDA Fortran. It is suitable for expert Fortran or CUDA C programmers who need to extract maximum performance from GPUs using an explicit GPU Fortran programming model. This talk introduces the CUDA Fortran language, and through examples, illustrates how to explicitly program NVIDIA GPUs in native Fortran 95/03 through creation of GPU kernel subroutines, management of host and GPU device memory, definition of CUDA grids and thread blocks, launching kernels on an NVIDIA GPU device, and use of the CUDA Fortran runtime API. This talk includes a live component with a Linux workstation containing a Tesla card, and the PGI CUDA Fortran compiler.

Key Topics	High Performance Computing, Tools
Summits	GPU Developers Summit, NVIDIA Research Summit
Speaker	Michael Wolfe (The Portland Group)

FRIDAY 1:00 PM – 2:00 PM
HILLSBOROUGH ROOM

1087 High Dynamic Range Spectral Imagery

X-Rite specializes in the construction of devices that acquire highly accurate color data. As part of our advanced development, we have engineered a process for generation of high dynamic range (HDR) spectral images from digital CCD cameras using highly characterized LED illumination sources. Our instrumentation involves multiple cameras and illumination sources arranged to produce a gonio-spectrophotometric set of images. Within this process, standard image processing techniques exacted a high computational cost and initially prevented this process from being attractive and useful. CUDA enabled us to achieve a roughly 140x computational improvement over the original implementation, both by parallelizing the computations via the GPU and by encouraging specific algorithm optimizations.

Key Topics	Film, Visualization
Summit	GPU Developers Summit
Speaker	Marc Ellens (Software Engineer, X-Rite, Inc.)

FRIDAY 1:00 PM – 2:00 PM
PIEDMONT ROOM

1098 NEXUS Lab

Experience the future of GPU Computing development today! NEXUS is NVIDIA's new GPU Computing development environment for Microsoft Visual Studio. In this hands-on lab, attendees will learn to use the NEXUS source debugger and profiler to efficiently solve common problems encountered when debugging and profiling CUDA C code, including

> Debugging logic errors using hardware breakpoints and variable inspection

> Reading and visualizing GPU memory directly in the Visual Studio IDE (no manual readbacks required!)

> Real-time expression evaluation on the GPU while your application is halted at a breakpoint

> Hardware-evaluated breakpoint conditionals for debugging a specific block or thread

> Detecting out-of-bounds memory errors

> Finding erroneous memory overwrites using data breakpoints

> Understanding the performance of your whole system using the Nexus Analysis tools

NVIDIA engineers will be available at each session to answer your NEXUS and general GPU Computing questions. Lab workstations are limited to the first 20 attendees per session. NVIDIA would like to thank HP and Parallels for their generous support of SLI Multi-OS workstations.

Key Topics	General Interest, Tools
Summit	GPU Developers Summit
Speakers	Brian Kretzler (NVIDIA); Greg Smith (Manager of Developer Tools, NVIDIA)

FRIDAY 1:00 PM – 2:00 PM
EMPIRE ROOM

1400 GPU Metaprogramming using PyCUDA: Methods & Applications

Writing reliable GPU codes that achieve peak performance in the face of changing requirements and hardware platforms can be a challenging task. In this talk, I will introduce the open-source PyCUDA toolkit, which assists in this task in a number of ways: Convenient, high-level interface PyCUDA binds all functionality in Nvidia CUDA to a convenient interface in the high-level scripting language Python. Resource management and error checking are automatic. Code Templates PyCUDA comes with tuned and debugged code for many common operations, such as vector math and reductions, which saves debugging and coding time. Metaprogramming PyCUDA allows GPU code to be generated at run-time code, which

makes many advanced programming techniques easy–such as empirical optimization, constant folding, and run-time specialization. Scalability PyCUDA covers "small-scale" and "large-scale" uses alike: It allows quick prototyping and experimentation, but it also integrates easily into large-scale computational software. Having introduced the toolkit, I will show how PyCUDA has supported a number of applications in computational science: First, we have successfully used PyCUDA in a high-performance discontinuous Galerkin finite element (DG-FEM) solver. The term DG-FEM describes a family of high-order accurate numerical methods for systems of partial differential equations that model real-world processes such as electromagnetic scattering or fluid flow. We found that these methods' algorithmic structure makes them very suitable for execution on a GPU, often achieving speedup factors on the order of 50 when compared to a single CPU core. PyCUDA and GPU metaprogramming were crucial in achieving this level of performance. As an added benefit, the resulting solver turned out to be very versatile with respect to equation types, domain dimensionality, and discretization parameters.

Second, I will discuss a recent effort seeking to automate the writing of high-performance GPU code for a large class of computational kernels that includes many of those needed for the numerical discretization of PDEs. Again, code generation and empirical optimization as provided by PyCUDA provide the basis for the approach that may make GPU performance possible even in situations where manual development is not economical.

Key Topic	Tools
Summit	GPU Developers Summit
Speaker	Andreas Kloeckner (Graduate Student, Brown University)

FRIDAY 1:00 PM – 2:00 PM
CRYSTAL ROOM

1405 Adobe Creative Suite 4 – Performance Enhancements Featuring GPU Acceleration

Plus a sneak peak at the Future of GPU Acceleration!

Many Adobe Creative Suite 4 applications have enhanced features designed to take advantage of nVidia GeForce and Quadro GPU cards to accelerate application performance. In this session, find out how Adobe Photoshop, After Effects and Adobe Premiere Pro get a distinct performance boost when handling large files or creating motion graphics and visual effects. Plus you'll get a sneak peek at what's coming in the future around GPU acceleration in Adobe Premiere Pro.

Key Topics	Computational Imaging, Visualization
Summit	GPU Developers Summit
Speaker	Kevan O'Brien (Open Workflows Specialist, Adobe Systems)

FRIDAY 1:00 PM – 2:00 PM
SACRAMENTO ROOM

1451 iray®: Interactive, Zero Tuning, Consistent Rendering in mental ray, RealityServer and other Applications

iray is a interactive, global illumination ray-traced rendering technology that generates photo real imagery without introducing rendering algorithm specific artifacts, and without requiring the use of renderer specific parameterizations. This "push-button" rendering capability can dramatically improve artist productivity in a wide variety of use cases, including architectural visualization, automotive styling and marketing, product styling, among others. iray takes full advantage of the CUDA programming model, allowing interactive performance on single NVIDIA GPUs, and scaling almost linearly on multi-GPU platforms. The iray rendering mode will be available in the next versions of mental ray (3.8) and RealityServer (2.4). We will demonstrate the key features of iray for end users and application developers, including the progressive refinement of images until maximum fine detail is reached, providing a single process which smoothly combines interactive pre-visualization and final frame rendering.

Key Topic	General Interest
Summit	GPU Developers Summit
Speaker	Michael Kaplan (Vice President of Strategic Development, mental images)

FRIDAY 1:00 PM – 2:00 PM
VALLEY ROOM

1460 Supercomputing Research Roundtable

This moderated Research Roundtable will follow up Thursday's "GPU Supercomputing Supersession" with a discussion among the presenters and audience members about the issues, challenges, and opportunities in GPU-powered "big iron" supercomputing. The roundtable will bring together researchers that have recently deployed or are interested in deploying GPU clusters to enable scientific research. Attendees and presenters will share their experiences in deploying GPU technology and discuss the future of this technology in supporting research to tackle the world's most challenging scientific problems. Topics will include

> what it takes to build a GPU cluster

> current issues affecting GPU cluster computing

> how to best utilize a GPU cluster for scientific research

> software tools for scientific computing on GPU clusters

> integrating compute, data management and visualisation on GPU clusters

Key Topic	High Performance Computing
Summit	NVIDIA Research Summit, GPU Developers Summit
Speaker	John Taylor (Leader, CSIRO Computational and Simulation Sciences, CSIRO)

FRIDAY 1:30 PM – 2:00 PM
GARDEN ROOM

1138 Rigid Body, Cloth and Fluid Physics for CUDA

Since NVIDIA's acquisition of the PhysX technology in 2008, we've been working hard to port Rigid body, Cloth and Fluid PhysX libraries to CUDA. This presentation describes some of the algorithms, their parallelization and optimization for GPUs using CUDA.

Key Topic	Physics Simulation
Summit	GPU Developers Summit
Speaker	Richard Tonge (Senior Software Engineer – PhysX, NVIDIA)

FRIDAY 1:30 PM – 2:00 PM
CALIFORNIA ROOM

1402 Applications of Graphics Processing Units to the Binary Black Hole Evolutions

We apply general-purpose computation on GPUs to obtain sizable speedups over a CPU in post-Newtonian evolutions of a binary black hole system. We discuss effective techniques for optimizing our GPU code on the CUDA architecture and present results demonstrating the speedups obtained. We also describe an MPI-based approach for scaling a large number of binary black hole simulations over multiple GPUs. This approach will allow us to complete the largest scientific GPU calculation to date using the NCSA Lincoln cluster.

Key Topic	Astro GPU
Summits	GPU Developers Summit, NVIDIA Research Summit
Speaker	John Silberholz (University of Maryland)

FRIDAY 1:30 PM – 2:00 PM
FAIRFIELD ROOM

2022 NextIO Inc. Startup Presentation

NextIO, Inc. is the leader in next-generation I/O virtualization solutions for today's dynamic data center in a variety of industries including enterprise, telecommunications, high performance computing, digital media and financial services. With its innovative Express Connect platform, NextIO offers the unique ability to virtualize I/O technology on any server, operating system, hypervisor and storage architecture. Leveraging PCI Express, NextIO offers true I/O virtualization for any end-point technology. Express Connect delivers unprecedented rack-level scalability, with I/O and server resources that can be scaled independently for 50-70% savings in capital, power, and cooling. NextIO's any-to-any I/O connectivity boosts performance and reliability while streamlining IT deployment, simplifying administration and reducing costs.

Key Topics	General Interest, Visualization
Summit	Emerging Companies Summit
Speaker	KC Murphy (CEO, NextIO, Inc.)

FRIDAY 1:30 PM – 2:00 PM
BELVEDERE ROOM

2046 BumpTop Startup Presentation

Bump Technologies Inc. is a Toronto, Ontario, Canada-based company that has created and brought to market BumpTop, an innovative 3D user interface enabling natural organization and sharing. The technology was demonstrated at TED, the annual, invitation-only "Technology, Entertainment, Design" conference where the world's leading thinkers and doers gather to find inspiration; based on its success there, it received angel and venture capital funding and incorporated in 2007. Investors include GrowthWorks Capital, Extreme Ventures and angel investors Andy Hertzfeld, Austin Hill and G. R. Heffernan. Bump Technologies also provides OEMs and third party developers with a multi-touch and 3D rendering framework to enhance the interactivity and productivity of their applications.

Key Topics	3D, General Interest
Summit	Emerging Companies Summit
Speaker	Anand Agarawala (Founder and CEO, BumpTop)

FRIDAY 1:30 PM – 2:00 PM
CUPERTINO ROOM

2048 Mersive Technologies Startup Presentation

Mersive is a software company that enables the creation of large seamless displays with resolutions exceeding high-definition. As a leader in visual technology, Mersive software simplifies the way high-resolution displays are designed, configured, and maintained to provide a cost-effective, true auto-alignment solution capable of merging any number of projectors into a seamless display of almost any size, shape, and resolution. Technology once considered practical only for select capital-rich industries is now made available by Mersive for a range of applications from simulation and training to museum displays and digital signage. Mersive's products serve corporate, government, and academic organizations. Mersive mSeries Displays are unique in that they maintain their visual accuracy over time using auto-calibration software and occupy a very small footprint – 36 inches deep – compared to traditional rear-projection, seamless displays. This small footprint enables use within existing spaces; it is not necessary to make any facility modifications. Design, engineering & marketing teams can view all the details of their work in lifelike brilliance on one seamless, high pixel count display. Because the display is easily installed and configured, maintenance costs are significantly reduced thereby opening new markets for large-scale, high resolution displays.

Key Topic	Visualization
Summit	Emerging Companies Summit
Speaker	Randall Stevens (President / CEO, Mersive Technologies)

FRIDAY 1:30 PM – 2:00 PM
GLEN ELLEN ROOM

2054 Evolved Machines Startup Presentation

We are pioneering the reverse engineering of brain circuitry to build intelligent machines.

Key Topic	General Interest
Summit	Emerging Companies Summit
Speaker	Paul Rhodes (President, Evolved Machines, Inc.)

FRIDAY 2:00 PM – 2:30 PM
ATHERTON ROOM

1058 Tridiagonal Solvers on the GPU and Applications to Fluid Simulation

This presentation will explore the efficient GPU implementation of direct numerical simulation of turbulent viscous incompressible fluid in 3D domain. We will discuss solving the full system of Navier-Stokes and energy equations using the Alternating Direction Implicit (ADI) numerical method, as well as implementation details of a fast tridiagonal matrix solver on CUDA . Finally we will compare the performance of GPU and CPU on a particular modeling problem in which the GPU outperforms the latest multicore CPUs by an order of magnitude in double precision on the whole solver.

Key Topic	Algorithms & Numerical Techniques, Tools, Astro GPU, Computational Fluid Dynamics
Summits	GPU Developers Summit, NVIDIA Research Summit
Speaker	Nikolai Sakharnykh (Developer Technology Engineer, NVIDIA)

FRIDAY 2:00 PM – 2:30 PM
GOLD ROOM

1084 The TotalView Debugger for CUDA

For applications to take advantage of the performance boost offered by GPU Computing technologies and techniques, developers need to either write or re-write their programs in CUDA or another language supported by the GPU. That development process is much more efficient if developers can rely on a good debugger to give them control over and visibility into the program that they are writing. This talk will discuss a preview version of the TotalView debugger specifically developed to support CUDA development on NVIDIA GPU hardware.

Key Topic	Tools
Summit	GPU Developers Summit
Speaker	Christopher Gottbrath (TotalView Technologies)

FRIDAY 2:00 PM – 2:30 PM
HILLSBOROUGH ROOM

1122 Maximizing GPU Efficiency in Extreme Throughput Applications

NVIDIA GPUs provide the massive data throughput needed to enable real-time processing of high data volumes, such as those encountered in high definition and 4K video applications. This talk will explore several important design considerations for extreme data rates including proper overlapping of data transfer and compute, multiple GPU approaches, and efficient interoperation with OpenGL.

Key Topic	Computational Imaging, Tools
Summit	GPU Developers Summit
Speaker	Joe Stam (NVIDIA)

FRIDAY 2:00 PM – 2:30 PM
CALIFORNIA ROOM

1441 New Insights into the Phase Space of Binary Black Holes Using GPU Technology

We perform ensemble studies of binary black hole inspirals. The binary black hole problem is of great interest for the cosmological community (merger of galaxies with BHs at the center) as well as the gravitational wave community (where the merger of BHs is the most important signal source). The full binary black hole merger problem is computationally very demanding and even with advanced numerical techniques ensemble studies are currently not possible. Using a standard approximate solution to Einstein's equation (the post-Newtonian equations) one can accurately model the inspiral until shortly before merger when the approximation techniques break down. Utilizing this approximation technique we study the 7-dimensional parameter space of the BH merger problem using a Monte-Carlo approach, which extends very naturally to GPUs.

Key Topic	Astro GPU
Summits	GPU Developers Summit, NVIDIA Research Summit
Speaker	Frank Herrmann (University of Maryland)

FRIDAY 2:00 PM – 2:30 PM
FAIRFIELD ROOM

2001 Visioglobe Startup Presentation

3D real time Visualization Engine available on
> Mobile Device
> Automotive
> Internet

To provide a real 3D navigation and local search, enhanced with real time information coming from
> Social networking
> Events
> Local based Advertising

Key Topics	3D, Embedded & Mobile, Visualization
Summit	Emerging Companies Summit
Speaker	Eric Bernard (CEO, Visioglobe)

FRIDAY 2:00 PM – 2:30 PM
GLEN ELLEN ROOM

2023 Motus Games Startup Presentation

Motus Games is a leader in Human-Computer Interaction, founded in 2001 by MIT graduates and experienced entreprenuers with the intial goal of creating a motion-based light saber game, clearly an obvious first mission for a bunch of gamers. While the prototype was successful and one of the first if not the first controller of its kind, the world didn't seem ready for motion-based control. So the prototypes went home with gamers and our focus turned to golf, where the team created the industry-leading iClub (www.iclub.net), a wireless peripheral for the golf club that analyzes the golf swing and shows a 3D "X-Ray" of the swing. In 2007, Motus Games returned to its roots with a mission to create and enable the best motion-based user experience for gamers. From the next-generation Darwin hardware to artificial intelligence and analysis middleware (codenamed Graffiti), Motus Games aims to drive the innovation behind how motion-based games are played and give developers the tools needed for fueling the next generation of motion-based interfaces for Visual Computing.

Key Topics	3D, General Interest
Summit	Emerging Companies Summit
Speaker	Satayan Mahajan (CEO, Motus Games)

FRIDAY 2:00 PM – 2:30 PM
CUPERTINO ROOM

2027 Multiverse Startup Presentation

Multiverse is a pioneering interactive entertainment company. In addition to providing a revolutionary development platform to independent creative teams, Multiverse also works closely with leading brands to translate their ideas into engaging online consumer experiences

Key Topic	General Interest
Summit	Emerging Companies Summit
Speakers	Corey Bridges (-); Rafhael Cedeno (-)

FRIDAY 2:00 PM – 2:30 PM
BELVEDERE ROOM

2035 DUX Soft Pvt. Ltd. Startup Presentation

DUX Soft Pvt. Ltd. has been founded to design, integrate and deliver an innovative mix of productivity and performance based products and solutions, for enterprises working with floating point computations or rich media technologies, and in the process establish itself as a leader in the High Productivity Computing, Digital Media Solutions, Artificial Intelligence based Storage & Workflow integration market. DUX serves its niche market by bringing together innovative ideas, scalable technologies, and a top-of-the-line team – with over 50 years of combined experience setting up major animation studios in India, designing and implementing successful pipelines, and managing scalable infrastructures. DUX strives to provide exceptional value and peace of mind to our customers, minimize their costs, maximize resource utilization, and allow them to concentrate freely on their core competencies. At the forefront of DUX's solutions is a unique digital workflow tool called SWIVEL which manages pipelines, rendering and assets and extends to provide digital mining capabilities for new media enterprises. DUX has also acquired a cutting edge Artificial Intelligence & Predictive Engine algorithm which it is using to create a self-learning and self-healing intelligent cache based Storage environment.

Key Topic	General Interest
Summit	Emerging Companies Summit
Speakers	Amit Srivastava (CEO / Chairman, DUX Soft Pvt. Ltd.) and Alex Santo (Consultant, DUX Soft Pvt. Ltd.)

FRIDAY 2:00 PM – 3:00 PM
CRYSTAL ROOM

1068 OpenCL Optimization

In this session, we will discuss how to optimize OpenCL programs on NVIDIA GPUs. Three main aspects are discussed: memory, execution configuration, and instruction throughput. On memory optimization, we will discuss how to increase bandwidth by global memory coalescing and using local memory. Then we will discuss the concept of occupancy and various considerations in specifying the execution configuration of a kernel. Finally, we discuss techniques for improving instruction throughput.

Key Topics	Film, Tools, Visualization, Medical Imaging & Visualization
Summit	GPU Developers Summit
Speaker	Peng Wang (NVIDIA)

FRIDAY 2:00 PM – 3:00 PM
SACRAMENTO ROOM

1073 Programming the Professional Video Pipeline

The Quadro Professional Video Pipeline enables the binding of SDI video input and output devices directly to the GPU. This provides applications with the extraordinary capability to performing capture, processing and then output live to air or other video device completely in video memory on the GPU. Gone are the days of needing a CPU and lots of system memory for video processing. Come learn how to capture SDI video directly to GPU memory for subsequent processing as Video Textures, Video Buffer Objects or in CUDA. Once in the GPU, the video processing possibilities are endless, color grading, transcoding, compositing, restoration, effects etc. Then, output the ultimate quality imagery directly from GPU memory live to air, the network, presentation system or other video recording or display device. This session will present the Quadro Profession Video API and solution stack with case studies and usage examples.

Key Topics	Computational Imaging, Visualization, Tools, Film
Summit	GPU Developers Summit
Speaker	Thomas TRUE (Applied Engineer, NVIDIA)

FRIDAY 2:00 PM – 3:00 PM
GARDEN ROOM

1079 Scalable Multi Agent Simulation on the GPU

Multi agent systems have been recently gaining increased attention by game AI developers, mainly in seeking a robust motion planning solution for non player characters. This session addresses some of the principal design considerations for a scalable graphics hardware implementation of collision avoidance simulation, using NVIDIA's CUDA. In particular, it highlights the tremendous performance payoff of the GPU when exploiting nested data parallelism. It also covers the topic of formalizing the integration of game AI and physics tasks on the GPU to form a collaborating and a more intuitive simulation framework to the game developer. Attendees will acquire a number of highly optimized algorithms in various areas of real-time, multi agent simulation.

Key Topics	Algorithms & Numerical Techniques, Machine Learning & Artificial Intelligence
Summits	GPU Developers Summit, NVIDIA Research Summit
Speaker	Avi Bleiweiss (NVIDIA)

FRIDAY 2:00 PM – 3:00 PM
PIEDMONT ROOM

1098 NEXUS Lab

Experience the future of GPU Computing development today! NEXUS is NVIDIA's new GPU Computing development environment for Microsoft Visual Studio. In this hands-on lab, attendees will learn to use the NEXUS source debugger and profiler to efficiently solve common problems encountered when debugging and profiling CUDA C code, including

> Debugging logic errors using hardware breakpoints and variable inspection

> Reading and visualizing GPU memory directly in the Visual Studio IDE (no manual readbacks required!)

> Real-time expression evaluation on the GPU while your application is halted at a breakpoint

> Hardware-evaluated breakpoint conditionals for debugging a specific block or thread

> Detecting out-of-bounds memory errors

> Finding erroneous memory overwrites using data breakpoints

> Understanding the performance of your whole system using the Nexus Analysis tools

NVIDIA engineers will be available at each session to answer

your NEXUS and general GPU Computing questions. Lab workstations are limited to the first 20 attendees per session. NVIDIA would like to thank HP and Parallels for their generous support of SLI Multi-OS workstations.

Key Topic	General Interest, Tools
Summit	GPU Developers Summit
Speakers	Brian Kretzler (NVIDIA); Greg Smith (Manager of Developer Tools, NVIDIA)

FRIDAY 2:00 PM – 3:00 PM
VALLEY ROOM

1100 Speed Limit Recognition System (Optimized for processing on the GPU)

In this study, we investigate the use of a programmable graphics processing unit (GPU) as an embedded processor for real-time recognition of speed limit signs on the road. The input to our system is a video sequence of the road taken from a moving vehicle. We process this video in real-time and determine if there are any speed limit signs present in the scene and, if so, we recognize and output the number indicated by the sign. The main goal of the recognition system is to operate in real time on a resource-constrained embedded system. Therefore, we first examine the merits and demerits of mapping algorithms often used for speed-limit recognition on to the GPU. Through this process, we find techniques that benefit significantly from the GPU architecture and eliminate algorithms that do not map efficiently on it. We then implement and analyze two sign detection schemes: one feature-based, one template-based. From the results of our experiments, we make several important conclusions about the trade-off between recognition rates and performance. We also make an estimate for the amount of hardware resources needed to perform the recognition in real-time."

Key Topics	Embedded & Mobile, General Interest
Summit	GPU Developers Summit
Speakers	Vladimir Glavtchev (Software Engineer Intern, NVIDIA, University of California, Davis); Pinar Muyan-Ozcelik (PhD Student, UC Davis); Jeff Ota (BMW Technology Office USA); John Owens (UC Davis)

FRIDAY 2:00 PM – 3:00 PM
REGENCY 2 ROOM

1142 Handheld Augmented Reality

Handheld Augmented Reality (AR) presents an exciting opportunity for mobile interaction and game designers. By moving experiences off the screen and out into the physical world, handheld AR has the potential to sidestep the limitations of small mobile displays by giving users the illusion that they are looking through a window into a larger 3D space merged with the world. And by attaching this virtual world to the physical world, AR creates new opportunities for physical and social interaction. In this talk, I will discuss a range of issues that must be dealt with when creating mobile AR experiences, and illustrate these points with a serious of AR game prototypes, with a focus on social, physical, and tangible interaction.

Key Topic	Embedded & Mobile
Summit	NVIDIA Research Summit
Speaker	Blair MacIntyre (Associate Professor, Georgia Tech, GVU Center)

FRIDAY 2:00 PM to 3:00 PM
EMPIRE ROOM

1440 Computational Based Modelling and Simulation

The impact of computational based Modelling and Simulation in providing industry with a competitive edge delivering higher performance and cost effective products is well documented and cannot be under estimated. This has, in the past been enabled by the year on year floating point performance improvements delivered by the processor and system manufacturers and has been the primary end user experience of Moore's law. The development of MultiCore technologies in all its wide variety of forms challenges this end user effect of Moore's law requiring a radical rethink of algorithms and their implementations just to maintain the same progress seen over the previous decades. The exploitation of GPU Computing promises to deliver a radical and affordable step change in simulation capability and this will be illustrated with details of a Computational Fluid Dynamics (CFD) solver (3D Finite Volume Unstructured) that demonstrates 20x performance benefits over the latest Intel Nehalem cores.

Key Topic	General Interest
Summit	GPU Developers Summit
Speaker	Jamil Appa (Aeronautical Engineer, BAE Systems)

FRIDAY 2:30 PM – 3:00 PM
ATHERTON ROOM

1059 Fast Tridiagonal Solvers on GPU

We study the performance of three parallel algorithms and their hybrid variants for solving tridiagonal linear systems on a GPU: cyclic reduction (CR), parallel cyclic reduction (PCR) and recursive doubling (RD). We develop an approach to measure, analyze, and optimize the performance of GPU programs in terms of memory access, computation, and control overhead. We find that CR enjoys linear algorithm complexity but suffers from more algorithmic steps and bank conflicts, while PCR and RD have fewer algorithmic steps but do more work each step. To combine the benefits of the basic algorithms, we propose hybrid CR+PCR and CR+RD algorithms, which improve the performance of PCR, RD and CR by 21%, 31% and 61% respectively. Our GPU solvers achieve up to a 28x speedup over a sequential LAPACK solver, and a 12x speedup over a multi-threaded CPU solver.

Key Topic	Algorithms & Numerical Techniques, Astro GPU, Computational Fluid Dynamics , Tools
Summits	GPU Developers Summit, NVIDIA Research Summit
Speaker	Yao Zhang (University of California, Davis)

FRIDAY 2:30 PM – 3:00 PM
GOLD ROOM

1085 Debugging Tools for CUDA

Explore the ease of debugging heterogeneous applications using the integrated CPU and GPU debugging environment of CUDA-GDB. CUDA-GDB is NVIDIA's very first native GPU debugger that lets you debug CUDA applications on the Linux

platforms. In this session you will learn about the capabilities of the CUDA-GDB debugger with a walk through of the simplicities of being able to debug complex parallel applications. CUDA-GDB is also a platform to built custom debuggers so if you plan to build a CUDA debugger this session will provide the details on building your very own CUDA debugger.

Key Topic	General Interest, Tools
Summit	GPU Developers Summit
Speaker	Satish Salian (CUDA Tools Manager, NVIDIA)

FRIDAY 2:30 PM – 3:00 PM
CALIFORNIA ROOM

1112 Visualizing the Universe: Raycasting Astrophysical Simulation Data

We use GPU-assisted raycasting to render large, three-dimensional time-dependent astrophysical AMR data sets at interactive frame rates on standard desktop computers. Our approach allows us to embed unstructured point datasets, like stars or galaxy splats, into the rendering of gaseous interstellar or intergalactic material. The approach supports a combined color-mapping of several input data fields and allows for a very flexible adaption to the special requirements of different types of simulations. Its interactivity makes it a useful tool for data analysis as well as for fast generation of high-quality animations from astrophysical datasets. We will show various resulting animations ranging from large scale structure formation in the early universe, to the evolution of the first stellar object and the cosmological reionization era. Finally, we will give an overview about lessons learned and opportunities for future work.

Key Topic	Astro GPU
Summits	GPU Developers Summit, NVIDIA Research Summit
Speaker	Ralf Kaehler (KIPAC/SLAC)

FRIDAY 2:30 AM – 3:00 AM
HILLSBOROUGH ROOM

1458 Raster Processing Using Multiple GPUs and CPUs

With the ever-increasing availability of satellite data streams,affordable storage space and communication bandwidth, there is a strong demand for processing increasingly large raster data sets in GIS and Remote Sensing applications. Although available computing power also increases day to day, these increases come mostly in the form of additional processing units. Utilizing these additional units effectively requires significant modifications to data processing algorithms. This talk describes an approach to creating a generalized pipeline for processing raster data that scales with the number of GPUs as well as the number of CPUs, and is suitable for implementing most of the classic raster algorithms.

Key Topic	General Interest
Summit	GPU Developers Summit
Speaker	Dimitri Rotow (Product Manager, Manifold)

FRIDAY 2:30 PM – 3:00 PM
FAIRFIELD ROOM

2016 Cooliris Startup Presentation

Cooliris, Inc. was founded in January 2006 with a simple mantra: "Think beyond the browser". We focus on creating products that make discovering and enjoying the Web more exciting, efficient, and personal. Our core products include Cooliris (formerly PicLens), which transforms your browser into an interactive, full-screen "cinematic" experience for web media, and CoolPreviews, which lets you preview links instantly. Cooliris has reached 12 million installs of the product, with tens of thousands downloading everyday. The latest version is available for the PC, Mac, and Linux platforms, and can be downloaded at www.cooliris.com. Headquartered in Palo Alto, California, the Cooliris team consists of seasoned executives, engineers and entrepreneurs from companies like Apple, eBay, Electronic Arts, Google, Microsoft, PayPal and YouTube. For more information, please visit www.cooliris.com.

Key Topic	General Interest
Summit	Emerging Companies Summit
Speaker	Austin Shoemaker (CTO, Cooliris)

FRIDAY 2:30 PM – 3:00 PM
BELVEDERE ROOM

2026 Foster Findlay Associates Startup Presentation

ffA originates and provides world leading 3D seismic analysis software and services to the oil and gas industry. Application of ffA Technology accelerates 3D seismic interpretaton, improves understanding of exploration risk and enables optimal hydrocarbon exploitation strategies to be defined.

Key Topics	3D, General Interest
Summit	Emerging Companies Summit
Speaker	Steve Purves (Technical Director, ffA)

FRIDAY 2:30 PM – 3:00 PM
GLEN ELLEN ROOM

2037 Allegorithmic Startup Presentation

Allegorithmic is the first company to propose a professional middleware for the authoring and on-the-fly rendering of textures. Substance allows content developers to produce texture assets twice as fast as usual, while Substance description files are typically 500-1000 times smaller than regular textures. Allegorithmic's new product "Substance" is poised to redefine the development and distribution of rich content for the next generation of online games. Allegorithmic is currently partnering with the most innovative companies of the field to provide industry innovators with the most advanced packages for content production.

Key Topic	Algorithms & Numerical Techniques
Summit	Emerging Companies Summit
Speaker	Sebastien Deguy (CEO, Allegorithmic)

FRIDAY 3:00 PM – 3:30 PM
SACRAMENTO ROOM

1114 Strategies for GPU Acceleration of Common Visual Effects for Broadcast and Post-Production

Since 1995 BorisFX has developed image processing and 3D graphics software for Broadcast and Post-Production, with a particular focus on staple visual effects such as 3D Text, Chroma Key, and Film Look. While achieving award-winning quality, these CPU bound effects have lacked the interactivity today's customers demand. With the advance of GPU hardware from NVIDIA the ability to accelerate these effects has become

a reality, and allowed for many previously impractical features to present themselves. In this presentation we will demonstrate how the use of the GPU has benefited our products in terms of performance and features in our creation of GPU-Rendered 3D Text, and discuss the strategies we employed to emphasize the benefits and minimize the drawbacks of the GPU in building the multi-pass Chroma Key and Film Look filters.

Key Topic	3D, Visualization, Film
Summit	GPU Developers Summit
Speakers	Boris Yamnitsky (BorisFX); David Yamnitsky (Engineer, Boris FX)

FRIDAY 3:00 PM – 3:30 PM
GOLD ROOM

1140 GPU Debugging Made Easy with DDT

One of the key components for exploiting the potential of GPUs is the debugging tool – the software through which problems with your GPU code can be resolved quickly and easily. Allinea Software's DDT is a powerful debugger for multi-threaded and parallel codes, known for its ease of use and scalability on large systems and applications. We will introduce Allinea DDT for NVIDIA GPUs which brings debugging of CUDA on real hardware – providing intuitive ways to control, visualize and debug CUDA threads.

Key Topic	Tools
Summit	GPU Developers Summit
Speaker	David Lecomber (CTO, Allinea Software)

FRIDAY 3:00 PM – 3:30 PM
BELVEDERE ROOM

2013 Forterra Systems, Inc Startup Presentation

Forterra provides enterprise grade virtual world software that enables our customers and partners to deploy their own secure, scalable 3D Internet solutions for the corporate, healthcare, government, and education industries. Our software and services enable organizations to collaborate, train, practice, and plan in compelling 3-D virtual worlds that integrate with existing media and business systems. Since 1998, the company has used its unique combination of experience in enterprise software, social networks, video game production, modeling and simulation, real-time graphics, and artificial intelligence to deliver innovative solutions around collaboration, training, and operations applications. Forterra's flagship product OLIVE (On-Line Interactive Virtual Environment) is a powerful client-server based virtual world platform that connects users in real-time over any network. Within the virtual world, users are represented by realistic 3D avatars that create natural human interaction and discussion. Independent of scale, solutions built on the OLIVE platform can be integrated with existing business systems and media, 3-D content, and data sets to create powerful collaborative environments. OLIVE's architecture is mature and proven, which is an important consideration for customers and partners making long term investment decisions. Forterra is a trusted business partner aligned with key industry leaders like IBM, Accenture, Nvidia, and Autodesk. www.forterrainc.com

Key Topic	3D, Visualization
Summit	Emerging Companies Summit
Speaker	David Rolston, PhD (CEO, Forterra Systems, Inc)

FRIDAY 3:00 PM – 3:30 PM
GLEN ELLEN ROOM

2019 Aqumin LLC Startup Presentation

Aqumin, founded in 2008, builds financial market interpretation software called AlphaVision. Using 3D visualization, sophisiticated data-aggregation, and high-end analysis algorithms, AlphaVision enables market paticipants to identify activity in financial markets that would otherwise be difficult or impossible to detect. Aqumin's technology platform provides seamless integration of large data streams from first tier providers such as ThomsonReuters, Bloomberg, NYSE, Nasdaq, ActiveFinancial and others; covering securities from all over the world and combining them into an interactive 3D visual and computational environment. Customers benefit by being able to bring diverse information together to answer questions about securites markets and navigate the results in global context. Much as Google Earth has enabled new classes of GIS applications through the use of XML APIs and mashups, Aqumin's AlphaVision is enabling new paradigms in the way professionals look at equity and equity deriviatives trading, risk management, order execution & buy-right, as well fundamental research. AlphaVision is sold as a monthly subscription service or as an add-on to the Bloomberg Terminal (see REG<GO>). Headquartered in Houston, Texas with a presence in New York and Chicago, Aqumin is currently generating revenue and is presently raising captial to fund expanded marketing and sales opportunities in the United States and select markets globally.

Key Topics	3D, Computational Finance
Summit	Emerging Companies Summit
Speaker	Sean Spicer (-)

FRIDAY 3:00 PM – 3:30 PM
FAIRFIELD ROOM

2055 Universal Robotics, Inc. Startup Presentation

The company's first product, Neocortex, to be sold in 2010, is the control system for a robotic box mover in the materials handling industry, a multi-billion dollar market . Neocortex enables robots to stack and unload pallets with mixed size boxes, a capability prospective customers have said is the "holy grail of automated material handling". Large distribution centers such as The Gap's, Gallatin, TN facility (a Neocortex alpha test site) will see significant reductions in material handling costs due to improvements in the speed of operation, a reduction in human errors, and increased safety and efficiency in warehouse operations, resulting in lowered green house gas emissions. Furthermore, because Neocortex is only the software, it will radically expand the use of automation into new hardware applications such as forklifts, mining equipment, and railroad locomotives. Until now, these applications have been unable to capitalize on automation efficiencies due to the current limitations of Artificial Intelligence.

Key Topics	Machine Learning & Artificial Intelligence
Summit	Emerging Companies Summit
Speakers	Alan Peters (CTO, Universal Robotics, Inc.); David Peters (CEO, Universal Robotics, Inc.)

FRIDAY 3:00 PM – 3:30 PM
CUPERTINO ROOM

2056 Kerosene and a Match Startup Presentation

Based in Orange County, CA, Kerosene and a Match (KaaM) is a software developer building tools that leverage the massively parallel, low cost computing power of commodity graphics processors to build ultra-high performance cloud computing platforms. The company's first product is a new multimedia discovery and indexing platform that combines GPU and cloud computing to give standard commercial search systems the ability to "discover" the content of "unsearchable" (aka, un-indexed and/or untagged by people) images, audio and video, and return the results just as it would text. Because upwards of 80% of all non-text content on the Internet is unsearchable, the KaaM technology offers organizations the ability to generate income from huge volumes of formerly unmonetized media. KaaM licenses its technology to enterprises, government agencies and other organizations with large scale networks housing enormous amounts of rich media. The company also plans to deliver a software as a service (SaaS) offering for smaller entities late in 2010.

Key Topic	General Interest
Summit	Emerging Companies Summit
Speaker	Cleve Adams (CEO, KaaM)

FRIDAY 3:00 PM – 4:00 PM
VALLEY ROOM

1002 Accelerating Energy Exploration and Production with Seismic Processing

The energy industry has always been at the forefront of technology and high-performance computing. Today, GPU Computing offers another step-change in the preceding benefits because high-capacity, high-throughput and massively parallel GPUs are well-applied to the compute intensive problems currently solved or proposed by the energy exploration and production industry. This talk will show how using GPUs for problems like seismic migration-particularly Kirchhoff time migration ("KTM" or "PSTM") and reverse time migration ("RTM")-and reservoir simulation can reduce the runtime of today"s important problems by an order of magnitude (10X) or more. In addition to the performance that the presenter has seen, in their customers' experience, it will also be possible to discuss how other organizations in the energy industry are harnessing GPUs.

Key Topic	Energy Exploration
Summit	GPU Developers Summit
Speaker	Ryan Schneider (Chief Technology Officer, Acceleware)

FRIDAY 3:00 PM – 4:00 PM
ATHERTON ROOM

1060 Graphcuts with CUDA and Applications in Image Processing

Abstract: Graph Cuts is a powerful and popular optimization approach to solve image processing problems such as image segmentation, stereo vision, image restoration and many more. In this talk, we present CUDA implementations of the push-relabel algorithm to compute Graph Cuts. Starting from the basic algorithm we discuss its parallel processing properties. Then different optimization strategies are explored and their strengths and weaknesses are evaluated. We conclude by exploring applications of Graph Cuts to solve image processing problems using GPUs.

Key Topics	Computational Imaging, Medical Imaging & Visualization
Summits	GPU Developers Summit, NVIDIA Research Summit
Speaker	Timo Stich (Developer Technology Engineer, NVIDIA)

FRIDAY 3:00 PM – 4:00 PM
CALIFORNIA ROOM

1065 Diesel-Powered GPU Computing: Enabling a Real-Time Radio Telescope in the Australian Outback

The Murchison Widefield Array (MWA) is a next-generation radio telescope currently under construction in the remote Western Australia Outback. The raw data rate is 5 to 20 GiB/sec, precluding offline processing. Since the computing budget for calibration and imaging is 20 TFLOP/sec, a real-time high-performance computer is required on-site. We describe a scalable heterogeneous computing pipeline implementation, exploiting both the high computing density and FLOP-per-watt ratio of modern GPUs. The architecture is highly parallel within and across nodes, with all major processing elements performed by the GPUs. Necessary scatter-gather operations along the pipeline are loosely synchronized and implemented in MPI. Our initial port to NVIDIA hardware shows a typical 10x improvement over the reference CPU implementation, with some portions showing even more substantial gains. The MWA will be a frontier scientific instrument and a demonstrator for planned peta- and exascale facilities.

Key Topic	Astro GPU
Summits	GPU Developers Summit, NVIDIA Research Summit
Speaker	Richard Edgar (Initiative in Innovative Computing, Harvard University)

FRIDAY 3:00 PM – 4:00 PM
CRYSTAL ROOM

1069 Zombies on Tegra: A Case Study in Mobile Augmented Reality

The NVIDIA Tegra family of ARM-based mobile application processors integrates extensive visual processing acceleration for amazing graphics, video, image processing and audio that can be combined in powerful ways. This presentation features the ARhrrrr augmented reality zombie game from Georgia Tech to illustrate how Tegra can seamlessly recognize features from a handheld camera image and composite the video with synchronized 3D graphics. The power of Tegra is accessed through open standard APIs including OpenGL ES, OpenMAX IL and OpenKODE contained in the Tegra software development kit. This session will provide real-world insights to demonstrate how leading-edge mobile silicon and API standards come together to create a powerful and compelling mobile visual computing platform.

Key Topics	Computer Vision, Visualization, Embedded & Mobile
Summit	GPU Developers Summit
Speakers	Lars Bishop (Engineer, NVIDIA); Blair MacIntyre (Associate Professor, Georgia Tech, GVU Center)

FRIDAY 3:00 PM – 4:00 PM
EMPIRE ROOM

1074 Computational Fluid Dynamics (CFD) for the GPU

The field of computational fluid dynamics (CFD) has far-reaching applications and displays a consistent need for larger and faster simulations. At EM Photonics we have been studying this field and its computational needs for two years. We have identified the GPU as a strong performer in the CFD field and as such have implemented solvers that harness the power of GPUs in the application of CFD formulations. We will present some background on these innovations in this summary discussion.

Key Topics	Astro GPU, Computational Fluid Dynamics, Physics Simulation
Summit	GPU Developers Summit
Speakers	John Humphrey (Senior Engineer, EM Photonics); Daniel Price (EM Photonics)

FRIDAY 3:00 PM – 4:00 PM
GARDEN ROOM

1080 The Art of Performance Tuning for the CUDA Architecture

The CUDA architecture and programming model has proven very successful in allowing programmers to easily express and accelerate fine-grained parallelism in their codes. CUDA owes much of its success to how quickly most programmers get a substantial speedup in their initial porting efforts. Like any architecture, however, extracting absolute maximum performance can be more involved. In this moderated roundtable NVIDIA engineers will lead a discussion on the art of understanding and optimizing performance of codes written for the CUDA architecture. Presenters and attendees will share their strategies, experiences, and questions for performance tuning on the CUDA architecture.

Key Topics	General Interest, Tools
Summit	GPU Developers Summit
Speaker	Paulius Micikevicius (NVIDIA)

FRIDAY 3:00 PM – 4:00 PM
HILLSBOROUGH ROOM

1090 Spatial Data Structures for Massively Parallel Computing

Many applications in physical simulation and rendering require the use of spatial data structures to achieve optimal performance. Spatial data structures can be used to accelerate collision detection, ray tracing and molecular dynamics simulations. This talk will describe how data structures such as uniform grids and bounding volume hierarchies can be efficiently constructed and traversed on the GPU using data parallel primitives such as sort and scan.

Key Topics	Algorithms & Numerical Techniques, Tools, Physics Simulation
Summits	GPU Developers Summit, NVIDIA Research Summit
Speaker	Richard Tonge (Senior Software Engineer – PhysX, NVIDIA)

FRIDAY 3:00 PM – 4:00 PM
PIEDMONT ROOM

1098 NEXUS Lab

Experience the future of GPU Computing development today! NEXUS is NVIDIA's new GPU Computing development environment for Microsoft Visual Studio. In this hands-on lab, attendees will learn to use the NEXUS source debugger and profiler to efficiently solve common problems encountered when debugging and profiling CUDA C code, including

> Debugging logic errors using hardware breakpoints and variable inspection

> Reading and visualizing GPU memory directly in the Visual Studio IDE (no manual readbacks required!)

> Real-time expression evaluation on the GPU while your application is halted at a breakpoint

> Hardware-evaluated breakpoint conditionals for debugging a specific block or thread

> Detecting out-of-bounds memory errors

> Finding erroneous memory overwrites using data breakpoints

> Understanding the performance of your whole system using the Nexus Analysis tools

NVIDIA engineers will be available at each session to answer your NEXUS and general GPU Computing questions. Lab workstations are limited to the first 20 attendees per session. NVIDIA would like to thank HP and Parallels for their generous support of SLI Multi-OS workstations.

Key Topics	General Interest, Tools
Summit	GPU Developers Summit
Speakers	Brian Kretzler (NVIDIA); Greg Smith (Manager of Developer Tools, NVIDIA)

FRIDAY 3:00 PM – 4:00 PM
REGENCY 2 ROOM

1123 Domain Specific Languages for Programming GPUs

Processors and computing systems are becoming very heterogenous. Imagine writing a program that runs on a cluster, a shared-memory multi-core processor, and a many-core GPU. In each case the application needs to be completely rewritten. The cluster will use MPI, the shared-memory system threads and locks, and GPU a data-parallel programming model like CUDA. The cost of maintaining three different versions of a program is often prohibitive, and limits the adoption of GPUs in many applications. My thesis is that the only way to program such diverse set of systems is to raise the level of abstraction and build domain-specific programming models that are able to automatically map applications to different types of platforms. Our goal is to make it much easier to adopt innovative new platforms. In this talk, I will present the case for domain-specific programming models, and outline several projects that we are doing at Stanford as part of the Pervasive Parallelism Laboratory.

Key Topic	General Interest
Summit	NVIDIA Research Summit
Speaker	Pat Hanrahan (Professor, Stanford University)

FRIDAY 3:30 PM – 4:00 PM
BELVEDERE ROOM

2017 Canesta Startup Presentation

Canesta is the inventor of revolutionary, low-cost electronic perception technology that enables ordinary electronic devices in consumer, security, industrial, medical, automotive, factory automation, entertainment, military, and many other applications, to perceive and react to objects or individuals in real time. When given true, fine-grained 3-dimensional depth perception with Canesta's unique CanestaVision electronic perception chips and software, such products can gain functionality and ease of use not possible in an era when such devices were blind. Numerous applications are under active development by Canesta's OEM customers and partners, including building automation, security, robotics, automotive, and others. Canesta was founded in April 1999, and is located in San Jose, CA. The company has 40 granted patents and counts multiple Fortune 100 companies as its customers. All trademarks and registered trademarks are those of their respective companies. The omission of a trademark or registered trademark symbol is inadvertent and is not intended to convey the status of any mark or contravene any claim.

Key Topics	3D, Visualization, General Interest
Summit	Emerging Companies Summit
Speaker	Jim Spare (President & CEO, Canesta, Inc.)

FRIDAY 3:30 PM – 4:00 PM
GLEN ELLEN ROOM

2039 HPC PROJECT Startup Presentation

HPC Project, under the Wild Systems brand, provides application-in-a-box solutions for demanding users who require intensive computational power. The company brings to the market the cumulative experience of experts in simulation, parallel computing and computer architecture. The company is dedicated to package turnkey products, providing outstanding computing performance in a desktop form factor for interactive use of complex dataset. Wild Systems Solutions are supported by best-of-breed technologies representing more than 25 years of R&D in the supercomputing field. This technology is based upon a formal analysis of source codes enabling correct automated transformations. These techniques are used to detect inherent parallelism in applications that could then be optimized to take advantage of the latest chipset architecture. Such a technology happens to be particularly handy when dealing with specialized architectures such as GPU. These circuits require a particular attention in the organization of data in memory and of communication between on-board memory and main memory. Hand-coding this type of structures could be very tedious and error-prone. Our compiler framework provides in this context the necessary alignment of data and the pipelining of their communication for a high-performance execution.

Key Topic	General Interest
Summit	Emerging Companies Summit
Speaker	Pierre FIORINI (CEO, HPC PROJECT)

FRIDAY 3:30 PM – 4:00 PM
CUPERTINO ROOM

2049 SceneCaster Startup Presentation

SceneCaster is a leading provider of 3D social media applications that enrich the online consumer experience and enable highly immersive and contextual advertising opportunities. SceneCaster's applications reach millions of users around the world and extend the reach and accessibility of 3D applications to a mass market of social networking and social media users. SceneCaster's mass market appeal and immersive 3D Web technology offers advertisers a unique, interactive and entirely performance-based advertising solution that reaches a global audience of highly engaged users.

Key Topic	3D
Summit	Emerging Companies Summit
Speaker	Alain Chesnais (Chief Technology Officer, SceneCaster)

FRIDAY 3:30 PM – 5:00 PM
SACRAMENTO ROOM

1447 The GPU Revolution in Film Production: Report from the Battlefield

Come join our panel of renowned industry experts working at major US film studios including Disney Feature Animation, Sony Pictures Imageworks and Digital Domain as they discuss and present their latest work from feature films and how the GPU is changing the way they make films by solving complex technical challenges and enhancing production pipeline worldwide.

Key Topic	Film
Summit	GPU Developers Summit
Speaker	Rob Bredow (Chief Technology Oficer, Sony Imageworks); Arthur Shek (Manager, Technology, Walt Disney Animation Studios); Chris Horvath (DigitalArtist, Industrial Light & Magic)

FRIDAY 3:30 PM – 5:30 PM
GOLD ROOM

1086 Advanced C for CUDA

This presentation covers the major CUDA optimizations. Topics will include: maximizing memory throughput, kernel launch configuration, using shared memory, and improving GPU/CPU interaction. We will also describe how the CUDA Visual Profiler and code instrumentation can be used to assess performance. While C for CUDA is used for illustration, the concepts covered will apply equally to programs written with OpenCL and DirectCompute APIs.

Key Topic	Tools
Summit	GPU Developers Summit
Speaker	Thomas Bradley (NVIDIA)

FRIDAY 4:00 PM – 4:30 PM
CRYSTAL ROOM

1070 Face Recognition for Photographs and Video

Face Recognition in Video (FRiV) is a challenging task. Successful solution of FRiV problem requires state of the art algorithms for visual analysis as well as their very optimized implementation in order to make the whole system working on

off the shelf hardware. Presentation shows how computational performance of Viewdle FRiV technology is improved by harnessing power of GPU based computing using NVidia CUDA technology. Transition from CPU based implementation to GPU implementation is described. Changes in computing model made for efficient execution of computationally intensive technology parts on GPU hardware are presented. Essential improvement of processing speed of Viewdle FRiV technology based on GPU computing is demonstrated. Problems arisen during transition are analyzed. Features needed from GPU to make FRiV technology working even faster are presented.

Key Topics	Computational Imaging, Computer Vision
Summit	GPU Developers Summit
Speaker	Yuriy Musatenko (Viewdle)

FRIDAY 4:00 PM – 4:30 PM
EMPIRE ROOM

1075 Reconstructing the Brain: Extracting Neural Circuitry with CUDA and MPI

Determining the detailed connections in brain circuits is a fundamental unsolved problem in neuroscience. Understanding this circuitry will enable brain scientists to confirm or refute existing models, develop new ones, and come closer to an understanding of how the brain works. Prof. Jeff Lichtman and Center for Brain Science (CBS) at Harvard launched the Connectome Project three years ago to determine the complete, detailed wiring diagrams of neural circuits from sequential high-resolution images of the central nervous system using electron microscopy (EM). These high-resolution, large-scale EM datasets pose very challenging computational problems for 3D segmentation and visualization in terms of developing suitable algorithms, coping with the ever-increasing data sizes, and maintaining interactive performance. <p> In this talk we will present our insights and lessons learned in using CUDA to reconstruct neural connections in high-resolution EM data. We will present technical details and non-trivial issues regarding the implementation of NeuroTrace, our system for semi-automatic segmentation and interactive visualization of terabytes of EM image data. The segmentation method is based on a sequence of 2D level set segmentations of cell membranes integrated with an image correspondence energy for robust transition between consecutive slices and a weighted path extrapolation method to trace a 3D centerline of a neural pathway along non-axis aligned slices. The visualization method employs view-dependent on-demand filtering and evaluation of a local histogram edge metric, as well as on-the-fly interpolation and ray-casting of implicit surfaces for segmented neural structures. We will describe the implementation details of both methods in CUDA for interactive performance. We will also present a novel client-server implementation of NeuroTrace that uses distributed processing with MPI on a CUDA compute cluster with eight S1070 Tesla and 16 compute nodes.

Key Topics	Computational Imaging, Medical Imaging & Visualization, High Performance Computing, Life Sciences
Summits	GPU Developers Summit, NVIDIA Research Summit
Speaker	Won-Ki Jeong (Initiative in Innovative Computing, Harvard University)

FRIDAY 4:00 PM – 4:30 PM
ATHERTON ROOM

1109 Building a Modern Engine for the Visual Simulation Industry

Blue Newt Software has been designing visual software for more than 5 years, and in late 2008 won a contract to redesign the visual rendering engine for a series of driving simulators for a large German automotive company. This company has both very high quality demands and firm performance constraints making this a challenging problem. This session will discuss the development process and graphics technologies employed in our modern visual rendering and simulation engine. To meet these needs we designed a component-based visual engine, a deferred-rendering graphics engine, and a fully-scriptable API. We provide support for full high-dynamic range rendering, a series of post-processing effects, and a data production toolchain for data ingest and usage. We use modern graphics rendering techniques including instancing, shading, buffered objects, and more. We'll describe the problem domain, our solution and product, and the technologies we employed in our approach. We'll describe how we took the requirements of an industry-leading customer and created a product ready for their future. We'll also detail how we researched and employed best-practices from the game industry and applied them to the tasks and requirements of the visual simulation industry.

Key Topics	Embedded & Mobile, Physics Simulation, Visualization
Summit	GPU Developers Summit
Speaker	Bob Kuehne (Blue Newt Software)

FRIDAY 4:00 PM – 4:30 PM
VALLEY ROOM

1141 Energy Industry Applications: Three Examples of GPU Acceleration in the Oil and Gas Industry Using CUDA

The energy industry has traditionally been one of the largest market sectors for High Performance Computing hardware, software, algorithms and related technologies. Many of the critical industry applications map well to GPU hardware and can benefit from the high floating point capability and lower price/performance ratios. This talk will address three representative industry algorithms and their mapping to the GPU i)Kirchhoff time migration for seismic processing ii) Sparse matrix solvers for reservoir simulation and iii) the split-step Fourier transform. An overview of these algorithms will be provided along with a discussion of their GPU implementations, reference performance numbers and a general discussion of price/performance comparisons for the GPU vs. traditional implementations.

Key Topics	Algorithms & Numerical Techniques, Energy Exploration, High Performance Computing
Summit	GPU Developers Summit
Speaker	William Brouwer (Senior Physicist, Stone Ridge Technology)

FRIDAY 4:00 PM – 5:00 PM
HILLSBOROUGH ROOM

1091 Clusters with GPUs under Linux and Windows HPC

Clusters with GPUs are becoming a popular choice in High Performance Computing. This talk will describe the steps

necessary to deploy clusters with GPUs both under Linux and Windows Server HPC,ranging from working with headless compute nodes to the interaction with the queuing system.

Key Topic	High Performance Computing
Summit	GPU Developers Summit
Speakers	Calvin Clark (Microsoft); Massimiliano Fatica (Senior Applied Engineer, NVIDIA)

FRIDAY 4:00 PM – 5:00 PM
PIEDMONT ROOM

1098 NEXUS Lab

Experience the future of GPU Computing development today! NEXUS is NVIDIA's new GPU Computing development environment for Microsoft Visual Studio. In this hands-on lab, attendees will learn to use the NEXUS source debugger and profiler to efficiently solve common problems encountered when debugging and profiling CUDA C code, including

> Debugging logic errors using hardware breakpoints and variable inspection

> Reading and visualizing GPU memory directly in the Visual Studio IDE (no manual readbacks required!)

> Real-time expression evaluation on the GPU while your application is halted at a breakpoint

> Hardware-evaluated breakpoint conditionals for debugging a specific block or thread

> Detecting out-of-bounds memory errors

> Finding erroneous memory overwrites using data breakpoints

> Understanding the performance of your whole system using the Nexus Analysis tools

NVIDIA engineers will be available at each session to answer your NEXUS and general GPU Computing questions. Lab workstations are limited to the first 20 attendees per session. NVIDIA would like to thank HP and Parallels for their generous support of SLI Multi-OS workstations.

Key Topics	General Interest, Tools
Summit	GPU Developers Summit
Speakers	Brian Kretzler (NVIDIA); Greg Smith (Manager of Developer Tools, NVIDIA)

FRIDAY 4:00 PM – 5:00 PM
GARDEN ROOM

1105 SQL/XML-IMDBg: A GPU In-Memory Database and Query Co-Processor

Attendees will learn about the architectural design decisions and implementation strategies chosen in migrating an in-memory database system and query execution engine to execute database workloads concurrently on CPU and GPU hardware. We will also give insights into the optimizer structure and query executer for executing the "split-work" plan simultaneous on CPU and GPU hardware.

Key Topic	Databases & Data Mining
Summit	GPU Developers Summit
Speaker	Harald Frick (CEO, QuiLogic)

FRIDAY 4:00 PM – 5:00 PM
CALIFORNIA ROOM

1066 Astronomy & Astrophysics

Astronomy, astrophysics, and cosmology face very large computational challenges. The physical accuracy of simulations depends on computing capacity, often with steep scalings. Moreover, astronomical observing systems that are in design and construction involve data rates of tens and hundreds of gigabits per second. Facilities like the Square Kilometer Array (SKA) telescope, in 2020, will dwarf even these, generating exabytes of data each day. These computational challenges push well beyond the capacity of traditional approaches, and solutions that draw upon high-efficiency parallel implementations and lower power consumption are required. This roundtable will discuss the use of GPU computing in the field of astronomy and astrophysics. This will include current GPU-enabled research areas, as well as potential new areas within the field.

Key Topic	Astro GPU
Summit	NVIDIA Research Summit, GPU Developers Summit
Speaker	Christopher Harris (Research Associate, The University of Western Australia); Allice Quillen (Associate Professor, University of Rochester)

FRIDAY 4:30 PM – 5:00 PM
CRYSTAL ROOM

1071 Face Detection

The task of face detection in various types of multimedia content has emerged quickly during the last decade to become essential element in many consumer and specialized applications, hardware, SoC. Traditionally the task of face detection has been executed on CPU; however the algorithms haven't changed significantly during the last few years. The latest trend of moving computationally expensive general purpose calculations to GPU couldn't have no influence on the particular task of Face Detection, which resulted in appearance of specialized Computer Vision libraries written for GPUs.

Key Topics	Computational Imaging, Computer Vision
Summit	GPU Developers Summit
Speaker	Anton Obukhov (Developer Technology Engineer, NVIDIA)

FRIDAY 4:30 PM – 5:00 PM
EMPIRE ROOM

1076 Optimizing Ion Channel Kinetics Using A Massively Parallel Genetic Algorithm on the GPU

Voltage-gated ion channels effect the integration of information in many neurons. Some neurons express over 10 voltage-gated channels that turn information processing into a highly non-linear affair. The only path allowing dissection of the role of each channel in neuronal physiology passes through kinetic modeling of the investigated channels. The currently popular analysis techniques suffer from various shortcomings that limit the ability of the researcher to rapidly produce physiologically relevant models of voltage-gated ion channels. We have recently shown that combining a stochastic search algorithm with ionic currents measured following multiple voltage-clamp protocols enables to semi-automatically constrain models of voltage-gated ion channels. This approach, while

numerically successful, is highly demanding computationally with optimization lasting typically for several days on a 160 CPU Linux cluster. To solve this computational bottleneck we have been converting our optimization algorithm to work on a GPU using CUDA. We have succeeded to parallelize the process on a GTX 295 giving a speed increase of roughly X100 over that of the CPU. This shortened simulation times to hours instead of days. More importantly, it allows interactive modification of the modeling process by the user. Once fully automated and expanded to machines containing several GPU cards, this approach may be used for rapid online analysis of voltage-gated currents. This in turn will allow researchers to greatly shorten the time and costs required for building models of neuronal physiology, which will facilitate our understanding of neuronal physiology.

Key Topics	Algorithms & Numerical Techniques, Physics Simulation, Visualization
Summits	GPU Developers Summit, NVIDIA Research Summit
Speaker	Alon Korngreen (Senior Lecturer of Cellular Neurophysiology, Bar-Ilan University)

FRIDAY 4:30 PM – 5:00 PM
VALLEY ROOM

1102 4D Volume Rendering

With advances in image acquisition and numerical simulation techniques, fields ranging from medical imaging to astrophysics are producing data with very large spatial and temporal resolutions. Interactive visualization techniques are crucial to understand and isolate features from the resulting large time dependent 4D volumetric data. This presentation explores the various rendering methods such as texture slicing, raycasting in graphics and cuda as well as hybrid approaches showing their promises and pitfalls. It is common for 4D data to exceed the graphics memory capabilities and approaches for efficiently streaming data such as PBO's and CPU/GPU asynchronous modes are explained. We conclude with a discussion on how other related solutions from NVIDIA can be integrated, specifically focusing on 3D Vision stereo and NVScale middleware to harness multiple GPU's for distributed rendering.

Key Topics	3D, Visualization, Medical Imaging & Visualization, Film, Energy Exploration
Summits	GPU Developers Summit
Speaker	Shalini Venkataraman (NVIDIA)

FRIDAY 4:30 PM – 5:00 PM
ATHERTON ROOM

1111 Industrial Simulation Using Massively Parallel Algorithms

The purpose of the talk is to show that there is an exciting industrial market for next generation supercomputing based tools, and what direction this is going in. This talk is a result of a 5 year research within the Oil&Gas sector, identifying new areas where Supercomputing will enhance existing use cases in a game changing manner. The cases are grounded in real needs from operational environments at the world leading Oil&Gas operator StatoilHydro.

Key Topics	Energy Exploration, Visualization, High Performance Computing
Summit	GPU Developers Summit
Speaker	Thorolf Tonjum (Director of R&D, Stormfjord)

VISUALIZE **GREEN**

VISUALIZE A GREEN EVENT

Place trash and recyclables in proper bins

Use public transportation during the show

In hotel, decline new sheets and towels

Also, unplug phone and laptop chargers

Offset your travel at www.cool-it.us/

Take only collateral/giveaways you will use

What We're Doing

- 100% of greenhouse gas generated is offset
- Composting and recycling
- Producers and vendors agree to green guidelines
- Minimizing printed materials
- Using recycled paper/non-toxic inks
- Monitoring energy usage

The world is banking on innovation. And so are we.

Every day Silicon Valley Bank clients push the boundaries of what was previously thought possible. Their innovations make our lives more efficient, connect people across geographies, make our environment cleaner and safer, eradicate disease and improve the quality of life. Their vision for the world will create a better future. We're banking on it.

John China, Head of Venture Capital Group
3000 Sand Hill Road, Building 3, Suite 150
Menlo Park, California 94025
Phone 415.512.4235 E-mail jchina@svb.com

SVB› *Find a way* | Silicon Valley Bank
svb.com

GPU TECHNOLOGY CONFERENCE

SPEAKER BIOS

Abbas-Turki, Lokman A.
PHD student in Applied Mathematics and Mathematical Finance, Paris-Est UMLV LAMA

Lokman A. Abbas-Turki is a SUPELEC Engineer in Signal Processing and MS degree in Applied Mathematics and Mathematical Finance. Currently PhD Student in Mathematical Finance at: Laboratory of analysis and applied mathematics

THURSDAY 4:00 PM – Derivative Price Modeling on the GPU

Adams, Cleve
CEO, KaaM

Cleve Adams has more than 27 years of computer/software industry experience, including over a decade in the Internet software marketplace. His business insight and executive management style have produced an impressive history of growing and managing companies from startup to more than $1.3 billion in market value. Before joining KaaM, Adams was CEO and president of High Tower Software. Two years prior he served in the same role for Puresight, where he successfully grew the company from a small organization in the security filtering space, into a technology leader in the mobile/wireless market. From the mid-1990's through the early 2000's he served as Vice President of Sales and Marketing for Websense, the global leader of web filtering solutions, where he led the company's transformation from a small security startup to a $750 million global corporation in the span of just four years. Adams brings a wealth of additional startup and global enterprise technology experience, having held executive and management roles at organizations such as Novell, Texas Instruments, Sequel Technologies, Acusoft, and Icon International. He holds an MBA from Pepperdine University.

FRIDAY 3:00 PM – Kerosene and a Match Startup Presentation

Agarawala, Anand
Founder and CEO, BumpTop

Anand Agarawala is creator of the BumpTop desktop interface and Founder/CEO of Bump Technologies, Inc. Originally developed as Anand's Master Thesis in Computer Science and Human Computer Interaction at the University of Toronto, BumpTop is centered on the idea that interacting with computers should be more visceral and human than the rigid, mechanical style imposed by many machines today. Anand demonstrated the technology at TED, the annual, invitation-only "Technology, Entertainment, Design" conference and based on its success there, Anand raised angel and venture capital funding. After the launch of its first public version in April 2009, BumpTop quickly gained attention throughout the industry and has become the most watched video of software in history with over 4 million views on YouTube. It has also been featured in the NYTimes, BBC World News, PC World, Digg, Venture Beat, the WSJ and hundreds of other publications around the globe. BumpTop has several distribution partnerships with leading OEMs, graphics card manufacturers, and multi-touch panel manufacturers. Anand holds several UI related patents and has been invited to speak at numerous conferences (including TED and CEBIT) as well as at institutions and universities across North America, Asia and Europe on topics ranging from next generation user interfaces to entrepreneurship. In 2009, Business Week magazine named Anand among its Top Young Tech Entrepreneurs.

FRIDAY 1:30 PM – BumpTop Startup Presentation

Albanese, Claudio
Prof, Kings College

Claudio Albanese gained a PhD in Physics from ETH Zurich in 1988 and pursued his academic career firstly in Mathematical Physics and, starting from the year 1994, in Mathematical Finance. His postdoctoral experience includes visits to UCLA, NYU and Princeton. He was regular faculty at the University of Toronto and at Imperial College London, achieving the level of full professor. He is currently working mostly as a consultant for financial institutions and teaches as Visiting Professor at King's College London.

THURSDAY 3:00 PM – OPLib: A GPL Library of Elementary Pricing Functions in CUDA/OpenCL and OpenMP

Allada, Veerendra
Graduate Research Assistant, Ames Laboratory / Iowa State University

Porting GAMESS to GPUs has been undertaken by PhD students Andrey Asadchev, Jacob Felder, and Veerendra Allada under the guidance of Prof Mark S Gordon , Prof Theresa Windus, and Scientist Brett Bode at Iowa State University. Andrey Asadchev is a PhD candidate in the Department of Chemistry at Iowa State University and leads the team. Jacob felder is a second year graduate student in the Department of Chemistry at Iowa State University. The speaker Veerendra Allada is a PhD candidate in Computer engineering at Iowa State University.

THURSDAY 5:30 PM – Uncontracted Rys Quadrature Implementation of Up to g Functions on GPUs

Alt, Thomas
CEO, Metaio, Inc.

Dr. Thomas Alt founded in February 2003 the company metaio GmbH together with Peter Meier and is serving the company as CEO in Munich, Germany. His first professional experience with Augmented Reality solutions were gathered at Volkswagen AG in Wolfsburg. During the two years he spent there, he was in diverse functions responsible for the virtual techniques used for manufacturing planning. Dr. Alt's dissertation that had the focus on Augmented Reality, which he graduated Magna Cum Laude in 2002, gave him the idea for a business plan. He proved his right at the Munich as well as the Futuresax business plan contest, which leaded to the foundation of metaio GmbH.

THURSDY 2:00 PM – Metaio, Inc. Startup Presentation

Anderson, David
Research Scientist, UC Berkeley

Dr. David P. Anderson received a Ph.D. in Computer Science from the University of Wisconsin – Madison in 1985. He taught in the Computer Science department at U.C. Berkeley, worked at several startup companies, then returned to U.C. Berkeley as a Research Scientist. His work focuses on Citizen Cyber-Science: using the Internet to involve the global public in scientific research. He leads the BOINC project, which develops middleware for volunteer computing, and he is also involved in creating new technology for distributed thinking and web-based education.

- **THURSDAY 3:00 PM – Volunteer Computing for GPUs: Petaflops for Free**

Ansorge, Richard
Senior Lecturer, University of Cambridge

Dr Richard Ansorge is a Senior Lecturer at the University of Cambridge Department of Physics and Fellow of Fitzwilliam College. He has more than forty years experience of applying computers to scientific data acquisition and processing. He worked in experimental High Energy Physics at CERN and elsewhere until about 1990. More recently he has been developing parallel programming applications for Medical Imaging, particularly MRI and PET.

- **THURSDAY 4:00 PM – Medical Image Registration with CUDA**

Appa, Jamil
Aeronautical Engineer, BAE Systems

Jamil Appa is an Aeronautical Engineer and has worked at BAE SYSTEMS Advanced Technology Centre for the last 12 years having joined from Cranfield University. Initially working on the Flite3D Computational Fluid Dynamics suite based on the then State of the Art Unstructured methods for Euler simulations. This involved the industrialisation, delivery and support of all aspects of the suite from geometry preparation to mesh generation to solver and post processing/visualisation. This was the main CFD suite that was used by the Airbus wing design team to perform the integrated design for the A380. This lead to heading up the technical work on Solar a UK Unstructured Navier Stokes CFD suite developed by BAE SYSTEMS, ARA, QinetiQ and Airbus, which uses innovative mesh generation and solver techniques with the primary aim of developing a robust capability that can handle complex Aerospace geometries. He is currently responsible for IT and HPC related research in the $36m CFMS research programme and leads the Integration Technologies Group in the Mathematical Modelling Department at the ATC. He is project manager and strategist for the BAE SYSTEMS corporate e-Engineering initiative aimed at developing and demonstrating Grid Solutions for a range of business activities both within the Company and with external collaborators, partners and suppliers.

- **FRIDAY 2:00 PM – Computational Based Modelling and Simulation**

Astrom, Mattias
CEO, C3 Technologies AB

Mr. Astrom is a serial entrepreneur who has successfully founded, developed and sold four technology companies. Prior to joining C3 Technologies, Mr. Astrom was CEO of Nordnav, a venture backed GPS spin-out from Stanford, which he took to a successful exit in 2007. Earlier in his career, Mr. Astrom worked in Technology Venture Capital in London, Sophia Antipolis and the Amsterdam. Mr. Astrom has a Master's degree from Stockholm School of Economics and an MBA from Massachusetts Institute of Technology.

- **THURSDAY 1:00 PM – C3 Technologies Startup Presentation**

Atkinson, Simon
Chief Executive Officer, Mirics Semiconductor Inc

After receiving a degree in Applied Physics from the University of Lancaster in 1984, Simon Atkinson joined Plessey Semiconductors as a design engineer specialising in the development of RF IC products for wireless applications. In 1988, he joined LSI Logic where as manager of the Radio Systems Group, he was responsible for the development of various RF/Mixed signal products for mobile telecommunications applications. In 1991, he co-founded Mosaic Microsystems Ltd, a company focussed on providing product development services for portable wireless and telecommunications applications. Between 1991 and 1996, Mosaic successfully developed many RF products covering such areas as DECT handsets, Digital DBS, Cable tuners, Wireless LAN and Keyless entry systems. Mosaic Microsystems along with its US subsidiary, was acquired by Analog Devices Inc. in 1996. At Analog Devices, as Director – Strategic RF Products, he was responsible for creating highly innovative RF products for DECT, GSM and 3G terminals. At the end of 2003, Mr Atkinson left ADI to form Mirics Semiconductor.

- **THURSDAY 5:00 PM – Mirics Semiconductor Inc. Startup Presentation**

Aubert, Dominique
Dr, Université de Strasbourg/Observatoire Astronomique

Dr. Dominique Aubert is a Lecturer at the University & Astronomical Observatory of Strasbourg (France) since 2006. Dominique's research Interests are the Formation of Galaxies, Galactic Dynamics & Cosmology. Methodology: Theoretical Dynamics. Numerical simulations, HPC and GPGPUs.

- **FRIDAY 10:00 AM – Numerical Cosmology Powered by GPUs**

Baliga,Avinash
Senior Engineer, Windows Nexus Debugger, NVIDIA

- **THURSDAY 5:00 PM – The Art of Debugging for the CUDA Architecture**

Beard, Brian
SSD Marketing Manager, Samsung Semiconductor, Inc.

- **FRIDAY 10:30 AM – The High Life for High Performance**

Bell, Nathan
Research Scientist, NVIDIA

Nathan Bell joined NVIDIA Research in August 2008. His current research interests include sparse linear algebra and programming models for parallel computing. Nathan contributes to several open source projects including Thrust, a high-level parallel template library, and PyAMG, a library of algebraic multigrid methods in Python. Nathan received a bachelor's degree in Computer Science from Georgia Tech and a PhD in Computer Science from the University of Illinois at Urbana-Champaign (UIUC).

- **FRIDAY 10:00 AM – Advanced Numeric Computing**
- **THURSDAY 2:00 PM – Thrust: A Parallel Template Library for CUDA**

Bernard, Eric
CEO, Visioglobe

After Airplane studies, a Business School and an MBA, Eric started to work as an Oracle Consultant for major companies in Europe, such as EA Sports, Atari, Michelin, Gemplus. He then joined SQLI, a Software service company in France, as an Account manager and finally worked for a leading web agency in Europe: Fi-System, as a South East France Director. He met Philippe in 2007 and decided to move on Visioglobe challenge.

- **FRIDAY 2:00 PM – Visioglobe Startup Presentation**

Bienert, Joerg
CEO, Empulse GmbH

Jörg Bienert Co-founder and managing director of empulse AG Cologne. He has many years of experience in IT strategy development, project management and entrepreneurship. After his studies in electrical engineering, he started his career as international IT project manager in a German insurance company before he became a manager at Accenture. In 2000 he joined the international internet agency Framfab as COO and focused on business development and sales. In 2006 he founded empulse together with Michael Hummel, a former colleague from accenture

- **THURSDAY 2:30 PM – Empulse GmbH Startup Presentation**

Bischof, Horst
Professor, Graz University of Technology

Horst Bischof received his M.S. and Ph.D. degree in computer science from the Vienna University of Technology in 1990 and 1993, respectively. In 1998 he got his Habilitation (venia docendi) for applied computer science. Currently he is Professor at the Institute for Computer Graphics and Vision at the Technical University Graz, Austria. H. Bischof is member of the scientific boards of the applied research centers ECV, VrVis and KNOW. H. Bischof is board member of the Fraunhofer Inst. für Graphische Datenverarbeitung (IGD). His research interests include object recognition, visual learning, motion and tracking, visual surveillance and biometrics, medical computer vision, and adaptive methods for computer vision where he has published more than 400 peer reviewed scientific papers. Horst Bischof was co-chairman of international conferences (ICANN, DAGM), and local organizer for ICPR'96. He was program co-chair of ECCV2006 and Area chair of CVPR 2007, ECCV2008, CVPR 2009, ACCV 2009. Currently he is Associate Editor for IEEE Trans. on Pattern Analysis and Machine Intelligence, Pattern Recognition, Computer and Informatics and the Journal of Universal Computer Science. Horst Bischof has received several awards among them the 29th Pattern Recognition award in 2002; the main price of the German Association for Pattern Recognition DAGM in 2007, the Best scientific paper award at the BMCV 2007 and the Best scientific paper award at the ICPR 2008.

- **WEDNESDAY 3:00 PM – Important Trends in Visual Computing (General Session)**
- **FRIDAY 1:00 PM – Computer Vision with Horst Bischof**

Bishop, Lars
Engineer, NVIDIA

Lars M. Bishop has been an engineer in the Mobile Developer Technologies group at NVIDIA for four years. He specializes in 3D user interfaces and works with external developers to leverage the power of NVIDIA's Tegra mobile platform in their next-generation media applications. He received a BS in Math/Computer Science from Brown University and an MS in Computer Science from the University of North Carolina at Chapel Hill. Prior to joining NVIDIA, Lars worked for nine years at NDL (now Emergent Game Technologies), leading the development of the Gamebryo3D cross-platform game engine. He is the co-author of the book Essential Mathematics for Games and Interactive Applications: A Programmer's Guide. His outside interests include photography and bass guitar.

- **FRIDAY 3:00 PM – Zombies on Tegra: A Case Study in Mobile Augmented Reality**

Blackman, Sam
CEO, Elemental Technologies

Sam Blackman, CEO and Chairman of Elemental Technologies, Inc., brings extensive management experience and video processing expertise to the Elemental team as chief executive officer. Prior to co-founding Elemental in 2006, Samuel specified and architected next-generation products as an IC Design Manager for Portland-based Pixelworks. He spent time in China organizing the company's Shanghai design center and was responsible for a wide variety of functional blocks on six ImageProcessor ICs. Prior to joining Pixelworks in 2000, Sam held engineering positions at Silicon Graphics and Intel Corporation. Sam holds an M.B.A from University of Oregon, an M.S.E.E. from University of California at Berkeley and a B.S.E.E. from Brown University in Providence, RI.

- **THURSDAY 1:00 PM – Elemental Technologies, Inc. Startup Presentation**

Bleiweiss, Avi
NVIDIA

Avi Bleiweiss joined NVIDIA Corporation in 2007 as a member of the architecture group with his main role of leveraging GPU computing to accelerate game AI workloads. He spans 24 years of R&D experience in the development of high end graphics systems. Previously, he worked for AMD/ATI where he led the software development of ASHLI, a GPU shading toolkit, and frameworks evaluating GPU performance for game physics and ray tracing. Before that he was a principal engineer at Silicon Graphics, responsible for simulator and driver implementation of a programmable geometry engine. Formerly, he was a distinguished engineer at Kubota Graphics Computer, overseeing the architecture of Denali's rasterization subsystem. And earlier, Mr. Bleiweiss was a visitor scientist at Hewlett Packard (HP) Laboratories, where he collaborated with Prof. Don Greenberg on a design of HP Precision Architecture's coprocessor for realizing radiosity algorithms. He contributed numerous publications to Siggraph and Graphics Hardware conferences. Avi Bleiweiss joined NVIDIA Corporation in 2007 as a member of the architecture group with his main role of leveraging GPU computing to accelerate game AI workloads. He spans 24 years of R&D experience in the development of high end graphics systems. Previously, he worked for AMD/ATI where he led the software development of ASHLI, a GPU shading toolkit, and frameworks evaluating GPU performance for game physics and ray tracing. Before that he was a principal engineer at Silicon Graphics, responsible for simulator and driver implementation of a programmable geometry engine. Formerly, he was a

distinguished engineer at Kubota Graphics Computer, overseeing the architecture of Denali's rasterization subsystem. And earlier, Mr. Bleiweiss was a visitor scientist at Hewlett Packard (HP) Laboratories, where he collaborated with Prof. Don Greenberg on a design of HP Precision Architecture's coprocessor for realizing radiosity algorithms.

➔ **FRIDAY 2:00 PM – Scalable Multi Agent Simuluation on the GPU**

Blewitt, Christopher
CEO, Luminova

Chris Blewitt throughout his career has developed a considerable reputation in lighting application and associated computer software development. In 1989, he led a technical development team producing one of the world's first computer based 3D Lighting Analysis programs, adopted by GE Lighting Systems (USA) as their primary in-house lighting design system. Also Chris has an extensive construction industry background, starting in 1974 as State Manager and Senior Lighting Engineer for Holophane Australia. In 1980, he became one of the founding partners of Versalux Pty. Ltd, distributor for GE Lighting Systems in the Australiasian region. In late1995, Chris Blewitt and Paul Arden formed a small development group that lead to the establishment of Luminova, the world's first company dedicated to high speed iterative modeling systems in conjunction with accurate and measurable physical based computer rendering technologies. Through his considerable experience of working directly with many different design, engineering and project management teams for large scale developments, combined with his deep knowledge of 3D systems and physical based rendering technologies, Chris has aquired a unique perspective and inate understanding of how these two diverse fields can be brought together most effectively. This has been clearly proven over the last 13 years, where under Chris's leadership, Luminova has successfully delivered over 3,000 projects for many of the world's largest commpanies.

➔ **THURSDAY 1:00 PM – Luminova Startup Presentation**

Bodin, François
CTO, CAPS entreprise

Francois Bodin cofounded CAPS (www.caps-entreprise.com) in 2002 while he was a Professor at University of Rennes I and since January 2008 he joined the company as CTO. His contribution includes new approaches for exploiting high performance processors in scientific computing and in embedded applications. Prior to joining CAPS, Francois Bodin held various research positions at University of Rennes I and at the INRIA research lab. He has published over 60 papers in international journals and conferences and he has supervised over 15 PhD thesis. Professor Francois Bodin holds a Master's in CS and a PhD in CS, both from University of Rennes I. He is a member of numerous program committees (PPoPP2010, HIPS2009, MUCOCOS09, MUCOCOS08, Europar 2007, EUROPAR99, ODE 2007, ICS'96, ICS'97, ICS'98, AGI'94, AGI'95, AGI'96 , HPCS'95, ...). He has been invited as a speaker to Hipeac 2009, IWOMP 2009 and Parco 2009.

➔ **THURSDAY 5:30 PM – Heterogeneous Multicore Parallel Programming**

Boone, Christopher
President & CEO, Visuvi Inc

Mr. Boone is Visuvi's President & CEO. He most recently served as President and CEO of AppCentral, a desktop virtualization software company and was formerly CEO of Cimbal Management Consulting where he worked as a telecom, media and technology expert advising portfolio managers, law firms, technology companies and private equity firms. He was previously a director of business and corporate development with Excite@Home, working with At Home's senior leadership team in negotiating its distribution agreements with cable and telecom customers and was a key member in the firm's bankruptcy process. He previously worked as an investment banker with JP Morgan's Structured Products Group in New York managing a billion dollar, global portfolio. He is a member and recognized scholar of the Gerson Lehrman Group's Technology, Media & Telecommunications council and holds four patents pending in virtualization and distributed networks. Mr. Boone holds a bachelor of science degree from the University of Notre Dame.

➔ **THURSDAY 5:00 PM – Visuvi, Inc. Startup Presentation**

Bozinovic, Nikola
VP Engineering, MotionDSP

Nikola Bozinovic is a VP of Engineering at MotionDSP. Under his technical lead, MotionDSP partnered with In-Q-Tel and NVIDIA to bring the next-generation video enhancement technology to both intelligence community and general consumers. Prior to building MotionDSP's R&D in 2006, he held engineering positions at Microsoft Research Asia and Veodia. He was awarded EURASIP best paper award (co-authored with Janusz Konrad) for his work on frequency-domain motion analysis. His other research interests include video reconstruction and restoration, motion modeling, and video analytics. He holds a Ph.D. in Electrical and Computer Engineering from Boston University, and a B.S.E.E. from University of Nis, Serbia.

➔ **FRIDAY 11:00 AM – Multi-Frame Analysis: The Future of Video Processing**

Bradley, Thomas
NVIDIA

Thomas Bradley MEng(Hons) MIEE graduated with a first-class MEng degree in Computer Systems Engineering from the University of Bristol, UK, in 2000, having also completed the final year of the Diplome d'Ingenieur at l'Ecole Nationale Supurieure de Telecommunications in Brest, France. He worked as processor architect for video encoding processors at STMicroelectronics before moving to ClearSpeed Technology plc to lead architecture development for general purpose parallel processors. Since then he has specialized in High Performance Computing software development at ClearSpeed and now at NVIDIA.

➔ **THURSDAY 11:00 AM – Overview of Computational Finance on the GPU**

➔ **FRIDAY 3:30 PM – Advanced C for CUDA**

Bridges, Corey
Executive Producer, Multiverse

Corey Bridges: Responsible for business development and strategic planning, Corey Bridges is helping build the world's leading network of interactive entertainment, which includes casual games, Massively Multiplayer Online Games (MMOGs), and 3D virtual worlds. Over the past year, Corey has directed development of two ground-breaking projects for the casual game market, Battle and

Team Battle, the world's first web-based, multiplayer, real-time combat series. Both titles artfully blend proven gameplay common to console systems and installed games with the immediacy and ubiquity of Flash. Since the company's beginning, Corey has also helped pioneer the concept of virtual worlds as a new medium that will support a range of applications, including entertainment, education, business collaboration, research, and simulation.

FRIDAY 2:00 PM – Multiverse Startup Presentation

Brouwer, William
Senior Physicist, Stone Ridge Technology

Dr. William Brouwer is a senior physicist at Stone Ridge Technology. He completed a BSc at the University of Queensland in 1998 a BSc (Hons) in 1999 and a PhD in physics from William and Mary in 2005. Dr. Brouwer was most recently a postdoctoral scholar at The Pennsylvania State University where he contributed to the NSF funded ChemXSeer cyberinfrastructure project. He has worked in computational quantum chemistry and is currently working on projects in computational geophysics using accelerating graphics hardware. He is a member of the Golden Key and National Society of Collegiate Scholars honor societies and is published in several fields including chemical physics, document analysis/algorithms and computational physics.

FRIDAY 10:30 AM – SCVF: A Framework for Visualization of Grid and Particle Based CUDA Simulations

FRIDAY 4:00 PM – Energy Industry Applications: Three Examples of GPU Acceleration in the Oil and Gas Industry Using CUDA

Butrashvily, Mordechai
CEO, CTO, GASS / Hoopoe Cloud

> Rich experience with HPC and GPU technologies for computing and visualizations in distributed environments
> Former member of IMOD
> Head of the GPGPU workgroup @ IGT (ICS Magnet program)
> Involved in HPC related activities in Israel governed by the Israeli Chief Scientist in the fields of AOI, image/video processing, data classification and more
> Involved with several education activities aimed at teaching K-12 and youngsters scientific computing and supercomputing with remote/web based technologies at Weizmann Science Institute

THURSDAY 5:00 PM – GASS Company for Advanced Supercomputing Solutions LTD Startup Presentation

Catuhe, David
CTO, Vertice

David Catuhe is a Microsoft Certified Technical Specialist, Microsoft Certified Solution Developer .NET, Microsoft Certified Trainer, a Most Valuable Professional .NET and the CTO of Vertice as well as a Consultant for Bewise. David Catuhe is, before all, the father of Nova. Fan of video games since his youth, he has therefore always been interested by real time 3D and then wanted to offer to the industry a tool to create easily, quickly and at an affordable cost real time applications answering their professional needs. David Catuhe began the development of Nova during his studies at the University Paul Sabatier in Toulouse. At the end of his studies, David has joined Bewise (sister company of Vertice) as one of the shareholders and managers, and continued to develop Nova at the rhythm of the latest tools made available to the development teams by the editors. In 2002, as demand was increasing on real time 3D tools, David and his associates have decided to create the Vertice company, editor and distributor of Nova. Since then, David manages and leads the development of Nova, always keeping in mind the aim of the product at the start of its creation which brought the success it has today.

THURSDAY 3:00 PM – Vertice Startup Presentation

Cazeaux, Jean Bernard
CEO, Visualization Sciences Group, Inc.

Founder in 1987 of the G5G Company who produced 3D graphics applications, Jean-Bernard Cazeaux lead the TGS group in 1998 as CEO. In May 2004, TGS joined Mercury Computer Systems group and Jean-Bernard handles the development and the strategy of the 3DViz activities BU. He has initiated and organized the many TGS acquisitions: Portable Graphics, Yonowat, Eovia and Indeed-Visual. Through a MBO operation, supported by IRDI-ICSO Private Equity, VSG is a stand-alone company since June 2009 led by Jean Bernard Cazeaux as CEO. Jean Bernard Cazeaux holds a Master's Degree in Computer Sciences from the University of Sciences of Bordeaux.

THURSDAY 2:30 PM – Visualization Sciences Group, Inc. Startup Presentation

Cedeno, Rafhael
CTO, Multiverse

Rafhael Cedeno: With a distinguished track record of creating fundamental technologies that made many large-scale applications possible, Rafhael Cedeno directs overall technology development at Multiverse. At companies like Netscape, Excite, and Kontiki he built technologies that defied convention and enabled entirely new business models to flourish. His proven ability to create scalable solutions has formed the foundation of several innovative new companies. Under his guidance, Multiverse has built a revolutionary virtual world development platform that enables new creative talent to enter an industry once limited to only extremely well-funded teams.

FRIDAY 2:00 PM – Multiverse Startup Presentation

Cheryauka, Arvi
Sr Sys Eng, GE Healthcare – Surgery

Dr. Cheryauka is a senior 3D IQ Systems Engineer with GE Healthcare– Surgery. He has MSc. in Computational Physics from Novosibirsk University, Russia, and PhD in Physics from the Russian Academy of Sciences. Dr. Cheryauka has worked in non-destructive evaluation and medical imaging, in both academia and industry. His interests include numerical design, forward and inverse simulations, system prototyping, multi-modality image registration / fusion, and image-guided, minimally-invasive interventions.

THURSDAY 3:00 PM – GPU-Accelerated Imaging on Mobile X-ray C-arm

Chong, Jike
Ph.D. Candidate, University of California, Berkeley

Jike Chong is currently a Ph.D. candidate at University of California, Berkeley, working on application frameworks in computational finance and automatic speech recognition to help domain experts efficiently utilize highly parallel computation platforms. Prior to Berkeley, he was a member of technical staff at Sun Microsystems, Inc,

where he designed core logic in the Niagara II processor. Jike received his B.S. and M.S. degrees from Carnegie Mellon University, Pittsburgh, Pennsylvania in 2001. He also conducted research at Intel Application Research Labs and Xilinx Research Labs. He is an Intel Ph.D. research fellow, a Berkeley Mayfield Fellow for Entrepreneurship, and a student member of IEEE, HKN, and TBP.

➔ **THURSDAY 2:30 PM – Using GPUs to Estimate the Value-at-Risk of Portfolios**

Chung, Jimmy
Technical Marketing Manager, Graphics Memory, Samsung Semiconductor, Inc.

Jimmy Chung has more than eight years of experience at Samsung Semiconductor in its Memory Division. During that time, he has handled a variety of responsibilities in technical marketing, OEM technical support and hardware development for computing technology. Prior to joining Samsung, he managed memory suppliers and served as a memory development engineer at Hewlett-Packard Company. Jimmy holds a Master of Science degree in Electrical Engineering from California State University, Northridge.

➔ **FRIDAY 11:30 AM – Next-Generation Graphics DRAM: Challenges and Opportunities**

Clark, Calvin
Microsoft

Calvin Clark joined Microsoft in 1996, starting in Redmond as a developer and later as serving as a developer lead on e-Commerce and B2B products. In 2001 he moved to the San Francisco Bay Area to work on the consulting side of the business, helping Microsoft's software and hardware partners build successful products on Microsoft's server platforms. Since 2006 his focus has been on High Performance Computing solutions, delivering trainings and consulting services to numerous ISVs, Solution Integrators, and OEMs in the HPC space. He lives in Menlo Park, CA with his wife and daughter.

➔ **FRIDAY 4:00 PM – Clusters with GPUs under Linux and Windows HPC**

Clegg, Don
Vice President, Marketing and Business Development, Super Micro Computer, Inc.

Heading up Supermicro's Marketing & Business Development groups, Don brings more than 20 years of direct experience in Sales, Marketing and Computer Design. Focused on long term growth opportunities, Don and his team are responsible for working with key partners to develop new business initiatives. Prior to Supermicro, Don was vice president of marketing for Netcell Corporation, a VC backed storage semiconductor start-up as well as vice president of marketing and strategic sales at Tyan Computer Corporation. Don has also held design and marketing positions at Opti, Inc., Chips and Technologies, Wyse Technologies and Northstar Computers. Don graduated with High Honors from Brigham Young University with a Bachelor's Degree in Electrical Engineering.

➔ **FRIDAY 11:00 AM – Integrated Design that Shatters 1U Server Performance Records**

Chiu, Ting-Wai
Professor, National Taiwan University

Being a high energy theoretical physicist, Prof. Ting-Wai Chiu's research focuses on the quantum field theory of strong interaction, Quantum Chromodynamics (QCD). The most promising approach to solve QCD nonperturbatively is to discretize the continuum space-time into a 4 dimensional lattice, i.e., lattice QCD, and to compute physical observables with Monte Carlo simulation. For lattice QCD with exact chiral symmetry, it often requires supercomputers (e.g., 10 racks of IBM BlueGene) to perform the simulations. Recently, the TWQCD Collaboration in Taiwan, led by Prof. Chiu, is the first lattice QCD group around the world to use a GPU cluster (120 GPUs) to perform large-scale unquenched simulations of lattice QCD with exact chiral symmetry, attaining 14 Teraflops (sustained) at a price of $200,000. Professor Chiu received his Ph.D degree in physics from University of Utah in 1980, and did his postdoctoral research at University of California, Irvine, 1980-1982. Moreover, Prof. Chiu was a Visiting Professor of Physics at California Institute of Technology, 1987-1989. Currently, Prof. Chiu is a Professor of Physics, and an Associate Director of the Center for Quantum Science and Engineering (CQSE) at National Taiwan University.

Collins, Stephen
CEO, Anark Corp

Stephen Collins is a co-founder of Anark and an 19 year industry veteran within the high-tech, entertainment, and enterprise 3D market sectors. He has extensive sales, marketing, business development and private equity financing experience. Stephen has guided Anark through its evolution from innovative 3D game development studio to an emerging market leader with advanced and practical 3D and data management solutions that enable customers to cost effectively realize the full potential of their CAD design data and other 3D graphics assets via a wide variety of collaboration, communication and visualization applications. Prior to founding Anark, Stephen worked for market leaders such as Xerox, Tandy's Business Systems Division, and Microage, where he consistently exceeded sales goals and company expectations. Stephen earned his BA from the University of North Carolina at Charlotte in 1989, where he majored in History and minored in Marketing.

➔ **THURSDAY 3:30 PM – Anark Corporation Startup Presentation**

Cordingley, Dan
President & CEO, Teradici Corporation

Prior to co-founding Teradici, Dan was a General Manager with Intel Corporation, where he oversaw the Digital Subscriber Line (DSL) silicon business unit, a group with leading market share in key segments of the highly competitive DSL chipset industry. Before Intel, Dan was a General Manager and Director at Level One Communications, a leading supplier of networking and telecom semiconductors that was acquired by Intel in 1999 for US $2.2B. Under Dan's leadership, Level One conceived, standardized, and delivered several important communications standards in widespread use today. Dan has also held senior marketing and design engineering roles at Mitel Semiconductor (now Zarlink), Spectrum Signal Processing, Nortel Networks, and IBM

➔ **FRIDAY 1:00 PM – Teradici Corporation Startup Presentation**

Coumans, Erwin
Physics Simulation Team Lead, Sony Computer Entertainment US R&D

Erwin Coumans is simulation team lead at Sony Computer Entertainment US R&D. After his study of

computer science at Eindhoven University, he worked for Guerrilla Games in the Netherlands and Havok in Ireland doing collision detection and physics engine development. Erwin is actively promoting the sharing of ideas and technology through forum discussion, open standards and open source. He is main author of the Bullet Physics Library, used by several game companies for games on PLAYSTATION 3, XBox 360, PC, Wii and iPhone.

- **FRIDAY 1:00 PM – Bullet: A Case Study in Optimizing Physics Middleware for the GPU**

Cox, Samuel
CEO, Milabra

Samuel R. Cox co-founded Milabra with Dr. Naveen Agnihotri, Ph.D. in October 2008 in order to commercialize neural network machine vision technology developed by Dr. Agnihotri. Sam is responsible for the company's vision to use cutting edge visual recognition techniques combined with visual research in visual semiotics in order to target advertising based on visual content, instead of text. The desire to use technology to not only find, but find meaning as well is a result of Sam's entrepreneurial and academic background. With a deep passion for linguistics, art and entrepreneurship, Sam originally studied Art History, Chinese Language and Business at Queen's University in Kingston, Ontario, where he started his first software and design firm. He later completed his MBA at Cass Business School in the United Kingdom, and has worked as both an entrepreneur and investor in China, Canada, England and the United States before joining Milabra. He is originally from New York City.

- **THURSDAY 2:30 PM – Milabra Startup Presentation**

Cui, Xiaohui
Oak Ridge National Laboratory

Dr. Xiaohui Cui is an associate staff in the Computational Sciences & Engineering Division of Oak Ridge National Laboratory. His research interests include Swarm Intelligence, Agent Based Social Modeling and Simulation, Emergent Behavior in Complex System, GPU computing, Text Mining, Information Retrieval and Knowledge Discovery. He is a member of North American Assoc. for Computational Social and Organizational Sciences, IEEE and ACM. His current research focuses in developing new computational algorithms inspired from biological models to meet multiple national crucial challenges. His current research works include high performance computing, text mining, collective intelligence of multi-agent system, parallel and distributed knowledge discovering, large scale agent based simulation, massively parallel implementations on GPU and PS3, terrorist threat-vulnerability analysis, and malicious cyber behavior detection.

- **THURSDAY 4:00 PM – Large-Scale Text Mining on the GPU**

Dally, Bill
Chief Scientist, NVIDIA

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.

- **WEDNESDAY 4:30 PM – Breakthroughs in High Performance Computing (General Session)**
- **THURSDAY 11:00 AM – Research Summit Fast Forward with Bill Dally**

Deguy, Sebastien
CEO, Allegorithmic

Dr. Deguy has a computer science background with a specialization in mathematics, random processes, simulation, computer vision and image synthesis. He is also an award-winning director and producer of traditional and animated short films.

- **FRIDAY 2:30 PM – Allegorithmic Startup Presentation**

Dixon, Matthew
Postdoctoral Fellow, UC Davis

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.

- **THURSDAY 2:30 PM – Using GPUs to Estimate the Value-at-Risk of Portfolios**

Domine, Sebastien
Sr. Dir. Dev Tech Tools, NVIDIA

Sebastien Domine is Sr. Director of Developer Technology Tools at NVIDIA. He runs various software engineering teams and oversees the development of software products dedicated to ease the developer's life and to foster the creation of more applications that can take advantage of the GPU. Prior to NVIDIA, he worked on PC games at GameFX/THQ and 3D digital content creation tools at Katrix and Nichimen Graphics. He holds a Diplome d'Ingenieur in Computer Science from EPITA, Paris, France.

- **THURSDAY 11:00 AM – NEXUS: A Powerful IDE for GPU Computing on Windows**
- **FRIDAY 1:00 PM – NEXUS: A Powerful IDE for GPU Computing on Windows**

Edgar, Richard
Initiative in Innovative Computing, Harvard University
Richard Edgar is a research associate at the Initiative in Innovative Computing at Harvard University. He is currently working on adding GPU acceleration to a number of projects, primarily the Murchison Widefield Array and medical image registration with Freesurfer. Richard obtained his undergraduate and graduate degrees at the University of Cambridge, concentrating in theoretical astrophysics.

- **FRIDAY 3:00 PM – Diesel-Powered GPU Computing: Enabling a Real-Time Radio Telescope in the Australian Outback**

El Dokor, Tarek
CEO and CTO, Edge3 Technologies

Tarek El Dokor is co-founder and CTO of Edge 3 Technologies, an off-shoot of the Embry-Riddle Machine Vision Lab. El Dokor is the Machine Vision Lab director and is one of the leading authorities on gesture recognition. He has co-authored and submitted various patents and publications in his field. El Dokor is also a faculty with the Electrical and Computer Engineering department at Embry-Riddle. He is the recipient of the 2007 Honda Initiation Grant for emerging technologies. He is also a 2008 University Partner with Nvidia. He is also the recipient of the 2008/2009 Honda R&D Americas and Honda Research Institute's Researcher of the Year Award.

- **THURSDAY 5:00 PM – Examples of CUDA-based Realizations in Artificial Vision and Soft Computing**
- **THURSDAY 1:30 PM – Edge 3 Technologies Startup Presentation**

Ellens, Marc
Software Engineer, X-Rite, Inc.

Marc Ellens is a software engineer currently working for the premier color company, X-Rite, Inc. He began his career at Lectra Systems architecting a suite of design and simulation applications for the textile industry. While at X-Rite, he was lead programmer for PersonalDesigner (used by retail consumers to change the paint color of their rooms) and has recently been working on GPU-based rendering from multi-angle spectrophotometer data. Now working on spectral reconstruction from image data, he has come to rely on CUDA. Marc holds a PhD in Computer Graphics and Modeling from the University of Utah.

- **FRIDAY 1:00 PM – High Dynamic Range Spectral Imagery**

Fiorini, Pierre
CEO, HPC PROJECT

Pierre founded HPC Project in 2007 with an experimented team of computer professionals. The company has developed the concept of desktop supercomputing appliances, a new generation of computing solutions leading the coming revolution of interactive simulation. Before this venture, Pierre has worked for the French Minister of Defence leading research in parallel computer design and programming. He was division head at the French Minister of Industry for the funding of R&D in telecoms, consumer electronics and multimedia industries. Pierre spent seven years as a partner of I-Source Gestion, the main VC firm in France for early-stage investments. Pierre Fiorini has graduated from Ecole Polytechnique in 1989. Pierre gets a PhD degree with Honor from Université Paris-Sud in 1995. Pierre is ACM and IEEE Member.

- **FRIDAY 3:30 PM – HPC PROJECT Startup Presentation**

Fatica, Massimiliano
Senior Applied Engineer, NVIDIA

Massimiliano Fatica is a Senior Applied Engineer at NVIDIA where he works in the area of GPU computing (high-performance computing and clusters). Prior to joining NVIDIA, he was a research staff member at Stanford University where he worked on applications for the Stanford Streaming Supercomputer. He holds a laurea in Aeronautical Engineering and a PhD in Theoretical and Applied Mechanics from the University of Rome "La Sapienza".

- **FRIDAY 4:00 PM – Clusters with GPUs under Linux and Windows HPC**

Frick, Harald
CEO, QuiLogic

Harald Frick is founder and CEO of QuiLogic Inc. a company specialized in providing in-memory database technologies. He is the principal architect of the SQL/XML-IMDB a memory database product line. After his study of mathematics and biochemistry at University Innsbruck (Master of Science) he owned an assistant professor position on the Institute for Computer Science at University Innsbruck. During his time at that position Harald was responsible for developing the hardware and the operating system infrastructures for a special parallel computer system developed and build at the Institute. He also conducted research on visual programming languages for easier application building in the field of microbiology and biochemistry.

- **FRIDAY 4:00 PM – SQL/XML-IMDBg: A GPU In-Memory Database and Query Co-Processor**

Fuchs, Ludwig A.
Board Member, Realtime Technology AG

Ludwig A. Fuchs is co-founder of Realtime Technology AG and a member of the Board of Directors with particular responsibility for finance, research and development, human resources and marketing.

- **THURSDAY 3:30 PM – Realtime Technology AG Startup Presentation**

Fung, James
Developer Technology, NVIDIA

James Fung's work has been in the area of applying GPU Hardware for parallel general purpose computing, including implementing Computer Vision on the GPU. He is founder of the OpeNVIDIA project, which won the ACM Multimedia 2005 Open Source Software Award. He authored "Computer Vision on the GPU" in the popular GPU Gems 2 series of graphics programming books. This work has achieved implementation of vision algorithms on the GPU, including projective image stitching, Chirplet detection, Radon Transforms and natural feature processing and matching. He holds a Ph.D. in Electrical and Computer Engineering from the University of Toronto and has authored over a dozen papers in IEEE and ACM conferences in the areas of parallel GPU Computer Vision and Mediated Reality. Merging technology with artistic endeavors, his work has appeared on the Discovery Channel and Canadian national radio. He currently works

at NVIDIA examining computer vision and image processing on graphics hardware.

- **WEDNESDAY 10:45 AM – DirectCompute (Pre-Conference Tutorial)**
- **THURSDAY 3:00 PM – Advances in GPU-based Image Processing and Computer Vision**

Gil, Laurent
CEO and President, Viewdle, Inc.

Laurent Gil is an internet pioneer and a serial entrepreneur, with activities from technology investment banker in Asia and the US to internet start-up in Latin America to non profit work in Africa. As President and CEO, Laurent leads the company's technology, marketing, business and corporate development strategy. Prior to joining Viewdle, Laurent was a senior investment banker that held various positions in corporate finance and capital markets at the French investment bank Credit Agricole, with assignments in New York, Tokyo and Paris. During his MBA at Wharton, Laurent co-founded a boutique investment bank in Brazil that completed several large transactions in the telecom sector in Latin America. He also co-founded and served as CFO of TAHO, a wireless internet service provider in Rio de Janeiro, which raised US $10MM in 2001 and subsequently launched a service of broadband wireless internet in communities. More recently, Laurent Gil spent few years in West Africa deploying wireless networks in high schools, for a non-profit organization that he started.

- **THURSDAY 1:00 PM – Viewdle Inc. Startup Presentation**
- **THURSDAY 4:00 PM – GPU Ventures, for the Entrepreneur**

Gan, Zhengting
Q-Chem, Inc.

- **THURSDAY 11:00 AM – Quantum Chemistry on the GPU: Accelerating DFT Calculations**

Glavtchev, Vladimir
Software Engineer Intern, NVIDIA, University of California, Davis

Vladimir Glavtchev received his BS in Computer Engineering from the University of California Davis in 2008. In the summer of 2008, he was a Software Engineering intern at the BMW Technology Office in Palo Alto, CA, where he prototyped a traffic-sign recognition system. Currently, he is a Master's student at the University of California Davis currently doing research in the area of computer vision using GPUs. Concurrently, he is working as a Software intern at NVIDIA in the Embedded Multimedia group working on applications targeted for automotive uses.

- **FRIDAY 2:00 PM – Speed Limit Recognition System (Optimized for processing on the GPU)**

Gryc, Andy
Senior Product Marketing Manager, Graphics and Multimedia Technologies QNX Software Systems

Andy Gryc has been a software developer and designer for more than 20 years. Prior to joining QNX Software Systems, he worked as the lead embedded software architect for GM OnStar; designed and implemented a speech recognition engine for a speech technology company; and served as a member of the Hewlett-Packard team that created the software for palmtops and the BIOS for the Omnibook notebooks. He currently works as the senior product marketing manager for graphics and multimedia technologies for QNX Software Systems.

- **FRIDAY 11:00 AM – Using OpenKODE to Create a Seamless Driving Experience**

Gupta, Sumit
Senior Manager, NVIDIA

Sumit Gupta is a Sr. Manager in the Tesla GPU Computing HPC business unit at NVIDIA since 2007. In this role, Sumit is responsible for marketing and business development of the CUDA-based GPU computing products. Prior to this, Sumit served in a range of positions such as Product Manager at Tensilica, Entrepreneur-in-Residence at Tallwood Venture Capital, Post-Doctoral Researcher at University of California San Diego and Irvine, chip designer at S3 Inc, software engineer at IBM and IMEC, Belgium. Sumit has a Ph.D. in Computer Science from the University of California, Irvine and a B.Tech. in Electrical Engineering from the Indian Institute of Technology, Delhi and has authored one book, one patent (awarded), several book chapters, and more than 20 peer-reviewed conference and journal publications.

- **THURSDAY 4:00 PM – GPU-based Libraries: Accelerating the Next Generation of Applications**

Han, Jeff
Founder, and Chief Scientist, Perceptive Pixel, Inc.

Jeff Han is the founder and chief scientist of Perceptive Pixel, a startup dedicated to the research, development, and marketing of multi-touch interfaces for the knowledge worker. Jeff is most recently credited for dramatically introducing to the general public the concept of the multi-touch interface at the TED 2006 conference, and the publication of his techniques have directly lead to a resurgence of practice and research in the field. Jeff's work has traditionally encompassed the fields of computer graphics and multimedia systems, but is now refocused towards machine learning and HCI (human-computer interaction). Since then Jeff has built Perceptive Pixel into a successful, profitable and rapidly growing business working with top-tier partners in such industries as government/defense, intelligence, medical imaging, broadcast, energy, finance, and education. A TED speaker in 2006, and among other honors, named to the Time 100 most influential persons list in 2008, Jeff continues to contribute frequently to research communities and publications. Jeff's formal training was in electrical engineering and computer science at Cornell University, where he worked on the innovative CU-SeeMe videoconferencing system.

- **THURSDAY 4:00 PM – Perceptive Pixel, Inc. Startup Presentation**

Hanchevici, Teodor C.
Technical Project Leader, PCI Geomatics

Teodor C. Hanchevici is a Technical Lead at PCI Geomatics and has been with the company since 2005 and he is currently leading the development of PCI GeoImaging Accelerator (GXL). Teodor has over 10 years experience in software development. His main focus is image processing and high performance computing. Teodor holds a Bachelor of Science in Applied Mathematics from Babes-Bolyai University (Romania), and a Professional Doctorate in Engineering from Technical University Eindhoven (the Netherlands).

- **THURSDAY 5:00 PM – Mapping Satellite Imagery on the GPU: Fast Orthorectification and Pan-Sharpening**

Hanrahan, Pat
Professor, Stanford University

Pat Hanrahan is the CANON Professor of Computer Science and Electrical Engineering at Stanford University where he teaches computer graphics. His current research involves visualization, image synthesis, virtual worlds, and graphics systems and architectures. Before joining Stanford he was a faculty member at Princeton. Pat has also worked at Pixar where he developed developed volume rendering software and was the chief architect of the RenderMan(TM) Interface – a protocol that allows modeling programs to describe scenes to high quality rendering programs. In addition to PIXAR, he has founded two companies, Tableau and PeakStream, and served on the technical advisory boards of NVIDIA, Exluna, Neoptica, VSee and Procedural. Professor Hanrahan has received three university teaching awards. He has received two Academy Awards for Science and Technology, the Spirit of America Creativity Award, the SIGGRAPH Computer Graphics Achievement Award, the SIGGRAPH Stephen A. Coons Award, and the IEEE Visualization Career Award. He was recently elected to the National Academy of Engineering and the American Academy of Arts and Sciences.

- **WEDNESDAY 3:00 PM – Important Trends in Visual Computing (General Session)**
- **FRIDAY 3:00 PM – Domain Specific Languages for Programming GPUs**

Hanweck, Gerald
Principal, Hanweck Associates, LLC

Gerald Hanweck, Jr., is founder and principal partner of Hanweck Associates. Previously, Mr. Hanweck served as JPMorgan's Chief Equity Derivatives Strategist from 2000 to 2003, and led the bank's U.S. Fixed-Income Derivatives Strategy team. He has taught master's-level business courses at Northwestern University's Kellogg Graduate School of Management and the Graduate School of Business at the University of Chicago, in addition to dozens of seminars on financial derivatives. Before joining JPMorgan in 1993, he worked as a derivatives researcher at Discount Corporation of New York Futures, and as a software developer at Microsoft. Mr. Hanweck holds a PhD in Managerial Economics and Decision Science from the Kellogg Graduate School of Management at Northwestern University, and an AB in Mathematics from Princeton University.

- **FRIDAY 1:00 PM – Hanweck Associates, LLC Startup Presentation**

Hardwick, Jim
Senior Software Engineer, TechniScan Medical Systems

Jim Hardwick is the Senior Software Engineer at TechniScan Medical Systems in Salt Lake City, Utah. He has a B.S. degree in Computer Science from Westminster College and previously worked at GE Healthcare and OEC Medical Systems. Jim's expertise is in medical device development, including software and hardware development and human-computer interaction.

- **THURSDAY 11:00 AM – Embedded Tesla-Using CUDA C and Tesla in a Medical Device**

Harris, Christopher
Research Associate, The University of Western Australia

Christopher Harris is a research associate with the International Centre for Radio Astronomy Research (ICRAR) at the University of Western Australia. His research investigates the processing of scientific algorithms using parallel computing architectures, in particular the Graphics Processing Unit (GPU). Christopher has worked with GPUs since the early GPGPU era, using OpenGL to implement fuzzy clustering algorithms. He has since gone on to explore the use of GPU computing to accelerate radio astronomy signal correlation, resulting in performance gains of x10 to x100. His current research application areas include radio signal processing and image synthesis.

- **FRIDAY 4:00 PM – Astronomy & Astrophysics**

Harris, Jerry
Sr Computer Scientist II, Adobe Systems

After finding accounting too mind numbing, Jerry Harris switched from an Accounting major to a Computer major during during his junior year at Harding University. Jerry graduated as a BA Business Systems Analyst in 1983. After a short stint as a Business System Analyst, Jerry found the work too dull and decided to pursue digital painting after being inspired by Apple's Lisas, and Macintosh computers. Shortly thereafter Jerry founded Pixel Resources which produced 6 versions of PixelPaint. PixelPaint was the first commercial color GUI paint program for the Macintosh II. Initially sold through SuperMac technologies, then later directly from the developer in 1995. Jerry worked on the Albert 2D/3D graphics system while at Apple on the Pink or Taligent project from 1989-1996. In 1996, formed VisualWare and created the Typefx technology. Visualware was acquired by Adobe in October of 1996, and TypeFX served as the basis for Layer Effects in Photoshop. Since joining Adobe, Jerry has participated in the engineering of the following features: painting, HDR, layer effects, warping, and GPU enablement.

- **THURSDAY 2:00 PM – Adobe Photoshop CS4 and GPUs**

Hayes, David
CEO, Innovative Converged Devices

David is the CEO of ICD and has a wealth of experience in the mobile phone and consumer electronics markets. David founded Velocity Mobile – a collaboration with Inventec in 2007. Prior to ICD, David was CEO of A Living Picture plc (ALP), a company he formed to develop Momento an advanced digital picture fame technology. ALP was acquired by i-mate in December 2006 and David subsequently became their Chief Technology Officer. Prior to forming ICD and ALP, David was Chief Executive of DAT plc who were one of the pioneers of over-the-air device management for Windows mobile devices. Whilst at DAT, David made strong links with the consumer electronic device community from ODM's, operators and retailers.

- **THURSDAY 3:00 PM – Innovative Converged Devices Ltd Startup Presentation**

Heck, Michael
Technology Advisor, Visualization Sciences Group (VSG)

Mike is the Technology Advisor at Visualization Sciences Group (formerly part of Mercury Computer Systems). He evangelizes the use of 3D visualization, works with customers to understand their requirements, helps customers apply visualization technology to meet current and future requirements and advises the development of visualization technology at VSG to meet future requirements. Mike has been involved in implementing, managing, teaching and applying 3D visualization software for over 20 years. During that time he has been a

speaker at conferences including SEG and the World Oil Visualization Showcase, an invited instructor for SIGGRAPH conference courses including Open Inventor and VRML, and author/co-author of several papers and technical articles on visualization and GPU computing.

- **FRIDAY 11:00 AM – Using the GPU for Gradient Reconstruction of Unstructured Meshes**

Heimberg, Ralf
CEO, ICIDO GmbH

> Industry Veteran with more than 30 years experience in engineering solutions
> Chairman of the Board at Triplan AG
> VP Central Europe Tecnomatix
> Previous positions at Adept Technology and General Motors Fanuc
> Mechanical Engineering degree from University of Kiel

- **THURSDY 4:30 PM – ICIDO GmbH Startup Presentation**

Helgason, David
CEO, Unity Technologies

David Helgason has served as CEO of the game technology company Unity Technologies since co-founding it in 2003. The vision is to democratize game development, and develop technology for the next generation of the industry. From in-browser MMO, through mobile, to social, casual, serious, and beyond. And it works: Unity is used by dozens of big game publishes and media companies, hundreds of smaller studios, and thousands of independent professionals, hobbyists, students, and 14 year old boys. David's background is in software, and an assortment of university dabbling.

- **THURSDAY 3:30 PM – Unity Technologies Startup Presentation**

Herbst, Jeff
Vice President of Business Development, NVIDIA

Jeff Herbst is the Vice President of Business Development at NVIDIA Corporation, the world leader in visual computing technologies (and inventor of the GPU) with yearly revenues of approximately $4.0 billion USD. In this role, which he has held since 2001, Jeff leads NVIDIA's worldwide business development efforts, including overall ecosystem development, mergers and acquisitions strategy, investments, partnerships and other strategic business relationships and transactions. Prior to NVIDIA, Jeff was the worldwide head of corporate and business development at AltaVista, and also served as general manager for a start-up focused on content delivery infrastructure for wireless networks. Earlier in his career, Jeff was a partner with the law firm of Wilson Sonsini where he specialized in corporate finance, joint ventures, mergers and acquisitions and other strategic business and intellectual property-related transactions. Jeff holds a B.S degree in Computer Science from Brown University, and a law degree from Stanford Law School.

- **THURSDAY 4:00 PM – GPU Ventures, for the Entrepreneur**

Herrmann, Frank
University of Maryland

- **FRIDAY 2:00 PM – New Insights into the Phase Space of Binary Black Holes Using GPU Technology**

Heyn, Toby
Research Assistant, University of Wisconsin – Madison

Toby is a graduate student at the University of Wisconsin -Madison working on his Master's degree. His is a research assistant for Prof. Dan Negrut, leader of the Simulation-Based Engineering Laboratroy. His research interests include high performance computing and massively multibody dynamics simulation.

- **FRIDAY 10:30 AM – Driving on Mars: Simulating Tracked Vehicle Operation on Granular Terrain**

Hinkel, Ed
SeniorEngineer – TotalView debugger, TotalView Technologies

- **THURSDAY 5:00 PM – The Art of Debugging for the CUDA Architecture**

Hook, James
Bloomberg

James Hook is a senior software engineer at Bloomberg L.P. He has a M.S. in Electrical and Computer Engineering from Carnegie Mellon University.

- **THURSDAY 2:00 PM – GPU Acceleration of Mortgage CMO Option Adjusted Spread**

Horvath, Chris
DigitalArtist, Industrial Light & Magic

- **FRIDAY 3:30 PM – The GPU Revolution in Film Production: Report from the Battlefield**

Huang, Jen-Hsun
CEO and Co-Founder, NVIDIA

Jen-Hsun Huang co-founded NVIDIA in 1993 and has served since its inception as president, chief executive officer, and a member of the board of directors. Under his leadership, NVIDIA invented—and led the development of—the graphics processing unit (GPU), pioneering its use in devices as varied as smart phones, PCs, cars, workstations, and supercomputers. NVIDIA GPUs deliver unmatched visual computing with breathtaking, interactive graphics that delight users, and massive parallel computing power that accelerates work on the world's most challenging technical problems. NVIDIA was named Company of the Year in 2007 by Forbes magazine and has ranked #1 over the past two years in Innovation in the Semiconductor industry by Fortune. Huang is a recipient of the Dr. Morris Chang Exemplary Leadership Award from the Fabless Semiconductor Association in recognition of his exceptional contributions to driving the development, innovation, growth, and long-term opportunities of the fabless semiconductor industry. He has received the Daniel J. Epstein Engineering Management Award from the University of Southern California, and an honorary doctorate from Oregon State University. Prior to founding NVIDIA, Huang held engineering, marketing, and general management positions at LSI Logic, and was a microprocessor designer at Advanced Micro Devices. He holds a BSEE degree from Oregon State University and an MSEE degree from Stanford University.

- **WEDNESDAY 1:00 PM – Opening Keynote**
- **FRIDAY 10:00 AM – Emerging Companies Summit Fireside Chat**

Humphrey, John
Senior Engineer, EM Photonics

John Humphrey is a member of the Accelerated Computing Solutions group at EM Photonics. He earned his MSEE degree from the University of Delaware studying the acceleration of electromagnetics algorithms using custom hardware platforms. At EM Photonics, he launched a GPU research effort in 2005 with a GPU-based FDTD solver based on OpenGL methods and then later explored working in CUDA. Since then, he has worked on accelerated algorithms in a variety of fields, including linear algebra solvers and computational fluid dynamics engines.

- **THURSDAY 2:00 PM – CULA: Robust GPU Accelerated Linear Algebra Libraries**
- **FRIDAY 3:00 PM – Computational Fluid Dynamics (CFD) for the GPU**

Kaehler, Ralf
KIPAC/SLAC

- **FRIDAY 2:30 PM – Visualizing the Universe: Raycasting Astrophysical Simulation Data**

Kaplan, Michael
mental images

Michael Kaplan is Vice President of Strategic Development at mental images where he leads iray development. Michael is one of the pioneers of 3D visualization and rendering with a distinguished 30-year career in the 3D industry, including leading roles at Adobe, Apple, Rocket Science Games and Ardent Computer prior to joining mental images in 2007.

- **FRIDAY 1:00 PM – iray®: Interactive, Zero Tuning, Consistent Rendering in mental ray, RealityServer and other Applications**

Karel, Ben
Stone Ridge Technology

Ben Karel is currently a graduate student in Computer Science at the University of Pennsylvania. He received his bachelors degree in Computer Science from the University of Delaware in 2009. Prior to starting at Penn, Ben worked as a summer intern for Stone Ridge Technology developing the Stone Ridge CUDA Visualization Framework. His interests include graphics, programming languages, and compilers.

- **FRIDAY 10:30 AM – SCVF: A Framework for Visualization of Grid and Particle Based CUDA Simulations**

Kartsaklis, Christos
Computational Scientist, Irish Centre for High-End Computing (ICHEC)

Christos Kartsaklis has received a BSc in Computing Science (1st Class Hons) from the University of Aberdeen (UK, 2002) and an MSc in High Performance Computing (Distinction) from the University of Edinburgh (UK, 2003). Christos has received an award in Software Engineering by ScotlandIS (2002) for his BSc thesis, has worked in the past for EPCC (UK, 2003-2005) as an HPC Applications Consultant, and is currently completing his PhD in Computer Science at the University of Manchester (UK, 2005-2009). He joined ICHEC as a Computational Scientist in May 2009. He has previously worked on distributed computing (peer-to-peer and agent-based), HPC library lifecycle (development, porting and optimisation) and HPC remote paging networks.

- **THURSDAY 11:00 AM – GPU Computing in Java: Generation and Scheduling of PTX Assembly**

Kato, Kimikazu
Researcher, Nihon Unisys, Ltd.

Kimikazu Kato received the M.S. in mathematical sciences, and the Ph.D in information technology from University of Tokyo in 1998 and 2008 respectively. Since 1998, he has been working for Nihon Unisys, Ltd. where he has been engaged in the research and development related to geometric computation and other mathematical algorithms. His research interests include computational geometry and discrete mathematics.

- **THURSDAY 3:00 PM – You Might Also Like: A Multi-GPU Recommendation System**

Kelmelis, Eric
CEO, EM Photonics

Eric Kelmelis is co-founder and CEO of EM Photonics, Inc. There his work has primarily focused on algorithm acceleration and solver development for the scientific computing, linear algebra, and image processing fields. Mr. Kelmelis received his Bachelors and Masters degrees in Electrical Engineering from the University of Delaware, where he currently serves as visiting faculty teaching high-performance computing and parallel programming. THURSDAY 2:00 PM – CULA: Robust GPU Accelerated Linear Algebra Libraries

- **THURSDAY 4:00 PM – GPU-based Libraries: Accelerating the Next Generation of Applications**

Kerris, Richard
CTO, Lucasfilm Ltd

Richard Kerris joined Lucasfilm in December of 2007. As Chief Technology Officer of one of the world's leading film and entertainment companies, Kerris oversees the company's technical operations and is responsible for the development and execution of technology strategy for Industrial Light & Magic, Skywalker Sound, LucasArts, Lucasfilm Animation, Lucasfilm Animation Singapore and Lucas Online. Prior to joining Lucasfilm, Kerris held key posts at Apple, Alias|Wavefront, Electric Image and Silicon Graphics.

- **FRIDAY 8:30 AM – Day 3 Keynote with Richard Kerris, Lucasfilm**

Keryell, Ronan
CSO, HPC Project

Ronan Keryell graduated from the École Normale Supérieure, in Paris, with a MSc in Computer Architecture (1988) and a PhD in Computer Science from ENS & Paris XI University (1992). While obtaining his PhD, he worked on the architecture of a SIMD parallel computer for high-end image synthesis (a non-integrated version of GPGPU in 1990...). From 1992 Ronan was associate professor at the CS research center (CRI) of Mines ParisTech where he worked on data parallel languages (HPF...), compilation and application optimizations within the PIPS project, a source-to-source compilation workbench specialized in interprocedural parallelization, using a polyhedral model. He was then associate professor from 1999 in the CS research center in the Telecommunication engineering school in Bretagne where he worked on secure processor architectures and compilers for parallel embedded architectures. He is currently CSO at HPC Project since

2007 and works on compilation and parallelization, application optimizations, for HPC and embedded systems.

- **THURSDAY 3:00 PM – Par4All: Auto-Parallelizing C and Fortran for the CUDA Architecture**

Kim, Taehoon
Co-founder, President, Nurien Software

Taehoon Kim is co-founder and president of Nurien Software, a leading developer and operator of 3D social networking and online games company. Taehoon is the visionary behind Nurien's innovative product, merging the social dynamics of Web with the appeal of massively multiplayer online game. Previously, Taehoon served as the President of Realtime Worlds Korea, the Korean branch of UK-based software technology and online entertainment company founded by David Jones, the creator of Grand Theft Auto and Lemmings. Prior to Realtime Worlds Korea, Taehoon served at Samsung Electronic's New Business Development team, where he spearheaded the creation of PC-Mobile cross-platform online gaming devices for the global market. Taehoon holds a B.S in Electrical and Computer Engineering and M.Eng.in Operations Research and Information Engineering from Cornell University.

- **THURSDAY 4:00 PM – Nurien Software Startup Presentation**

Kloeckner, Andreas
Graduate Student, Brown University

Together with his advisor Jan Hesthaven, Andreas is working on efficient solvers for the Vlasov-Maxwell model of plasma physics. His research aims to provide a particle-field coupling that reaps the benefits of high-order unstructured Discontinuous Galerkin (DG) Maxwell field solvers while remaining faithful to the complicated phase space behavior displayed by the model. A collaboration with Tim Warburton provided the opportunity to explore the use of GPUs (Graphics Processing Units) for DG solvers. This work led to the discovery of techniques that allow a single commodity GPU to perform DG computations around 50 times faster than a single core of a conventional CPU. One cornerstone of this performance gain is GPU metaprogramming, a technique supported by Andreas's PyCUDA library. After completing his degree in applied mathematics in Karlsruhe (Germany) working on the simulation of photonic crystals, Andreas joined Jan Hesthaven's group at Brown University to work towards a doctoral degree.

- **FRIDAY 1:00 PM – GPU Metaprogramming using PyCUDA: Methods & Applications**

Klug, Michael
CTO, Zebra Imaging

Michael Klug is the Chief Technology Officer and Co-Founder of Zebra Imaging. He received a Bachelor of Science degree and a Master of Science degree from MIT Architecture in 1989 and 1991 respectively. From 1991 until 1997, Mr. Klug worked as a Research Scientist at the MIT Media Laboratory focusing on design and development of 3D displays and holographic systems. During his tenure at MIT, he was responsible for basic technology development, prototype construction and demonstration, proposal generation and program budgeting and renewal, sponsor recruitment, coordination, presentations and demonstrations, and laboratory resource management and establishment of group research goals. He developed the basic technological predecessors of the systems now productized by Zebra Imaging, and has served as a consultant to various companies such as Polaroid Corporation, in related fields. He is recognized internationally as one of a handful of experts in the field of automated hologram printing technology. At Zebra Imaging, Mr. Klug is responsible for overall technology strategy, intellectual property development and management, and integration of R&D, business development, market research, and technology and product trajectories.

- **FRIDAY 4:00 PM – Zebra Imaging Startup Presentation**

Korngreen, Alon
Senior Lecturer of Cellular Neurophysiology, Bar-Ilan University

Dr. Alon Korngreen was born and raised in Israel. He received his Ph.D. in chemistry from Ben-Gurion University of the Negev, Israel in 1997. He then spent three years specializing in neurophysiology and electrophysiology in the lab of Nobel Laureate Prof. Bert Sakmann in Germany. He joined the faculty of Life Sciences at Bar-Ilan University in 2001 where he is currently a Senior Lecturer of Cellular Neurophysiology. His research touches one of the basic yet still unresolved questions in neuroscience: How do neurons process information? What is the neuronal code at the cellular level? During the last years he has been developing new computational techniques that will allow investigating these questions. One technique under development in the lab is the use of stochastic optimization algorithms for constraining compartmental models of neurons with non-homogenous distributions of ion channels and Markov chain models of voltage-gated ion channels.

- **FRIDAY 4:30 PM – Optimizing Ion Channel Kinetics Using A Massively Parallel Genetic Algorithm on the GPU**

Kuehne, Bob
Blue Newt Software

- **FRIDAY 4:00 PM Face Recognition for Photographs and Video**

Kretzler, Brian
NVIDIA

- **THURSDAY 11:00 AM, 2:00 PM, 3:00 PM, 4:00 PM, 5:00 PM – NEXUS Lab**
- **FRIDAY 10:30 AM, 1:00 PM, 2:00 PM, 3:00 PM, 4:00 PM – NEXUS Lab**

Kuo, Wesley
President, Ubitus Inc.

Wesley is a successful entrepreneur who previously founded iaSolution in 2000, a mobile software company in Asia with some leading clients including China Mobile, HTC, Lenovo etc.
In 2004, he accomplished a USD 70 Million merger with the Japan listed embedded software company - Aplix Corporation, and became the Board and Chief Strategy Officer of Aplix, primarily responsible for international sales and marketing and helped the company to conclude major OEMs like Samsung and Nokia. In 2007 May, he founded Ubitus Inc. as CEO focusing on cloud computing for rich video/graphics processing till today.

- **FRIDAY 2:30 PM – Ubitus, Inc. Startup Presentation**

Le Grand, Scott
Senior Engineer, NVIDIA Corporation

Scott Le Grand is a senior engineer on the CUDA software team at NVIDIA. His previous commercial projects include the game BattleSphere for the Atari Jaguar; Genesis, the first molecular modeling system for home computers, for the Atari ST; and Folderol, the first distributed computing project targeted at the protein folding problem. More recently, he ported the Folding@Home and AMBER molecular modelling codebases to CUDA and he is deeply involved in applying CUDA to problems in computational biology. In a former life, he picked up a B.S. in biology from Siena College and a Ph.D. in biochemistry from The Pennsylvania State University. In addition, he wrote chapters for ShaderX and GPU Gems 3 and coedited a book on computational methods of protein structure prediction.

➔ **THURSDAY 2:30 PM – GPU Accelerated Computational Biophysics in NAMD**

Lecomber, David
CTO, Allinea Software

Dr David Lecomber is one of the founders of Allinea and leads the research, development and support teams behind its software products. David's history in High Performance Computing began with the Oxford BSP group in 1993, working on an early alternative model for parallel programming to the emerging but complex MPI standard. He obtained a DPhil in Parallel Computing, producing work on the simulation of shared-memory systems using, and formal semantics for, distributed-memory clusters. He subsequently continued to work at Oxford in post-doctoral and teaching positions, researching parallel libraries and languages. After two years developing software for high-volume on-line services on clusters of Java machines, he returned to the fold of High Performance Computing at Allinea, researching the development tools needed for parallel programming.

➔ **FRIDAY 3:00 PM – GPU Debugging Made Easy with DDT**

Lichtenbelt, Barthold
NVIDIA

Barthold Lichtenbelt received an MSEE from the University of Twente, The Netherlands, in 1994. Starting that year, he has worked for HP, Dynamic Pictures and 3Dlabs where he has been involved in various aspects of computer graphics and driver development. He joined NVIDIA in 2005, where he currently manages the core OpenGL driver team. Barthold is also the chair of the Khronos OpenGL ARB Working Group and has been involved in enhancing OpenGL since version 1.3. Barthold is the author of the book "Introduction to Volume Rendering", and a contributor to the book "OpenGL Shading Language".

➔ **WEDNESDAY 9:00 AM – OpenGL: The Train Has Left the Station (Pre-Conference Tutorial)**

Lionetti, Fred
Researcher, University of California, San Diego

Fred Lionetti is the lead developer of Continuity 6, a problem-solving environment for multi-scale modeling in bioengineering and physiology. Working with Prof. Andrew McCulloch, principle investigator of the Cardiac Mechanics Research Group of the University of California-San Diego, Fred has developed this modeling resource for over 5 years. He is also a graduate student in Prof. Scott Baden's research group at UCSD. Fred focuses his research on scientific computation and parallel programming, particularly the applicability of graphics processors to cardiac modeling.

➔ **THURSDAY 5:00 PM – GPU Accelerated Solvers for ODEs Describing Cardiac Membrane Equations**

Ljung, Patric
Research Scientist, Siemens Corporate Research

Patric Ljung received his MS degree in Information Technology in 2000 from Linköping University, Sweden, and the PhD degree in Scientific Visualization in 2006 from Linköping University. Between 2000 and 2007 he was a faculty staff member at Norrköping Visualization and Interaction Studio (NVIS) and the Visual Information Technology and Applications (VITA) group, also collaborating with CMIV at Linköping University. Since 2007 he is a Research Scientist in the Imaging & Visualization Department at Siemens Corporate Research in Princeton, NJ, USA. From 1989 until 1995 he worked as a Software Engineer with embedded and telecom systems. His research interest involves interactive visualization of large data sets, volume rendering, graphics system design, and software design and engineering. He is a member of the IEEE Computer Society, Eurographics, and ACM/SIGGRAPH.

➔ **THURSDAY 2:30 PM – Efficient Integration of CUDA & OpenGL Proccessing in XIP**

Luebke, David
Senior Manager, NVIDIA Research

David Luebke helped found NVIDIA Research in 2006 after eight years on the faculty of the University of Virginia. Luebke received his Ph.D. under Fred Brooks at the University of North Carolina in 1998. His principal research interests are GPU computing and real-time computer graphics. Luebke's honors include the NVIDIA Distinguished Inventor award, the NSF CAREER and DOE Early Career PI awards, and the ACM Symposium on Interactive 3D Graphics "Test of Time Award". Dr. Luebke has co-authored a book, a SIGGRAPH Electronic Theater piece, a major museum exhibit visited by over 110,000 people, and dozens of papers, articles, chapters, and patents.

➔ **WEDNESDAY 3:00 PM – Important Trends in Visual Computing (General Session)**

Lutes, James
Software Developer (Imaging Geometry), PCI Geomatics

James Lutes is originally from New Brunswick, in Eastern Canada. He has a BScEng and MScEng in Geodesy and Geomatics Engineering from the University of New Brunswick, where he specialized in Engineering Surveys. Since 2001, he has been working in the field of satellite photogrammetry, first with Space Imaging in Denver (now known as GeoEye) where he supported photogrammetric production and satellite sensor calibration for the IKONOS and GeoEye-1 satellites. Currently, James works with PCI Geomatics in Toronto, Canada where he is developing new sensor support and re-implementing key image processing modules to take advantage of modern hardware.

➔ **THURSDAY 5:00 PM – Mapping Satellite Imagery on the GPU: Fast Orthorectification and Pan-Sharpening**

Lu, Peter
Post-Doctoral Research Fellow, Harvard University

Dr. Peter J. Lu is a post-doctoral research fellow in the deparment of physics and SEAS, Harvard University, Cambridge, MA. He received his AB summa cum laude in physics (2000) from Princeton University, and MA (2002) and PhD (2008) in physics from Harvard University. His main focus is on the physics of attractive colloids, with a current focus on high-performance image and data analysis with GPGPU. In addition, he has also published papers on the mathematics of islamic architectural tilings, early precision compound machines from ancient China, the first use of diamond in prehistoric China, and the discovery of the first natural quasicrystal. In addition, he has developed a real-time target-locking system for use in confocal microscopy.

- **THURSDAY 5:30 PM – Directing Experiments in the International Space Station With GPU-Assisted Image Analysis**

Lynch, Michael
CEO, Right Hemisphere

Michael Lynch leads Right Hemisphere. He is a successful software executive who helped grow 7th Level from 20 employees to a $70 million IPO. Michael has produced over thirty award winning Internet, education and entertainment software products for Learn2.com, Disney, Real Networks, IBM, Microsoft and Sony. Michael has been CEO of Right Hemisphere since June 2000 and transformed the company from a small tools provider for graphics professionals to the global leader in visual product Communications and collaboration for Global 1000 manufacturers.

- **FRIDAY 1:00 PM – Right Hemisphere Startup Presentation**

MacIntyre, Blair
Associate Professor, Georgia Tech, GVU Center

Blair MacIntyre is an Associate Professor in the School of Interactive Computing at the Georgia Tech, where he has directed the GVU Center's Augmented Environments Lab for 10 years. He has been conducting Augmented Reality research since 1991, with the goal of understanding the potential of AR as a new medium for games, entertainment, education and work. He has collaborated on a variety of AR gaming and entertainment projects over the years, and in the past few years has focused on handheld AR game design, interaction and evaluation. His research is supported by a variety of technology and content companies, including MediaPower, Turner Broadcasting, Sun Microsystems, Hewlett-Packard, Alcatel-Lucent, NVIDIA, Apple, Texas Instruments, Samsung, Qualcomm and Motorola and Nokia. He has written a wide range of academic papers, taught tutorials on Handheld Augmented Reality, and been Program Chair for the International Symposia on Wearable Computing (2000), Mixed and Augmented Reality (2003), and User-Interface Software and Technology (2003). He served as the guest editor of a 'Mixed Reality' special issue of IEEE Computer Graphics and Applications. He received a Ph.D. from Columbia University in the City of New York in 1998, and B.Math and M.Math degrees in Computer Science from the University of Waterloo in Canada in 1989 and 1991.

- **WEDNESDAY 3:00 PM – Important Trends in Visual Computing (General Session)**
- **FRIDAY 2:00 PM – Handheld Augmented Reality**
- **FRIDAY 3:00 PM – Zombies on Tegra: A Case Study in Mobile Augmented Reality**

MacLean, W. James
Senior Engineer, MDA Space Missions

Dr. MacLean has over 10 years of experience as a professor in Computer Science and Electrical & Computer Engineering at the University of Toronto. His research interests are in the field of Computer Vision, with special interests in motion analysis, stereo and high-performance implementations of computer vision algorithms. In addition to having published nearly 40 research papers, he has acted as program chair for the Workshop on Spatial Coherence for Visual Motion Analysis (European Conference on Computer Vision, 2004), and been the guest editor for a special issue of Computer Vision and Image Understanding on the same topic. He is a Senior Member of the IEEE and a member of Professional Engineers Ontario. He is currently a Senior Engineer with MDA Space Missions in Brampton, Canada, where his work involves computer vision and signal processing for both space and terrestrial robotics applications.

- **THURSDAY 4:30 PM – Particle Filters and Sequential Monte Carlo Techniques Using GPUs**

Madsen, Justin
Research Assistant, University of Wisconsin – Madison

Justin is a graduate student at University of Wisconsin-Madison, working as a Research Assistant for Prof. Dan Negrut at the Simulation Based Engineering Lab. He has worked on vehicle simulation and multi-body simulation of tracked vehicles. He also works as a Teaching Assistant for Mechanical Engineering undergraduate courses.

- **FRIDAY 10:30 AM – Driving on Mars: Simulating Tracked Vehicle Operation on Granular Terrain**

Mahajan, Satayan
CEO, Motus Games

Saty Mahajan, is a pioneer in the field of motion and Human-Computer Interaction. He has created a recognition and creation architecture for sensors that abstracts away the mundane details of sensor math and physics allowing the developer to focus purely on the user experience, ulitmately creating a single unifying software platform for the world of motion that is independent of the device and sensors. Saty worked prevously at Intel, Customachines, and studied Computer Engineering at the University of California, San Diego. In his spare time, Saty enjoys writing and playing music, dabbling in 3D modelling and visual effects and dreaming up the next big thing.

- **FRIDAY 2:00 PM – Motus Games Startup Presentation**

Malcolm, James
Chief Architect, AccelerEyes

James Malcolm is cofounder of AccelerEyes LLC and chief architect of the Jacket project, where he is responsible for software strategy and development. He holds degrees in Mathematics (BS), Computer Science (BS, MS), and Electrical Engineering (MS) from Georgia Institute of Technology. He can be reached at malcolm@accelereyes.com.

- **THURSDAY 4:30 PM – Advanced CUDA Development in MATLAB with Jacket**

Marrin, Mark
Chief Engineer, Zaxel Systems

Mark Marrin is the Chief Engineer at Zaxel Systems. He has 10 years experience developing soft real time systems on Windows. After obtaining his B.S in Computer Science from Carnegie Mellon University Mr. Marrin worked extensively on distributed scene capture and clustering software. This software supported the Virtual Viewpoint system – one of the world's first viable telepresence systems. Mr. Marrin has worked on the cutting edge of high resolution video for several years, leading a series of software projects that push the limits of PC hardware. These projects included a distributed HD playback system integrated with PowerPoint, a high end DDR, and a client/server based filter driver for transparent image compression.

- **FRIDAY 10:30 AM – Exploiting the GPU in Ultra High-End 4K Video Servers**

Mazhar, Hammad
Research Assistant, University of Wisconsin Madison

Hammad is an undergrate member of the Simulation Based Engineering Laboratory. His recent work has included assembling high performance computing hardware and implementing a fast multi-GPU collision detection algorithm. He intends on attending graduate school in the future.

- **FRIDAY 10:30 AM – Driving on Mars: Simulating Tracked Vehicle Operation on Granular Terrain**

McGovern, Myles
CEO, Immersive Media

Myles McGovern's leadership has provided strategic organizational vision and identified sustainable market opportunities, guiding IM's transformation from an early stage technology into a world class visual medium across a variety of markets. Prior to joining Immersive Media in 2004, Myles was the founder, president and chief executive officer of Centrinty where he spearheaded the Company's impressive growth in North America, Europe and Asia. Myles worked at Xerox for ten years, developing significant experience in technology. He has twice been nominated as Canadian Entrepreneur of the Year (1999, 2001).

- **THURSDAY 3:00 PM – Immersive Media Startup Presentaiton**

Melonakos, John
CEO, AccelerEyes

John is co-founder and CEO of AccelerEyes, located in Atlanta, GA and founded in 2007. AccelerEyes sells software which makes code run faster for scientists, engineers and financial analysts. Its product, Jacket, enables MATLAB to run on the GPU (graphics processing unit).John also co-founded DivEyes, an Atlanta-based innovation generator for high-performance digital imaging and vision technologies. To date, DivEyes has produced two companies: AccelerEyes and CellularEyes. CellularEyes is developing software imaging tools for pathologists.During school, John worked at various large companies in summer stints, including General Electric, Hewlett-Packard, Lockheed Martin Aeronautics Company, and Corning Cable Systems. He was also an early employee of startup Torion Technologies. From 1999 to 2001, John was a missionary for the Church of Jesus Christ of Latter-day Saints in Maracaibo, Venezuela. John holds a BS in EE (electrical engineering) from Brigham Young University and MS & PhD ECE (electrical & computer engineering) degrees from the Georgia Institute of Technology. He is a graduate of Georgia Tech's TI:GER program and Sam Nunn Security program. He was part of the winning team in the 2008 Georgia Tech and Georgia Bowl Business Plan Competitions as well as Georgia's 2009 GRA/TAG Business Launch Competition.

- **THURSDAY 5:00 PM – MATLAB GPU Computing Essentials Tutorial**

Messmer, Peter
VP Space Applications, Tech-X Corporation

Dr. Messmer received his Ph.D from the Institute of Astronomy, ETH Z rich, Switzerland. Dr. Messmer's research interests and skills include numerical plasma physics and its application to astronomy, parallel and distributed computing, algorithms, software tools and applications, grid technologies and Web technologies.

- **THURSDAY 4:00 PM – GPU-based Libraries: Accelerating the Next Generation of Applications**

Micikevicius, Paulius
NVIDIA

Dr. Paulius Micikevicius is a Developer Technology Engineer at NVIDIA, focusing on parallel computing. His main interests include parallel algorithms, application of many-core architectures, and graph theory. Prior to joining NVIDIA, he was an assistant professor of Computer Science at Armstrong Atlantic State University as well as a research associate at the Media Convergence Laboratory at UCF. Paulius holds a PhD in Computer Science from the University of Central Florida.

- **THURSDAY 2:00 PM – C for CUDA – Advanced**
- **FRIDAY 10:30 AM – Implementing 3D Finite Difference Codes on the GPU**
- **FRIDAY 3:30 PM – The Art of Performance Tuning for the CUDA Architecture**

Murphy, KC
CEO, NextIO, Inc.

- **FRIDAY 1:30 PM – NextIO, Inc. Startup Presentation**

Muyan-Ozcelik, Pinar
PhD Student, UC Davis

- **FRIDAY 2:00 PM – Speed Limit Recognition System (Optimized for processing on the GPU)**

Nack, John
Photoshop Product Mgr, Adobe Systems

John Nack is Principal Product Manager for Adobe Photoshop. Prior to joining Adobe, John was a Web designer and animator at AGENCY.COM New York, where we developed online content for clients such as Gucci, Nike, British Airways, and Coca-Cola. That experience convinced him that design tools could and had to be better, leading him to join Adobe in 2000, working on the Web animation tool LiveMotion. John has been on the Photoshop team since 2002 and has worked on features such as Adobe Camera Raw, Smart Objects, Vanishing Point, the File Browser, and Adobe Bridge.

- **THURSDAY 2:00 PM – Adobe Photoshop CS4 and GPUs**

Negrut, Dan
Professor, University of Wisconsin – Madison

Dan Negrut received his Ph.D. in Mechanical Engineering from the University of Iowa in 1998. He spent six years working for Mechanical Dynamics, Inc., a software company in Ann Arbor, Michigan. In 2004 he served as an Adjunct Assistant Professor in the Department of Mathematics at the University of Michigan, Ann Arbor. He spent 2005 as a Visiting Scientist at Argonne National Laboratory in the Mathematics and Computer Science Division. At the end of 2005 Dan joined the Mechanical Engineering faculty at the University of Wisconsin, Madison. His interests are in Computational Science and he leads the Simulation-Based Engineering Lab at Wisconsin. Ongoing projects focus on large scale multibody dynamics, uncertainty quantification, numerical integration methods for dynamic systems, virtual prototyping in engineering design, and real-time simulation of ground vehicle systems. For his research and educational initiatives Dan received in 2009 a National Science Foundation Career Award.

- **FRIDAY 10:30 AM – Driving on Mars: Simulating Tracked Vehicle Operation on Granular Terrain**

Neighoff, Todd
System Architect, The Johns Hopkins Applied Physics Laboratory

Mr. Todd Neighoff is a system architect and systems engineer performing research and development at the Johns Hopkins University Applied Physics Laboratory since 1997. Todd specializes in real time systems and imaging sensor systems. He graduated with a B.S. in Computer Science from The University of Maryland in 1984 and an M.S. in the same field from The Johns Hopkins University Whiting School of Engineering in 1989 THURSDAY 4:30 PM – Real Time Multi-Channel, Multi-Camera Image Processing using CUDA

Ni, Tianyun
3D Graphics Engineer, NVIDIA

Tianyun Ni received her Ph.D. from the University of Florida in 2008. At school her research interests include modeling and computing with geometry, real-time higher order surfaces construction on the GPUs. She has published a number of journal and conference papers on these topics. She now works as a 3D Graphics engineer in NVIDIA's Developer Technology team where she develops new graphics techniques and helps game developers to incorporate these techniques into their games. Her recent work involves in finding applications for Direct3D 11 and developing advanced technologies to harness the computing power of next-generation GPUs, especially in the area of hardware tessellation. Her current research interests also include general-purpose computation on GPUs and real-time global illumination.

- **THURSDAY 3:00 PM – DirectCompute**

Oberbrunner, Gary
Chief Scientist, GenArts, Inc.

An MIT graduate, Gary has been developing graphics software for more than 25 years. A visionary developer with a brilliant mind, Gary is responsible for the continued evolution of Sapphire and oversees all aspects of technology and development at GenArts. He started his career in 1984 at Masscomp as a hardware engineer, but his interest in parallel supercomputer graphics soon led him to Thinking Machines Corporation. In 1993 he moved to Advanced Visual Systems, where as Chief Architect he developed high-end scientific visualization software and visual programming systems. He joined Karl at GenArts in 1997 and they have worked closely together ever since. In addition to his work at GenArts, he is a Director of the Open Effects Association. His research interests include image processing, parallel programming systems, visualization algorithms and advanced computer architectures.

- **FRIDAY 10:30 AM – Real-time Visual Effects in Film and Video: the GPU Changes Everything**

OBrien, Sean
President, NavTech

- **THURSDAY 5:30 PM – Visual Innovation in the Classroom**

O'Brien, Kevan
Open Workflows Specialist, Adobe Systems

Kevan joined Adobe systems after nearly 17 years in the post production industry and for the last five years has been working on promoting Adobe's film and video tools. Kevan's views on open workflows and technology uses are often radical and allow users to deliver award winning results for film and TV at a variety of budget levels. When not "selling snake oil" or "repairing broken dreams", Kevan can be found working on his own production pieces as well as working out how the US health system actually works from an end user perspective.

- **FRIDAY 10:30 AM – Driving on Mars: Simulating Tracked Vehicle Operation on Granular Terrain**

Obukhov, Anton
Developer Technology Engineer, NVIDIA

Anton Obukhov is a Developer Technology Engineer at NVIDIA Corporation since 2008. His field of interests include multimedia processing, video encoding, and computer vision. He graduated from Moscow State University in 2008 with a Masters Degree in Computer Science from the Computational Mathematics and Cybernetics Department in Russia. Before joining NVIDIA, he conducted research and development in the Graphics and Multimedia Lab at Moscow State University while also working at YUVsoft Corporation.

- **THURSDAY 11:00 AM – Performance Primitives for Video Codec and Image Processing**
- **FRIDAY 4:30 PM – Face Detection**

Okoniewski, Michal
Interim President & CEO, Acceleware

Chief Scientist since the company's inception, Michal is a Professor in the Electrical & Computer Engineering Department, University of Calgary. He is holding the Alvin Libin Ingenuity Chair in Bio-medical Engineering, and the Canada Research Chair in Applied Electromagnetics. Michal is a renowned expert in applied computational electrodynamics. Michal's research group is also working on 3D imaging technology that can look into living cells and micro-machined devices that can provide dramatic improvement to RF devices, such as cellular phones. Dr. Okoniewski has a Ph.D. in Electrical Engineering from the Technical University, Gdansk.

- **THURSDAY 11:00 AM – Sparse Linear Algebra – Iterative Solvers and Preconditioners on GPUs for Engineering Simulations**
- **THURSDAY 5:00 PM – Acceleware Startup Presentation**

Ota, Jeff
BMW Technology Office USA

Jeff Ota is a Senior Advanced Technology Engineer for the BMW Technology Office Palo Alto. His research topics include machine vision, embedded systems, displays, and advanced diagnostics. For the past 3.5 years, Jeff has focused on scouting, prototyping, and integrating non-automotive technologies in these fields into BMW's cars.

- **FRIDAY 2:00 PM – Speed Limit Recognition System (Optimized for processing on the GPU)**

Owens, John
UC Davis

John Owens is an associate professor of electrical and computer engineering at the University of California, Davis, where he leads a research group with a focus on GPU computing, both fundamental primitives and applications. He is also one of the lead developers of the popular CUDPP GPU computing library. At UC Davis, he earned the Department of Energy Early Career Principal Investigator Award (2004), an NVIDIA Teaching Fellowship (2006), and his department's Graduate Student Association Award for Graduate Teaching and Mentorship (2009). John earned his Ph.D. in electrical engineering in 2003 from Stanford University and his B.S. in electrical engineering and computer sciences in 1995 from the University of California, Berkeley.

- **FRIDAY 2:00 PM – Speed Limit Recognition System (Optimized for processing on the GPU)**

Paladini, Gianluca
Program Manager, Siemens Corporate Research

Gianluca Paladini is a Program Manager in the Imaging & Visualization Department at Siemens Corporate Research (SCR) in Princeton, NJ, where he heads the Imaging Architectures research program. He specializes in 3D/4D medical imaging, and he is the software architect responsible for the design of IVT, the core Imaging and Visualization Toolkit used by all Siemens Healthcare modalities to develop their Syngo imaging applications. He is responsible for XIP (eXtensible Imaging Platform) rapid application framework funded by the National Cancer Institute within caBIG (Cancer Bio-Informatics Grid), and he is currently managing numerous Corporate Technology projects in high performance GPU-accelerated parallel computing, with applications in reconstruction, registration, segmentation and visualization for diagnostic imaging, cellular imaging, pre-clinical imaging and clinical trials.

- **THURSDAY 2:00 PM – State of the Art in GPU-Accelerated Techniques for Medical Imaging**

Parker, Steven
Principal Research Scientist, NVIDIA

Dr. Steven Parker is a Research Scientist at NVIDIA and leads the OptiX development team. He has been involved in interactive ray tracing for over a decade, and has been chief architect several ray tracing systems culminating in OptiX. In addition, he has published over 100 technical articles on graphics, high-performance computing and software engineering. He is adjunct faculty and teaches a graduate course on ray tracing at the University of Utah. Dr. Parker has a Ph.D. in Computer Science from the University of Utah and a B.S. in Electrical Engineering from the University of Oklahoma.

- **THURSDAY 2:00 PM – Interactive Ray Tracing with the OptiX ray tracing engine**

Passenger, Josh
Project Leader, CSIRO

Josh Passenger is the Project Leader for the Colonoscopy Simulation Project which aims to provide a no-risk-to-patient surgical simulation of the colonoscopy procedure for training and certification of gastroenterologists. Josh Passenger holds a Bachelor of Science from University of Queensland and has over 9 years of commercial software design and development experience. Josh's previous roles have included Technical Lead at lastminute.com and Technical Architecture at Gartmore Investment Management, London.

- **THURSDAY 4:30 PM – Harnessing the GPU for Surgical Training and Preoperative Planning**

Patel, Sanjay
CTO / Professor, MultiCoreWare / University of Illinois

Sanjay J. Patel is an Associate Professor of Electrical and Computer Engineering at the University of Illinois at Urbana-Champaign and the Chief Technology officer of MulticoreWare, Inc. At University of Illinois, he has an active research program investigating high-performance, massively parallel architectures, programming tools, and visual computing applications. At MulticoreWare, he is helping with the technology strategy for creating high-value tools and IP for the accelerated computing eco-system. He was previously the Chief Architect and Chief Technology Officer at AGEIA Technologies, prior to its acquisition by Nvidia Corp in 2008. Patel has a significant experience in multicore chip architecture, high performance computing architectures, parallel systems and programming tools having worked at Digital Equipment Corporation, Intel Corporation, HAL Computer Systems, and AGEIA Technologies. Patel earned his Bachelor (1990), Master of Science (1992) and Ph.D,(1999) in Computer Science and Engineering from the University of Michigan, Ann Arbor.

- **FRIDAY 10:30 AM – CUDA Kernel Performance Analysis Using NVPerfMon**

Peddie, Jon
Pres, Jon Peddie research

Dr. Jon Peddie is one of the pioneers of the graphics industry, starting his career in computer graphics in 1962. After the successful launch of several graphics manufacturing companies, Peddie began JPA in 1984 to provide comprehensive data, information and management expertise to the computer graphics industry. In 2001 Peddie left JPA and formed Jon Peddie Research (JPR) to provide customer intimate consulting and market forecasting services. Peddie lectures at numerous conferences on topics pertaining to graphics technology and the emerging trends in digital media technology. Recently named one of the most influential analysts, he is frequently quoted in trade and business publications, and contributes articles to numerous publications including as well as appearing on CNN and TechTV. Peddie is also the author of several books including Graphics User Interfaces and Graphics Standards (1992), High Resolution Graphics Display Systems, and Multimedia and Graphics Controllers, and a contributor to Advances in Modeling, Animation, and Rendering.

- **FRIDAY 10:00 AM – Fireside Chat with Jen-Hsun Huang and Jon Peddie**

Perlman, Steve
Founder and CEO, OnLive

Steve Perlman, OnLive founder, president & CEO, is an entrepreneur and inventor devoted to pioneering Internet, entertainment, multimedia, consumer electronics and communications technologies and services. Best known for the development of QuickTime®, WebTV® and Mova® Contour™ facial capture technologies, he has over 30 years of technology development experience, 19 years of start-up experience and a track record of bringing media-rich products and services quickly to market. In addition to having founded and operated multiple startup companies, Steve has been a Microsoft division president and a principal scientist at Apple Computer. Steve's technology work is built into all iPhones, video iPods, Macs and most PCs, and has been deployed by DirecTV, Dish, Comcast, Time Warner, Charter and Adelphia cable TV and satellite TV networks. Consumer products incorporating Steve's work have also been retailed by Sony, Philips, RCA, Panasonic, Mitsubishi, Samsung, Fujitsu, Motorola, Scientific Atlanta, Sega and Nintendo. Steve's recently patented facial motion capture technology, Mova Contour, was used for the production of the Academy Award®-winning photoreal computer-generated face of Brad Pitt in "The Curious Case of Benjamin Button" (2008), and of Edward Norton and Tim Roth in "The Incredible Hulk" (2008). Steve holds 90 US patents, and has more than 100 additional patents pending.

THURSDAY 2:00 PM – OnLive Startup Presentation

Pesavento, Gerry
CEO, IQ Engines

Gerry Pesavento is a Founder and CEO of IQ Engines. Prior to that he was the CEO of Teknovus, Inc. and Alloptic, Inc. He received his MBA from UC Berkeley.

THURSDAY 2:00 PM – IQ Engines Startup Presentation

Peters, Alan
CTO, Universal Robotics, Inc.

Dr. Alan Peters, Founder, Chief Technical Officer
The principal architect of Neocortex, Alan supervises engineering and software development at Universal. Alan also is an Associate Professor of Electrical Engineering at Vanderbilt University and Director of the Intelligent Robotics Laboratory, where he supervises research on ISAC, a humanoid robot. He has more than 50 publications and has secured several million dollars for research funding from the US Department of Defense and NASA. He has been part of a prestigious team of roboticists from University of Southern California, the Massachusetts Institute of Technology, the Georgia Institute of Technology and University of Massachusetts who conducted extensive research on NASA's humanoid robot, Robonaut, at the Johnson Space Center in Houston, TX. He holds a Bachelor of Arts from Oberlin College and an M.S.E.E and Ph.D. in Electrical Engineering from the University of Arizona.

FRIDAY 3:00 PM – Universal Robotics, Inc. Startup Presentation

Peters, David
CEO, Universal Robotics, Inc.

David Peters, Founder, Chief Executive Officer and Chairman of the Board. David founded Universal Robotics with his brothers Alan and Jon, launching operations in the spring of 2008. Under his leadership, the company has closed two rounds of financing, developed a market-shifting vision application, secured a channel partnership with a major manufacturer, and reached Alpha stage on its signature product. Prior to founding Universal, he was a successful entrepreneur in the entertainment industry for 17 years. As a producer, David has first-hand experience handling the operations details of motion picture construction, as well as the complex financing deal structures common to film funding. He has produced a dozen motion pictures, creating films which have generated over $100 Million in worldwide sales. He has a Bachelor of Fine Arts from the Cleveland Institute of Art.

FRIDAY 3:00 PM – Universal Robotics, Inc. Startup Presentation

Pfister, Hanspeter
Professor, Harvard University

Hanspeter Pfister is a Gordon McKay Professor of the Practice of Computer Science at Harvard's School of Engineering and Applied Sciences. His research lies at the intersection of visualization, computer graphics, and computer vision, and spans a wide range of topics, including scientific visualization, computational photography, point-based graphics, appearance modeling, 3D television, and face recognition. Previously, he spent 11 years at Mitsubishi Electric Research Laboratories, where he was most recently Associate Director and Senior Research Scientist.

THURSDAY 9:00 AM – Day 2 Keynote with Hanspeter Pfister, Harvard University

Phillips, James
Senior Research Programmer, University of Illinois at Urbana-Champaign

James Phillips is a Senior Research Programmer in the Theoretical and Computational Biophysics Group at the Beckman Institute for Advanced Science and Technology at the University of Illinois at Urbana-Champaign. He has a Ph.D. in Physics from the University of Illinois and a B.S. in Physics and Mathematics from Marquette University. Since 1999, James has been the lead developer of the highly scalable parallel molecular dynamics program NAMD, for which he received a Gordon Bell Award in 2002. His research interests include improving the performance and accuracy of biomolecular simulations through parallelization, optimization, hardware acceleration, better algorithms, and new methods.

THURSDAY 3:00 PM – Supercomputing Super Session

Pinto, Nicolas
MIT

Nicolas Pinto is a second-year PhD Student in Computational Neuroscience at MIT. He is currently a member of the DiCarlo Lab and the Sinha Lab at MIT, and the Visual Neuroscience Group at Harvard. His research interests lie at the intersection of Brain and Computer Sciences. The overarching goal of his research is to dramatically accelerate the development of computational theories of how the visual cortex accomplishes object recognition. In addition to advancing our understanding of how the brain works by generating new experimentally-testable hypotheses, this approach also holds great promise for the development of new artificial vision systems. A key innovation in his work is the ability to leverage the computational power of disruptive technologies like NVIDIA's GPUs to provide new insights

into this fundamental problem. Before coming to the United States, Nicolas studied in Brazil, South Korea, Switzerland and received two M.S. in Computer Science from France.

- **THURSDAY 5:00 PM – Unlocking Biologically-Inspired Computer Vision: a High-Throughput Approach**

Popkiewicz, Mark
CEO, MirriAd

Mark has extensive senior leadership experience in high growth companies like Braid Systems, Eicon Networks, SDX Business Systems, Lucent Technologies, Mobile Media, BBC Ventures Group and BBC Vecta. Originally with a technology and media background Mark has raised investment capital, grown businesses from small to large, set up operations globally (30 around the world) most recently in China and India and achieved market leadership with successful exits both public and trade. He is currently CEO of MirriAd a new company located in London, Mumbai, New York and Los Angeles, delivering new in-video embedded advertising services specifically aimed at content from the studio, broadcast and digital on line platforms.

- **THURSDAY 4:30 PM – MirriAd Startup Presentation**

Price, Daniel
EM Photonics

Dan Price is a member of the Accelerated Computing Solutions group at EM Photonics. After receiving an MSEE from the University of Delaware, he joined EM Photonics to work on accelerating computationally intense problems using commodity hardware platforms. His research has included the implementation of computational electromagnetic algorithms on GPUs, image processing algorithms for atmospheric compensation, and most recently dense linear algebra solvers using CUDA.

- **THURSDAY 2:00 PM – CULA: Robust GPU Accelerated Linear Algebra Libraries**
- **FRIDAY 3:00 PM – Computational Fluid Dynamics (CFD) for the GPU**

Pryor, Gallagher
CTO, AccelerEyes

Gallagher Pryor is the CTO and a founding member of AccelerEyes LLC. He is the inventor of Jacket product for GPU computing in MATLAB. Currently he oversees future product development and strategy. He is an expert on computer vision with an emphasis on high performance tracking for military applications. He has a BS in Computer Science from Georgia Tech. He is also a Ph.D. candidate in the Computer Science program at Georgia Tech.

- **THURSDAY 4:00 PM – GPU-based Libraries: Accelerating the Next Generation of Applications**

Purves, Steve
Technical Director, ffA

From an Electrical & Electronic Engineering background, Steve has expertise in software engineering, image processing and algorithm development applied to 3D seimsic data analysis. He has worked accross a spectrum of technical roles within ffA and has hands on experience of working with geoscientists in ffA's E&P partner organisations to define, lead development of and deliver innovative 3D seismic analysis within ffA's software products. He is responsible all aspects for ffA's research, technology development and software engineering from ffA's Newcastle upon Tyne base in the UK.

- **FRIDAY 2:30 PM – Foster Findlay Associates Startup Presentation**

Quillen, Alice
Associate Professor, University of Rochester

Alice Quillen is Associate Professor of Physics and Astronomy at the University of Rochester and is currently Vice Chair of the American Astronomical Society's Division of Dynamical Astronomy. Alice's research background is in observational infrared astronomy including the study of active galactic nuclei, the dynamics of stars and gas in galaxies and the dynamics of planetesimals and gas in protostellar disks. Alice recently has focused on understanding planet/disk dynamical interactions to predict the properties of planets and other bodies that reside in circumstellar disks. Alice's prediction of a planet in the Fomalhaut system in 2008 was confirmed with the discovery in 2008 of a planet in this system. She has a PhD in Physics from Caltech and BA in Physics from Harvard.

- **THURSDAY 5:00 PM – Multiparticle Simulation**
- **FRIDAY 4:00 PM – Astronomy & Astrophysics**

Ramey, Will
Product Manager, NVIDIA

As NVIDIA's Senior Product Manager for GPU Computing, Will helps define and promote platforms, libraries and developer tools for the CUDA architecture. Prior to joining NVIDIA in 2003, he managed an independent game studio and developed advanced technology for the entertainment industry as a product manager and software engineer. He holds a BA in Computer Science from Willamette University and completed the Japan Studies Program at Tokyo International University. Outside of work, Will learns something new every day, usually from his two kids. He enjoys hiking, camping, swimming, spending time with his wonderful wife, and playing The Game.

- **WEDNESDAY 9:00 AM – Languages, APIs and Development Tools for GPU Computing (Pre-Conference Tutorial)**

Ramos, Henry
Senior Field Technical Specialist, Lenovo

- **FRIDAY 10:30 AM – The High Life for High Performance**

Rasul, Kashif
Nomad Labs

Kasihf is currently finishing his Phd in Mathematics from Frei University, Berlin, Germany. He has previously worked as a software engineer for Visage Imaging, on the amira scientific visualization software. Prior to that he worked in the Cactus numerical relativity group at the Max-Planck Institute (Albert Einstein Institute) in Potsdam, as part of a European wide Grid-computing initiative, Gridlab. He is also the maintainer of the open-source mathematica cuda plugin for Mathematica.

- **THURSDAY 2:00 PM – Using CUDA within Mathematica**

Reid, Ian
Numerical Algorithms Group (NAG)

As Chief Commercial Officer for the NAG Group, Ian has responsibility for driving all aspects of commercial strategy. Ian has been with NAG for over 20 years and has

held various technical and commercial positions within the company during this time. Most recently he was Vice President for Business Development and he continues to lead a worldwide team with responsibility for developing strategic partnerships with software and hardware organisations. Ian is a member of the IFIP Working Group 2.5 on Numerical Software has acted as technical advisor on various high-end computing procurements. He earned his Ph.D. in Computer Arithmetic and his B.Sc. (Hons) class I in Mathematics and Computer Science at Lancaster University. Ian holds several non-executive directorships.

THURSDAY 5:00 PM – Banking on Monte Carlo

Reil, Torsten
CEO, NaturalMotion Limited

Torsten graduated with a BA in Biology from Oxford University and holds an MSc in Evolutionary and Adaptive Systems from Sussex University. Prior to founding NaturalMotion, Torsten was researching for a PhD in Complex Systems at Oxford University. In 2003, he was named as one of the world's top 100 innovators by MIT's Technology Review. Torsten developed NaturalMotion's core technology together with Colm Massey.

THURSDAY 2:00 PM – NaturalMotion Limited Startup Presentation

Rhodes, Paul
President, Evolved Machines, Inc.

Dr. Rhodes has been a Visiting Scholar at Stanford University and is the founder of Evolved Machines, Inc., a research organization pioneering the synthesis of artificial neural circuits and their application to olfaction and visual object recognition. The group is developing the first synthetic neural arrays which wire themselves by simulating neural circuit growth in 3-dimensions, and is the first company to harness the power of programmable GPU's for the simulation of neural computation, now achieving > 100-fold acceleration of the computing power of conventional cores. The company's goal is the development of the first generation of devices truly based on brain circuitry, pioneering the fusion of neuroscience and engineering to develop new categories of machines which embed some of the capacities of biological neural systems. Dr. Rhodes received a Ph.D. and M.S. in Neuroscience under Dr. Rodolfo Llinas at NYU Medical School, and received an M.S. in Physics from Stanford University after graduating with an A.B. in Physics, Magna cum Laude, from Harvard University.

FRIDAY 1:30 PM – Evolved Machines Startup Presentation

Robinson, David
President & CEO, TechniScan, Inc.

David Robinson is has spent virtually his entire career focused on the development and commercializing services and products based in the hard sciences. Mr. Robinson is currently the President & CEO of TechniScan, Inc. a medical device company engaged in the research, development, and commercial production of an ultrasound breast imaging system designed to better detect and diagnose breast cancer. The Company's first product, the SVARA Warm Bath Ultrasound is now in large scale clinical testing at clinics in the U.S at the University of California San Diego and at the Mayo Clinic in Rochester Minnesota. The Company's European partner Esaote is overseeing clinical testing in Germany at the University Clinic in Freiburg. SVARA produces three separate images; two images from sound transmission properties (one for the speed of sound and one for the attenuation of sound) and a 3D reflection image. These images, which have never before been available to radiologists, contain unique information about bulk tissue properties within the breast that may aid physicians in diagnosing breast cancer. TechniScan has applied for FDA 510(k) approval of the SVARA imaging system in anticipation of beginning commercial installations in the Unites States in mid 2010. Since 2001 TechniScan has raised more than $24 million in private equity and more than $5 million in Federal Grants from the National Institutes of Health and National Cancer Institute (NIH/NCI) which has supported product development.

THURSDAY 2:30 PM – TechniScan, Inc. Startup Presentation

Ramos, Henry
Senior Field Technical Specialist, Lenovo

FRIDAY 10:30 AM The High Life for High Performance

Rollin, Phillippe
Applied Engineer, NVIDIA

Philippe Rollin is currently Applied Engineer in the Professional Solution Group at NVIDIA where he is investigating new tools and technologies to bring the best out of our products. In the past, Philippe was a Technical Lead in the Developer Tools group at NVIDIA, working on a full featured realtime shader authoring and debugging environment FX Composer 2.x. Philippe graduated with an MS in Information Technology from EPITECH, Paris.

THURSDAY 5:00 PM – High Performance Remote Graphics

Rolston, PhD, David
CEO, Forterra Systems, Inc.

Dave Rolston has more than 35 years of experience in high tech. His experience spans a broad spectrum of industries, applications, and technologies including simulation and training, graphics applications, imagery, gaming, artificial intelligence, entertainment, and the early Internet. During his career, Dave has performed in various roles, including technical, business, operational, and general management assignments. His roles include VP of Engineering for ATI, CEO of MultiGen-Paradigm, and GM of the Advanced Graphics Division for Silicon Graphics. Dave has a BS in civil engineering, MS in industrial engineering, and PhD in computer science with emphasis in simulation and artificial intelligence. He holds several patents, has published a large number of technical papers and a best-selling book on artificial intelligence.

FRIDAY 3:00 PM – Forterra Systems, Inc. Startup Presentation

Rotow, Dimitri
Product Manager, Manifold

Dimitri Rotow is a product manager for Manifold.Net with a focus on high performance computing applications in GIS, Remote Sensing and DBMS. Prior to joining Manifold he was a General Manager at Intel Corporation where he was responsible for hardware and software projects involving workstations. While at Intel, Rotow also launched a project to re-write mathematics algorithms for Intel's supercomputer platforms and which also created Intel's first massively parallel mathematics

libraries. Rotow was one of the founders of Bell Technologies, which was acquired by Intel.

- **THURSDAY 2:30 – Raster Processing Using Multiple GPUs and CPUs**

Rubin, Eri
Head of CUDA Project, OptiTex, LTD.

Erin Rubin is head of the CUDA project for OptiTex Ltd., with 15 years of experience as software developer. He worked as a Senior Graphics Developer for IDT-E Toronto, on two feature animation movies and TV specials. Eri holds a Master of Science in Computer Graphics, Hebrew University in Jerusalem, Israel, a Bachelor of Science in Computer Science & Life Science [Magna Cum Laude], Hebrew University in Jerusalem, Israel. Eri also studied Animation for 3 years at Bezalel Arts Academy, Jerusalem.

- **THURSDAY 3:30 PM – OptiTex Startup Presentation**
- **FRIDAY 10:30 AM – Real Time Cloth: 3D Triangular Mesh Algorithms on the GPU**

Sahrai, Laleh-chloe
International Project Manager, ENODO

Laléh Sahrai is the International Project Manager for ENODO, which she joined in April 09. She is responsible for the market development of ENODO's next generation industrial simulation and real-time visualization tool, a technology adapting to a diversifying spectrum. Prior to her involvement with ENODO's enthusiastic team, she worked for over 10 years in Computer graphics for Entertainment, Film and Broadcast, contributing on TV series, Video Games, and Full3D feature films in prestigious post-production facilities in Montréal, Paris and Ho Chi Minh City (Vietnam) , as a 2D/3D graphic artist first , later as a VFX producer , and Head of Production. She now brings her share of experience, a hybrid combination of rendered and real time 3D, to the expending and energetic start-up.

- **THURSDAY 4:00 PM – ENODO Startup Presentation**

Sakharnykh, Nikolai
Developer Technology Engineer, NVIDIA

Nikolai Sakharnykh is a developer technology engineer at NVIDIA. He has worked with game developers providing support for graphics technology content. Recently he focused on GPU compute and CUDA. Currently he is working on CFD-related projects and supporting CUDA customers. His interests include computational fluid dynamics, sparse matrix solvers and visualization techniques. Nikolai graduated with honours from Moscow State University, the department of Computational Mathematics and Cybernetics as a specialist in applied mathematics and informatics.

- **FRIDAY 2:00 PM – Tridiagonal Solvers on the GPU and Applications to Fluid Simulation**

Salian, Satish
CUDA Tools Manager, NVIDIA

Satish Salian serves as the Manager for CUDA tools at NVIDIA, where he is responsible for the strategy, direction and development of CUDA tools and support for CUDA developers. He joined NVIDIA in 2001 and in his various management roles been responsible for the development of NVIDIA graphics tools – the NVIDIA's display control panel, the NVWMI tool and the NVAPI SDK. He holds a BS degree in Computer Engineering from Pune University, India.

- **THURSDAY 5:00 PM – The Art of Debugging for the CUDA Architecture**
- **FRIDAY 2:30 PM – Debugging Tools for CUDA**

Santo, Alex
DUX Soft Pvt. Ltd.

Since 1999, Alex's strong mix of creative and technical skills have led him to work with companies across the globe, from Italy, India, USA, Japan, and also to Canada, where he currently resides. His career has been focused in graphics related research, development, and production for film and television. Alex acted as the CTO and R&D Supervisor at Paprikaas in Bangalore. His work included the development of core libraries that helped integrate pipelines for in-house tools and applications with the entire infrastructure setup. He then did TD and consulting work in Italy and USA, implementing a variety of innovative tools and scripts into creative production environments. At DamnFx, he pioneered the design of a node based rendering pass system and shader creation tool that integrated with Maya, XSI, and 3DSMax for Mental Ray and 3Delight. Alex's recent research is focused on the design and implementation of production pipeline workflow management and real time rendering systems. Alex joined DUX Soft as CSA in Sept, 2008, focusing on the design and development of core libraries and the development of SWIVEL

- **FRIDAY 2:00 PM – DUX Soft Sratup Presentation**

Santos, Joe
Co-Founder, CEO & CTO, Xulu Entertainment

Joe Santos is a technologist and software architect focused on large scale software systems with architecture design and development methodologies that lead to high-performance and maintainability, while supporting extensive third party development. He has expertise in 3D graphics, real-time programming, networking, parallel processing, database technology, computational geometry, numerical methods, and applied physics. Joe leads development of Xulu's virtual world platform which includes DirectX10 graphics, multi-body physics simulation and modeling, highly scalable networking, and advanced tools that allow professionals and novices alike to generate highest quality content. Joe was an early employee at Cadence Design Systems, a world-leading electronic design automation company, where was a lead developer of one of its most profitable products, Virtuoso. He also created and led the development of OpenAcess, an open database and framework that is becoming an EDA industry-wide standard.

- **THURSDAY 4:30 PM – Xulu Entertainment Startup Presentation**

Sasaki, Gary
President, DIGDIA

Gary has been involved in the digital consumer electronics industry, plus various parts of its value chain for over fourteen years. Areas have included digital imaging/video, digital media, digital displays, digital cinema, digital cable, home integration, audio/video networking, telematics, and handheld products. Gary formed DIGDIA in 2004, and is focused on helping companies find and analyze growth opportunities in markets empowered by digital media. Customers include Samsung, Panasonic, Philips, Siemens, Intel, HP, Canon, Pioneer, Bosch, Bain & Co. and other companies. Gary has spoken at Digital

Hollywood, CABA, CableNEXT, Home Entertainment Show, Strategy Institute, Pepperdine University, Santa Clara University, IBDNetwork, SDForum and other events.

THURSDAY 11:00 AM – Digital 3D Entertainment

Selzer, Geoff
CEO, Emergent Game Technologies

Prior to accepting the leadership role at Emergent, Geoff consulted to video game and other media companies, specifically on the development and exploitation of cross-media intellectual property, fund raising and production for video game and other media companies. From 1995 through 1998, Geoff served as VP, Creative Development and Production for Disney Interactive, where he built, managed, and created the vision for a group of over 100 professionals responsible for the design, development and production of the interactive software titles at Disney Interactive. In addition to his work at Disney, Geoff enjoyed successful careers in publishing and investment banking with an emphasis on mergers and acquisitions. Geoff completed his MBA (with a concentration in finance/international business) at Northwestern University's Kellogg Graduate School of Business. He holds a Bachelor of Arts in Philosophy awarded by Beloit College in Wisconsin.

THURSDAY 1:30 PM – Emergent Game Technologies Startup Presentation

Shoemaker, Austin
CTO, Cooliris

Austin Shoemaker is the chief technical officer for Cooliris. Austin was a master's student in Computer Science at Stanford University specializing in artificial intelligence, and has stopped out to lead technology and product development for the Cooliris platform. Austin is fluent in Spanish and Mandarin Chinese. Prior to his involvement with Cooliris, Austin worked at Apple Computer for seven years, contributing to product development efforts in several divisions. In 1995, Austin began testing pre-release software and hardware for the company on a voluntary basis. In 1999 Austin joined Apple as a software intern in the HR division, and subsequently moved to the Mac OS X engineering team in 2001, where he worked with the Human Interface group to develop several user interface prototypes that substantially influenced product design over the years. Austin was an engineer on the team that created iPhoto, Apple's consumer digital photography solution, and later wrote the timeline and clip shelf UI for iMovie, Apple's consumer video editing software. Austin rejoined the operating system team in 2004 to work on Spotlight, the desktop search engine built into Mac OS X 10.4. Before working at Apple, Austin co-founded a dial-up Internet service provider that grew to span the west coast of Florida. The company, NetLine Communications, was sold in 1996 to Comcast Cable. As an undergraduate, Austin was on the rowing team and won Stanford's first rowing national championship in 2005.

FRIDAY 2:30 PM – Cooliris Startup Presentation

Silberholz, John
University of Maryland

John Silberholz is a math and computer science major in the honors program at the University of Maryland, College Park, and is a researcher for the Center for Scientific Computation And Mathematical Modeling at the university. His current research involves using GPUs to speed up difficult numerical relativity calculations. His research interests also include the development of effective metaheuristics for difficult combinatorial optimization problems, and he recently received a prestigious Barry M. Goldwater scholarship for his research in this area.

FRIDAY 1:30 PM – Applications of Graphics Processing Units to the Binary Black Hole Evolutions

Smith, Greg
Manager of Developer Tools, NVIDIA

Greg has managed the Nexus Analysis tools for almost 2 years. Prior to joining NVIDIA and the GPU revolution he worked as a software architect on CodeWarrior and Radix development tools for Sony and Nintendo platforms, and was a system and software architect on numerous embedded multimedia processors and platforms.

THURSDAY 11:00 AM, 2:00 PM, 3:00 PM, 4:00 PM, 5:00 PM – NEXUS Lab

FRIDAY 10:30 AM, 1:00 PM, 2:00 PM, 3:00 PM, 4:00 PM – NEXUS Lab

Solomon, Nanci
CEO, Xulu Entertainment

Co-Founder and CEO, Xulu Entertainment
Nanci Solomon has been an entrepreneur in technology and media for more than 20 years, and has been actively involved in virtual reality, user-interface, and high-end gaming since 1992. Prior to co-founding Xulu, Nanci was co-founder of Smart Machines, a robotics company. She was founder, publisher and CEO of ASIC Technology & News, a high-tech media company producing an award winning magazine, executive conferences, marketing software, and a tradeshow.
Other roles include industry analyst at Dataquest, ASIC applications and marketing at Fairchild Semiconductor, and founder of Tomorrow, Inc. a new media company for college students and young adults. She has also organized sports and celebrity fundraising events, and led several community/youth groups.

THURSDAY 4:30 PM – Xulu Entertainment Startup Presentation

Spare, Jim
President & CEO, Canesta, Inc.

Mr. Spare's expertise is in successfully launching new, innovative technology products and companies. He leverages his prior experience in product management, marketing, and business development to lead the Canesta team. Previously, Mr. Spare served as business development director for the Microsoft TV Platform Division where his team created and managed key alliances with OEMs and network operators including AT&T, Rogers Communications, and Motorola. In this role, Mr. Spare was a key driver of Microsoft's investments and strategic product plan for the cable and satellite industry in North America. Prior to Microsoft, Mr. Spare led a product management organization at General Instrument (now Motorola) that exceeded sales and profit expectations for the company's flagship line of new, interactive set-top boxes. His team grew the business from concept to $300Million in revenue in less than 3 years, propelling the business to the highest grossing and most profitable product line in the company. Most recently, Mr. Spare served as chief operating officer of INBOXTV Corporation where he managed marketing, engineering, business development and operations for this Internet startup. Mr. Spare holds an MBA from the Amos Tuck

School of Business at Dartmouth and a BSEE from North Carolina State University.

FRIDAY 3:30 PM – Canesta Startup Presentation

Spicer, Sean

Executive Vice President / CTO, Aqumin LLC

Sean Spicer has over 10 years of experience designing, implementing, and managing high-performance computing and computer graphics software technology systems. Sean started his career at Silicon Graphics developing some of the world's most advanced Volume-Visualization and Interpretation tool kits; enabling Biomedical, Aerospace, and Geophysical applications to handle data sets of massive scale. He has been awarded numerous US and international patents in the field of computer graphics and has held posts with ACM SIGGRAPH and the IEEE Subcommittee on Visualization and Parallel Graphics. Immediately prior to joining Aqumin, Sean spent seven years with Magic-Earth/Halliburton in various technology, management, and senior leadership positions. He left the company as Sr. Mergers & Acquisitions Manager at Halliburton's Landmark Division, where he led the company's software technology corporate development efforts. Sean holds a BSE in Mechanical Engineering and Physics from Duke University, an MS in Mechanical Engineering from Stanford University, and an MBA in Finance and General Management from the Tulane University Freeman School of Business.

FRIDAY 3:00 PM – Aqumin LLC Startup Presentation

Spitzer, John

Sr Director of Content Infrastructure, NVIDIA

John Spitzer is the Senior Director for Content Infrastructure at NVIDIA. Working his entire career in the sphere of 3D hardware development, John has worn many hats: OpenGL Architectural Review Board representative, 3D driver developer, 3D benchmark (SPECviewperf/SPECglperf) author, inter-corporate graphics performance committee chairman, application performance consultant, and worldwide developer technology director. Six years ago and not speaking a single word of Russian, John had the crazy idea of moving to Moscow and creating the world's most advanced testing automation system. His dream realized, John currently serves as general director of the NVIDIA branch there, which consists of 70 employees and a testing farm of over 500 computers. John holds bachelor's and master's degrees in computer science from Rice University.

THURSDAY 11:00 AM – Designed for CUDA Program Overview

THURSDAY 12:00 PM – Compute Application Testing

Srivastava, Amit

Chairman / CEO, DUX Soft Pvt. Ltd.

> A seasoned entrepreneur and technology enthusiast
> Started his first company as a junior in college that dealt with P2P and next generation internet technologies
> Business Development Manager for one of the biggest technology publications in Silicon Valley, Upside Media
> Teamed up with founder of PC, PC World, Mac World magazines to form Foundational Net, Inc. as a next generation internet technologies pioneering company with focus on pay-per-use software delivery and service oriented architectures
> Business graduate from University of California, Santa Barbara

FRIDAY 2:00 PM – DUX Soft Pvt. Ltd. Startup Presentation

Steele, Michael

General Manager, Visual Consumer Solutions, NVIDIA

Michael Steele is General Manager of the NVIDIA's Visual Consumer Solutions group, enabling a wide range of new GPU computing applications using CUDA technology. Steele has nearly twenty years of high-tech marketing and business experience in the PC and networking system and semiconductor industry. Prior to NVIDIA, Steele was VP of Marketing at AGEIA leading PhysX software and hardware technology enabling immersive video gaming with real-time physical simulations. Steele has also held a variety of leadership positions at RMI, AMD and NexGen.

THURSDAY 11:00 AM – Designed for CUDA Program Overview

Stevens, Randall

President / CEO, Mersive Technologies

Randall co-founded Mersive Technologies in 2004 along with Christopher Jaynes and Stephen Webb. Randall brings over 18 years of business operations, software development, and industry partnership knowledge to Mersive. Randall is also the Founder and Chairman of ArchVision, a software development firm specializing in 3D graphics technology for the design industry. Through ArchVision, Randall built relationships with the industry's leading design visualization software companies including Autodesk, Bentley Systems, Robert McNeel & Associates, Maxon, Lightworks, Alias, Newtek, EON Reality and SGI. Randall received his BA from the University Of Kentucky College Of Architecture.

FRIDAY 1:30 PM – Mersive Technologies Startup Presentation

Stich, Timo

Developer Technology Engineer, NVIDIA

Timo Stich is a Developer Technology Engineer for NVIDIA Corporation. His focus is on image processing applications of Graphics Processors. Prior to joining NVIDIA he was research staff in Computer Graphics and Image Processing at the Max-Planck-Institute for Computer Science and the Computer Graphics Lab of Brunswick University, Germany. He received a diploma degree in Computer Science from Mannheim University, Germany and a Ph.D. degree from the Brunswick University, Germany.

WEDNESDAY 10:45 AM – C on the GPU (Pre-Conference Tutorial)

FRIDAY 3:00 PM – Graphcuts with CUDA and Applications in Image Processing

Stone, John

Senior Research Programmer, University of Illinois at Urbana-Champaign

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.

- **THURSDAY 3:00 PM – GPU Accelerated Visualization and Analysis in VMD**

Sullivan, Michael
CEO, Affine Systems

Michael Sullivan is a co-founder of Affine Systems and has been CEO since its inception. He holds an S.M. in Applied Mathematics and an A.B. in Computer Science from Harvard University. Before Affine, Mike worked at the Institute for Defense Analyses (IDA) in Alexandria, VA, where he performed research and technical analysis in support of the Office of the Secretary of Defense and the Intelligence Community.

- **FRIDAY 1:00 PM – Affine Systems Startup Presentation**

Taylor, John
Leader, Computational and Simulation Sciences, CSIRO

Dr Taylor is currently science & business leader, CSIRO Computational and Simulation Sciences. Dr Taylor has written more than 140 articles and books on computational and simulation science, climate change, global biogeochemical cycles, air quality and environmental policy, from the local to the global scale, spanning science, impacts and environmental policy. Dr Taylor has worked as a Computational Scientist and group leader both at the Mathematics and Computer Science Division, Argonne National Laboratory and at the Atmospheric Science Division at Lawrence Livermore National Laboratory. Dr Taylor was Senior Fellow in the Computation Institute at the University of Chicago. Dr Taylor has served on the Advisory Panel of the Scientific Computing Division of US National Center for Atmospheric Research (NCAR) and the US National Energy Research Scientific Computing Center NUGEX Advisory Committee. Dr Taylor is a Fellow of the Clean Air Society of Australia and New Zealand.

- **THURSDAY 3:00 PM – Supercomputing Super Session**
- **FRIDAY 1:00 PM – Supercomputing Research Roundtable**

Tellerman, Shanna
CEO, Wild Pockets

Shanna Tellerman is the Founder and CEO of Sim Ops Studios, Inc., a spin-off company from Carnegie Mellon University formed in January 2006 and the developers of the innovative new platform Wild Pockets. Shanna has overseen operations, development, fund raising, and strategic relationships since the formation of Sim Ops. She has several years of experience working in the video game and emerging serious games industry with a focus on breaking down the barriers for creative development in 3D. Graduating with a Masters degree in Entertainment Technology from Carnegie Mellon University in May 2005, she was hired on as commercialization specialist at Shanna Tellerman is the Founder and CEO of Sim Ops Studios, Inc., a spin-off company from Carnegie Mellon University formed in January 2006 and the developers of the innovative new platform Wild Pockets. Shanna has overseen operations, development, fund raising, and strategic relationships since the formation of Sim Ops. She has several years of experience working in the video game and emerging serious games industry with a focus on breaking down the barriers for creative development in 3D. Graduating with a Masters degree in Entertainment Technology from Carnegie Mellon University in May 2005, she was hired on as commercialization specialist at the Entertainment Technology Center to continue development and investigate commercialization opportunities for the highly publicized simulation training platform technology Hazmat Hotzone. This technology platform launched the formation of Sim Ops Studios, for which Shanna negotiated an exclusive technology license from CMU. Shanna has experience in a variety of roles including strategic development, business development, project management, user-centered design and research, and 3D artistic production. She worked as a production intern for Electronic Arts on The Sims 2, the sequel to the best-selling PC game ever.

- **THURSDAY 4:00 PM – Wild Pockets Startup Presentation**

Tian, Xusheng
Bloomberg

Dr. Tian is a quant finance developer at Bloomberg L.P. He has a PhD in Electrical Engineering from Rensselaer Polytechnic Institute. Dr. Tian has worked various areas previously, including communication systems, stochastic modeling of network traffic, and video coding.

- **THURSDAY 2:00 PM – GPU Acceleration of Mortgage CMO Option Adjusted Spread**

Tombroff, Michel
CEO, SOFTKINETIC S.A.

Michel joined Softkinetic after 17 years in the software industry where he held senior engineering and management positions with companies in the start-up, pre-IPO and public stages. Prior to Softkinetic Michel spent 8 years at TIBCO Software (NASDAQ:TIBX) in Palo Alto (CA), Paris and London, where his last position was Vice President Sales. Before that, Michel spent 7 years at real-time operating system company Chorus Systems (acquired by Sun Microsystems in 1996), where his last role was Head of Engineering. Most recently Michel was Vice President Europe at Mathon Systems, the information risk management company founded by John Mathon (co-founder of TIBCO Software.) Michel received a B.S. in Electrical Engineering from University of Brussels and a Masters Degree in Computer Science from University of California, Santa Barbara.

- **THURSDAY 4:30 PM – SOFTKINETIC S.A. Startup Presentation**

Tonge, Richard
Senior Software Engineer – PhysX, NVIDIA

Richard Tonge is a Senior Software Engineer in the PhysX group at NVIDIA. He is the founder and technical lead of the GPU Rigid Body Dynamics team and was previously involved in adapting and optimizing the PhysX algorithms for the GPU using CUDA. Prior to NVIDIA he worked at AGEIA where he contributed to the architecture and algorithms of the PPU, the first custom chip for accelerating games physics. Prior to AGEIA he performed pioneering work on iterative Linear Complementarity Problem solvers for game physics at Mathengine plc.

Richard holds a BA/MA in computation from the University of Oxford.

- **FRIDAY 1:30 PM – Rigid Body, Cloth and Fluid Physics for CUDA**
- **FRIDAY 3:00 PM – Spatial Data Structures for Massively Parallel Computing**

Tonjum, Thorolf
Director of R&D, Stormfjord

Thorolf Horn Tonjum leads the research department of Stormfjord, a company specialising in leveraging the power of the Gpu to develop next generation tools for the Oil industry. The Stormfjord team, with 20 years background in the Oil industry, and 10 years in the programming field, has developed a competitive portfolio of leading technologies, that drive novel research projects in the Oil industry. By using simulation and visualisation of large data-sets, as information nexuses, where 3D is the information carrier, Stormfjord deliver tool sets to: plan, collaborate, and control complex scenarios. Stormfjord's key strength lies in the ability to bridge results from leading research with powerful industrial needs. Thorolf Horn Tonjum, holds an advanced research degree in Artificial Intelligence, a background in interactive use of 3D data, from game developer FunCom, and research and development of Gpu programming at NITH Polytechnic college.

- **THURSDAY 1:30 PM – Stormfjord Startup Presentation**
- **FRIDAY 4:30 PM – Industrial Simulation Using Massively Parallel Algorithms**

Trevett, Neil
NVIDIA

Neil has spent over twenty five years in the 3D graphics industry and is responsible at NVIDIA for enabling and encouraging compelling applications on mobile devices and Smartbooks. Previously, as Vice President of 3Dlabs, Neil was at the forefront of the silicon revolution bringing interactive 3D to the PC. Neil is currently the elected President of the Khronos Group where he initiated the OpenGL ES and OpenKODE working groups and chairs the OpenCL and EGL working groups that are defining industry standards for advanced compute, graphics and media processing on a wide range of mobile, embedded and desktop systems. Neil was elected President for eight consecutive years of the Web3D Consortium dedicated to creating open standards for communicating real-time 3D on the Internet.

- **WEDNESDAY 10:45 AM – OpenCL on the GPU (Pre-Conference Tutorial)**

True, Thomas
Applied Engineer, NVIDIA

Tom is an Applied Engineer in NVIDIA's Professional Solutions Group where he focuses on the use of GPUs in broadcast, video and film applications ranging from pre-visualization to post production and live to air. Prior to joining NVIDIA, Tom was an Applications Engineer at SGI. Thomas has a M.S. degree in Computer Science from the Graphic Lab at Brown University and a B.S. Degree from the Rochester Institute of Technology.

- **FRIDAY 2:00 PM – Programming the Professional Video Pipeline**

Uzzan, Bruno
CEO, Total Immersion

Bruno Uzzan oversees operations and business development for Total Immersion. He is principally responsible for building the company's client roster, including Renault, Peugeot, BMW, Disney, EADS, CBS, Thomson and SGI Japan. Before establishing Total Immersion, Uzzan served as a consultant for Pierre Henri Scacchi and Associates (Price Waterhouse Group). He holds a masters degree in management from the University of Paris Dauphine.

- **THURSDAY 3:00 PM – Total Immersion Startup Presentation**

Varah, Sean
CEO, MotionDSP Inc.

Dr. Sean Varah is CEO of MotionDSP, having founded the company in 2005. Previous to MotionDSP, he was co-founder and COO of Q Media Partners, a San Francisco-based television development company which closed first-look deals with HBO and ABC television. As Director of consumer technology investments at Sony Music's 550 Digital Media Ventures, he sourced and led the Series A investment in Keyhole Inc., which was acquired by Google in 2004 and is now Google Earth. Dr. Varah received a bachelor's degree from Stanford University and a doctorate from Columbia University.

- **THURSDAY 1:30 PM – MotionDSP Inc. Startup Presentation**

Venkataraman, Shalini
NVIDIA

Shalini Venkataraman is an applied engineer with NVIDIA's professional solutions group where she focuses on using GPU's to solve graphics and visualization problems in the medical and oil & gas communities. Prior to joining NVIDIA, she was a research staff in scientific visualization at several institutions including the Center for Computation and Technology at LSU and in Singapore, at the Institute of High-Performance Computing and the Center for Information-Enhanced Medicine. Her interests include scalable graphics and display environments, large volume visualization and higher bit depth rendering . She earned her Master's degree from the Electronic Visualization Lab at the University of Illinois-Chicago and B.Sc from the National University of Singapore.

- **FRIDAY 4:30 PM – 4D Volume Rendering**

Vetter, Jeffrey
Oak Ridge National Laboratory and Georgia Tech

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.

- **THURSDAY 3:00 PM – Supercomputing Super Session**

Wang, Peng
NVIDIA

Peng Wang is a member of the Developer Technology group at NVIDIA, where he works on GPU computing technologies. He received a Ph.D. in Computational Physics from Stanford University, where he worked on developing multi-physics computational fluid codes and

simulating astrophysical turbulences. He also got a M.S. in Physics and a B.S. in Scientific Computing from Nankai University, China.

- **FRIDAY 1:00 PM – Astrophysical Fluid Simulation Using Adaptive Meshes**
- **FRIDAY 2:00 PM – OpenCL Optimization**

Warburton, Timothy
Associate Professor, Rice University

Tim Warburton is an associate professor in the department of Computational and Applied Mathematics at Rice University. He has a number of research interests related to the numerical solution of partial differential equations. In particular: Analysis of high-order, unstructured, finite-element methods; Development of associated efficient algorithms, including parallel computational techniques and advanced preconditioning methods for solving linear systems; Scientific computing applications for these methods including electromagnetics, acoustics and computational fluid dynamics in complex geometrics; simplification and introduction of these techniques to academic, defense and industrial environments; development of artificial radiation boundary conditions for the time-domain Maxwell's equations.

- **FRIDAY 10:00 AM – Advanced Numeric Computing**

Wegener, Albert
CTO, Samplify Systems

Al Wegener is the CTO, chairman and founder of Samplify Systems, a venture-funded fabless semiconductor start-up whose compression products reduce bandwidth and storage bottlenecks in sampled data systems. Samplify's efficient algorithms effectively compress sampled data streams from A/D and to D/A converters by up to 8:1 at sampling rates up to 60 Gsamp/sec. Samplify's products offer lossless, fixed-rate and fixed-quality compression modes that users control in real time. Samplify customers are found in a range of applications such as medical imaging, wireless, and measurement systems where no compression solutions were previously available. Mr. Wegener is a DSP engineer, technical manager, and inventor with 25 years of experience in defense electronics, professional and consumer audio, and wireless applications. Mr. Wegener holds a BSEE from Bucknell University and an MSCS from Stanford University. Mr. Wegener holds fifteen issued U. S. patents and is named on seventeen additional, pending Samplify Systems patent applications.

- **THURSDAY 3:30 PM – GPU-based Decompression for Medical Imaging Applications**

Williams, Ian
NVIDIA

Ian Williams manages the Professional Graphics Applied Engineering group at NVIDIA. He has worked in NVIDIA's professional graphics applied engineering team for the last eight years and has been heavily involved in the design and development of NVIDIA's Quadro industry solutions such as Gsync and Quadro Plex. Prior to NVIDIA he worked at Silicon Graphics for 8 years in a variety of roles including sales engineering and application performance. For the past 5 years he has also chaired the SPECgpc committee which is responsible for development of the SPEC Viewperf benchmark. He is a Chartered Engineer with the Institute of Mechanical Engineers (UK) and holds a B.S. degree in Engineering Science and Technology from Loughborough University, UK as well as an Executive MBA from Pepperdine University.

- **THURSDAY 11:00 AM – Large Scale Visualization**
- **THURSDAY 2:00 PM – Have You Heard What a GPU Can Do? – A Revolution in Audio**

Wolfe, Michael
The Portland Group

Michael Wolfe has worked on compilers for the high performance and parallel computing market in the commercial and academic world for over 30 years. He joined The Portland Group, Inc., as a compiler engineer in 1996; his responsibilities and interests include deep compiler analysis and optimization. He was an associate professor at the Oregon Graduate Institute from 1988 until 1996, and was a cofounder and lead compiler engineer at Kuck and Associates, Inc., prior to that. He has published one textbook, "High Performance Compilers for Parallel Computing", a monograph, "Optimizing Supercompilers for Supercomputers," and a number of technical papers.

- **FRIDAY 11:00 AM – Programming NVIDIA GPUs using PGI Accelerator Fortran and C Compilers**

Yamnitsky, Boris
BorisFX

Boris Yamnitsky is the founder and the president of BorisFX, a Boston-based software firm that brings effects and graphics tools to the video post-production and broadcast markets. An emigrant from the former Soviet Union in the Seventies, Boris Yamnitsky received a degree in Mathematics from Boston University where he concentrated in Theoretical Computer Science and Linear Programming. His research paper with Prof. Levin, "An Old Linear Programming Algorithm Runs In Polynomial Time," was accepted to the 1982 IEE FOCS conference. In the early nineties Boris Yamnitsky, working as an engineer at Media 100, took part in the first release of Media 100 version 1.0. In 1995, Boris Yamnitsky started Artel Software, creator of the popular Boris FX plug-in products. Boris FX became the first plug-in package for systems such as Adobe Premiere and Media 100 launching the vibrant NLE plug-in market. In 1998, in cooperation with Avid, Boris FX introduced the first-ever AVX plug-in product for Avid systems. The company has grown to serve thousands of video editors worldwide. All company products benefit from tight integration with video editing and compositing market leaders such as Adobe, Avid, Apple and Sony. In 2005, BorisFX merged with Media 100 and moved its headquarters to Marlboro, MA.

- **FRIDAY 3:00 PM – Strategies for GPU Acceleration of Common Visual Effects for Broadcast and Post-Production**

Yamnitsky, David
Engineer, Boris FX

David Yamnitsky has worked for Boris FX for the past three years concentrating in GPU-Accelerated Image Processing and Video Plug-in design. He designed and developed Boris Continuum Shaders, the set of GPU-based filters for Apple's FxPlug plug-in architecture for Final Cut Pro and Motion. Among the hallmarks of the set is BCS Chroma Key Studio, the GPU-based real-time Chroma Keyer whose implementation has formed the basis of a patent application. The set also features BCS

Film Effect, the GPU- based implementation of the traditionally cumbersome Film Effect filter and an interactive preset thumbnail browser.

- **FRIDAY 3:00 PM – Strategies for GPU Acceleration of Common Visual Effects for Broadcast and Post-Production**

Zhang, Yao
PhD, University of California, Davis

Yao Zhang is a PhD student in the Department of Electrical and Computer Engineering at University of California, Davis. Zhang received his BS in electrical engineering from the Beijing Institute of Technology. His research interests are in the area of GPU computing, especially in parallel algorithms for numerical linear algebra, and the GPU architecture/software optimization.

- **FRIDAY 2:30 – Fast Tridiagonal Solvers on GPU**

Seizing the opportunity
Accelerating change in the digital computing era

Digitization is having a profound impact on technology companies, as the entire value chain behind media and entertainment is transformed. The emergence of new media is radically transforming current business models, reshaping ecosystems and driving enormous demand for new capabilities to support the digital media marketplace.

As companies race to adapt to new rules and capitalize on new opportunities, they are increasingly turning to Deloitte. Nobody else delivers the range of professional services that help our clients accelerate the evolution of their business models in this dynamically changing marketplace. Our broad solution portfolio and integrated delivery model brings the multiple disciplines of strategy, operations, human capital and technology consulting, risk, governance and audit, tax, and financial advisory services. Coupled with our deep industry insight and global scale, we help companies ranging from emerging to well-established tackle the challenges and opportunities of the digital computing age.

Learn more at www.deloitte.com/us/technology

Deloitte.

GPU TECHNOLOGY CONFERENCE

NVIDIA RESEARCH SUMMIT POSTER LISTING

NVIDIA FELLOWS

Poster #1 – Diamos, Gregory
Georgia Tech (United States)
Extending Register Renaming To Memory Renaming

Poster #2 – Lin, Yen-Tzu
Carnegie Mellon (United States)
Reducing Die-Scrap Cost through Utilization of GPUs

Poster #3 – Meng, Jiayuan
University of Virginia (United States)
Dynamic Warp Subdivision for Non-Speculative Runahead SIMT Gather

Poster #4 – Patney, Anjul
University of California, Davis (United States)
Interactive Subdivision of Smooth Surfaces on GPUs

Poster #5 – Pinto, Nicolas
MIT (United States)
A High-Throughput Screening Approach to Discovering Good Forms of Biologically-Inspired Visual Representation

Poster #6 – Sintorn, Erik
Chalmers University (Sweden)
Radix sort of line primitives in CUDA for real-time Self-shadowing and Transparency in Hair

Poster #7 – Ufimtsev, Ivan
Stanford University (United States)
Quantum Chemistry on GPU: Towards Real World Applications

Poster #8 – Vo, Huy
University of Utah (United States)
Towards A Parallel Streaming Framework for Multi-Processing-Unit Systems

AI, VISION, SPEECH

Poster #9 – Coates, Daniel
University of California, Berkeley (United States)
Probabilistic Visual Computation with Spiking Neurons, Population Codes, and CUDA

Poster #10 – De Souza, Alberto
Laboratorio de Computacao de Alto Desempenho (Brasil)
Implementation in C+CUDA of Multi-Label Text Categorizers

Poster #11 – De Souza, Alberto
Laboratorio de Computacao de Alto Desempenho (Brasil)
Implementation of a Biologically Inspired Stereoscopic Vision Model in C+CUDA

Poster #12 – Degenaar, Patrick
Imperial College (United Kingdom)
Mobile augmented vision platforms to aid the visually impaired

Poster #13 – Domanski, Luke
CSIRO (Australia)
3D Image Deconvolution on Heterogeneous Compute Stations

Poster #14 – Domanski, Luke
CSIRO (Australia)
Parallel van Herk/Gil-Werman image morphology on GPUs using CUDA

Poster #15 – Fassold, Hannes
Joanneum Research (Austria)
Fast GPU-based KLT feature point tracking using CUDA

Poster #16 – Gallup, David
University of North Carolina at Chapel Hill (United States)
Real-Time Local Stereo Using CUDA

Poster #17 – Gupta, Kshitij
University of California, Davis (United States)
Towards High Quality Speech Recognition on Low-End GPUs

Poster #18 – Houzet, Dominique
Grenoble INP / GIPSA-Lab (France)
Parallel Implementation of a Visual Saliency Model

Poster #19 – Ivanov, Paul
University of California, Berkeley (United States)
Estimating the Entropy of Natural Scenes from Nearest Neighbors using CUDA

Poster #20 – Kider, Joseph
University of Pennsylvania (United States)
High-Dimensional Planning on the GPU

Poster #21 – Muyan-Ozcelik, Pinar
University of California, Davis (United States)
Speed-limit Sign Recognition with GPU Computing

Poster #22 – O'Donovan, Adam
University of Maryland (United States)
The Audio Camera for Real Time Audio Imaging and Analysis

Poster #23 – Plyer, Aurelien
Onera (France)
FOLKI-GPU: A powerful and versatile CUDA code for real-time optical flow computation

Poster #24 – Sorci, Matteo
Ecole Polytechnique Federale de Lausanne (Switzerland)
Robust Facial Expression Recognition

Poster #25 – Srinivasan, Balaji Vasan
University of Maryland (United States)
Scaling kernel machine learning algorithm via the use of GPUs

MEDICAL

Poster #26 – Cheryauka, Arvi
GE Healthcare (United States)
Interactive GPU IP Chain Simulator for Mobile X-Ray Imaging

Poster #27 – Dowsey, Andrew
Imperial College (United Kingdom)
A Real-time Simulation, Guidance and Visualisation Platform for Intra-operative Minimally Invasive Robotic Surgery

Poster #28 – Gipp, Markus
University of Heidelberg (Germany)
Haralick Texture Features Computation Accelerated by GPUs in Biological Applications

Poster #29 – Hissoiny, Sami
Polytechnique Montreal (Canada)
Radiotherapy Dose Calculation Acceleration with CUDA

Poster #30 – Jeong, Won-Ki
Harvard University (United States)
Neural Processes Segmentation in Confocal Microscopy Datasets using the GPU

Poster #31 – Marcus, Guillermo
University of Heidelberg (Germany)
Using GPUs for Lossless-JPEG real-time image decompression of high-throughput microscope data

Poster #32 – Richter, Rainer
Cleveland Clinic (United State)
SNP Generic Data Analysis

Poster #33 – Righetti, Raffaella
Texas A&M (United States)
Implementation of Ultrasound Elastography Techniques on GPGPU

Poster #34 – Roberts, Mike
PARC (United States)
CUDA Accelerated Sparse Field Level Set Segmentation on Large Medical Data Sets

Poster #35 – Schreier, Ludwig
University of Central Lancashire (United Kingdom)
GPU implementation of Cell Dynamics simulation for block copolymer systems

MISC

Poster #36 – Clark, Joshua
University of Minnesota (United States)
Particle Dispersion Modeling and Visualization in Urban Environments Using CUDA

Poster #37 – Hollemeersch, Charles-Frederik
Ghent University (Belgium)
Accelerating Virtual Texturing Using CUDA

Poster #38 – Smithwick, Quinn
MIT (United States)
Update on Realtime shader rendering of holographic stereograms

Poster #39 – Zhao, Yong
Brown University (United States)
Real-Time High Definition Stereo on GPGPU using Progressive Multi-Resolution Adaptive Windows

SIMULATION

Poster #40 – AbuZayyad, Tareq
University of Wisconsin – Madison (United States)
Monte Carlo Simulation of Photon Propagation and Detection by the IceCube Neutrino Detector

Poster #41 – Caruso, Francesco
Politecnico di Milano (Italy)
First steps towards real garment simulation on massive parallelism hardware

Poster #42 – Cheung, Mark
Lockheed Martin (United States)
Topological Studies of Solar Magnetic Fields

Poster #43 – Chiu, Ting-Wai
National Taiwan University (Taiwan)
Simulation of Lattice QCD with a GPU Cluster

Poster #44 – Clark, Mike
Boston University (United States)
Mixed-precision GPU Krylov solver for lattice QCD

Poster #45 – Deng, Yangdong
Carnegie Mellon (United States)
Using GPUs for Internet Routing Processing

Poster #46 – Faircloth, Daniel
Georgia Tech (United States)
Implementation and Issues Concerning Higher-Order Moment Method Simulations on CUDA-Enabled GPUs

Poster #47 – Fu, Zhisong
University of Utah (United States)
Solving EikonalEquations on Triangulated Surface Mesh with CUDA

Poster #48 – Ge, Wei
Institute of Process Engineering, Chinese Academy of Sciences (China)
Multi-scale HPC system for multi-scale discrete simulation --- Development and Application of China´s first supercomputer with one petaflops peak performance in single precision

Poster #49 – Herrmann, Frank
University of Maryland (United States)
Integrating Post-Newtonian Equations on Graphics Processing Units

Poster #50 – Hobiger, Thomas
National Institute of Information and Communications Technology (Japan)
A real-time GPS software receiver running on the GPU

Poster #51 – Hobiger, Thomas
National Institute of Information and Communications Technology (Japan)
Computation of atmospheric delay corrections on GPUs

Poster #52 – Jacobsen, Dana
Association for Computing Machinery (United States)
Massively Parallel Incompressible Navier-Stokes Computations on the NCSA Lincoln Tesla Cluster

Poster #53 – Lee, Paul
University of Virginia (United States)
Design Space Exploration for Integrated CPU-GPU Chips

Poster #54 – Lee, R. Todd
Georgia Tech (United States)
Strategies for Accelerating a Finite-Difference Time-Domain Code Using CUDA

Poster #55 – Li, Shaojing
University of California, San Diego (United States)
Fast Evaluation of Helmholtz Potential on GPUs

Poster #56 – Lin, Jim Ching-Rong
Landmark Halliburton (United States)
Interactive 3D seismic attribute volume generation with parallel graphics hardware

Poster #57 – Lu, Jesse
Stanford University (United States)
Nanophotonic Simulation using CUDA

Poster #58 – Mudigere, Dheevatsa
GE Global Research (Germany)
Optimized CUDA Implementation of a Navier-Stokes based flow solver for the 2D Lid Driven Cavity

Poster #59 – Peters, Amanda
Harvard University (United States)
Model of Evolutionary Game Dynamics Accelerated Using GPU Hardware

Poster #60 – Tarbutton, Joshua
Clemson University (United States)
GPU Accelerated Path Planning for Multi-Axis Machine Tool Enhancement

Poster #61 – Tzeng, Stanley
A Paradigm for Divide and Conquer Algorithms on the GPU and its Application to the Quickhull Algorithm

Poster #62 – Wakatani, Akiyoshi
Konan University (Japan)
CUDA Implementation of a LBG method for VQ

Poster #63 – Wang, Shuang
University of Oklahoma (United States)
Noise Adaptive LDPC decoding using Particle filtering based on parallel GPU architecture

Poster #64 – Wayth, Randall
Curtin University of Technology (Australia)
A GPU-based Real-time Software Correlation System for the Murchison Widefield Array Prototype

Poster #65 – Zaspel, Peter
Universitat Bonn (Germany)
Towards a multi-GPU solver for the three-dimensional two-phase incompressible Navier-Stokes equations

TOOLS

Poster #66 – Battaglino, Casey
MIT (United States)
GPUOctave – Enabling GPU Computing for GNU Octave

Poster #67 – Brock, James
Northeastern University (United States)
Task and Data Parallel Framework for GPU Computing

Poster #68 – Davidson, Andrew
University of California, Davis (United States)
CUDA Data Parallel Prefix (CUDPP) Library: A Powerful Toolkit for GPU Computing

Poster #69 – Dittamo, Cristian
Universita di Pisa (Italy)
Filling the gap between GPUs and Virtual Machinges computational models

Poster #70 – Dominguez, Rodrigo
Northeastern University (United States)
Improving the Open64 Backend for GPUs

Poster #71 – Duraiswami, Ramani
University of Maryland (United States)
FLAGON: A library for using the GPU from Fortran 9X

Poster #72 – Edgar, Richard
Harvard University (United States)
SciGPU-GEMM: A Wrapper for Matrix-Matrix Multiplications on Low-End GPUs

Poster #73 – Feng, Wu
Virginia Tech (United States)
Accelerating Applications with Inter-Block GPU Communication via Fast Barrier Synchronization

Poster #74 – Gosink, Luke
University of California, Davis (United States)
Data Parallel Bin-based Indexing Strategies for Answering Queries on Multi-Core Architectures

Poster #75 – Lauer, Tobias
University of Freiburg (Germany)
Palo+GPU: A CUDA-powered In-memory OLAP Server

Poster #76 – Lever, Jake
Harvard University (United States)
PCA-Augmented Vector Quantization on the GPU

Poster #77 – Malony, Allen
University of Oregon (United States)
Performance Tool Integration in Programming Environments for GPU Acceleration: Experiences with TAU, TAUcuda, and HMPP

Poster #78 – Maruyama, Naoya
Tokyo Institute of Technology (Japan)
A High-Performance Fault-Tolerant Software Framework for Memory on Commodity GPUs

Poster #79 – Nanjundappa, Mahesh
Virginia Tech (United States)
SCGPSim: A fast SystemC simulator on GPU´s

Poster #80 – Padalikar, Sudnya
NFinTes (United States)
GPU-RPC: Exploiting The Latency Tolerance of CUDA Applications

Poster #81 – Rehman, Tauseef
Accelereyes (United States)
Multi-GPU Support for MATLAB using Jacket

Poster #82 – Alcantara, Dan
University of California, Davis (United States)
Real Time Parallel Hashing on the GPU

Poster #83 – Stuart, Jeff
University of California, Davis (United States)
DCGN – Message Passing on GPUs

Poster #84 – Taufer, Michela
University of Delaware (United States)
Improving Numerical Reproducibility and Stability in Large-Scale Numerical Simulations on GPUs

Poster #85 – Wu, Ren
HP Labs (United States)
Very Large Scale Clustering on GPUs via Streaming

Poster #86 – Zhang, Yao
University of California, Davis (United States)
Fast Tridiagonal Solvers on GPU

Poster #87 – Zhao, Kaiyong
Hong Kong Baptist University (Hong Kong, China)
Implementation of Multiple-Precision Modular Multiplication on GPU

Poster #88 – Zou, Yi
University of Califo rnia, Los Angeles (United States)
Parallel Multi-level Analytical Global Placement on Graphics Processing Units

WEBCOR

Building Better

We are **hands-on builders** with experience on hundreds of challenging projects. This is the exceptional value that we bring to every project and it is what distinguishes us in our field. Our experience allows us to actively engage in the preconstruction of a project, providing design analysis and recommendations that make the building systems more efficient, less costly and of higher value.

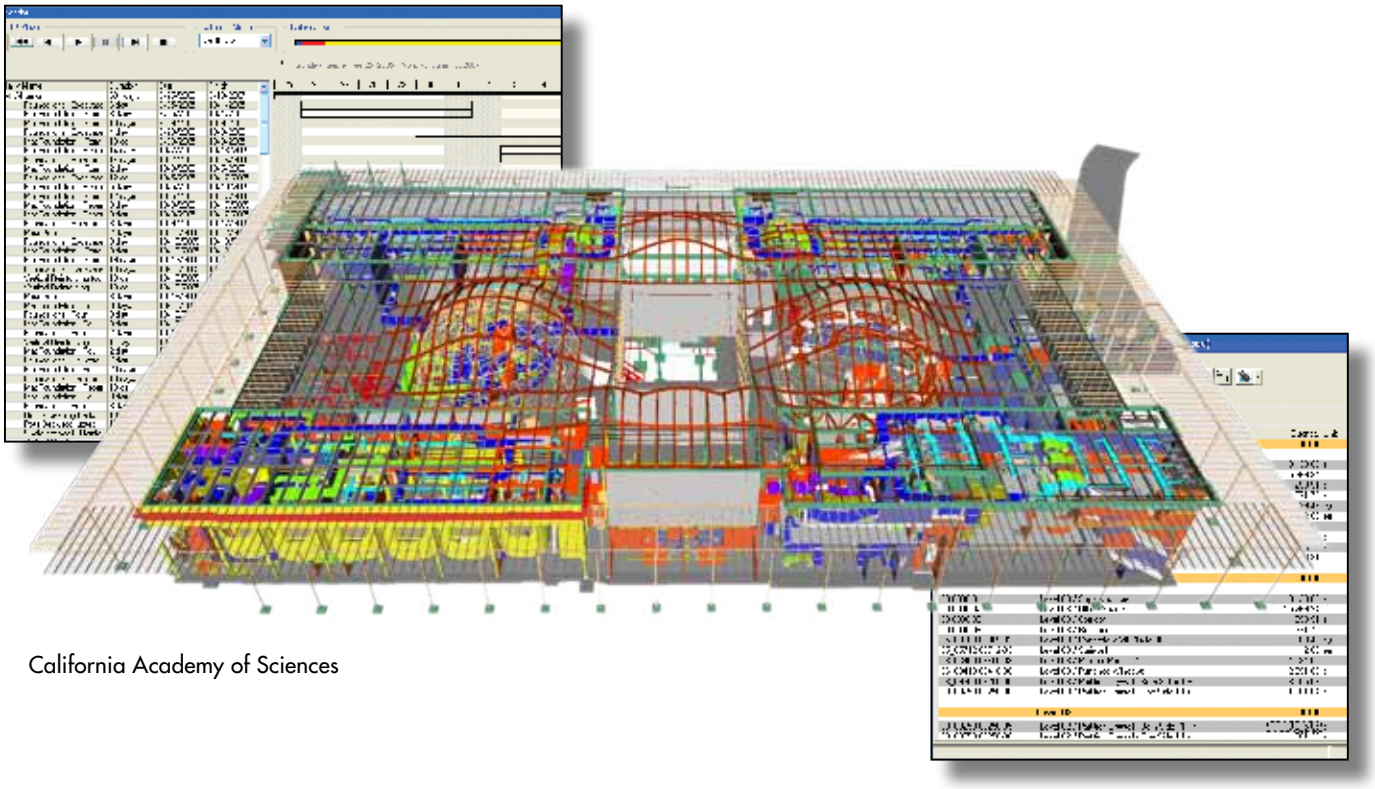

California Academy of Sciences

Smarter Tools

Webcor is the industry leader in developing and implementing new technology. By working closely with world-class institutions such as Stanford's Center for Integrated Facilities Engineering, we help continue the advancement and adoption of **Building Information Modeling**. Virtual building tools allow us to test new methods, create more realistic schedules, produce more complete estimates and identify conflicts before they appear in the field. The result – less time, lower costs, better buildings.

Let us show you the new era of construction!

SAN MATEO • SAN FRANCISCO • HAYWARD • LOS ANGELES • SAN DIEGO www.webcor.com

WEBCOR BUILDERS

CL #899242 A, B, C-8

GPU TECHNOLOGY CONFERENCE

SPONSORS AND EXHIBITORS

PLATINUM SPONSORS

Cooley Godward Kronish, LLP

Cooley Godward Kronish is a national law firm for the converging worlds of high technology, high finance and high-stakes litigation. We are counselors, strategists and advocates for the foremost private and public companies and investors in all major technology fields. Our Emerging Companies practice has a long tradition of representing emerging and high-growth companies worldwide. The GPU space is an exciting growth area in the technology arena, and Cooley Godward Kronish has been at the forefront, advising both established and start-up companies on the issues facing businesses in this industry. Our attorneys' extensive experience in intellectual property protection and business counseling along with the Firm's deep roots in the technology sector give us a unique perspective on the issues facing our clients. Cooley's team consists of experienced counselors and litigators that are equally skilled at representing and advising clients on the protection and commercialization of their intellectual property in a wide range of areas, including copyright, trademark, patent, technology licensing, privacy, electronic security and electronic commerce. We are dedicated to offering comprehensive and creative legal support, utilizing the full resources of the Firm.

Samsung

For over 30 years, samsung has been dedicated to making a better world through diverse businesses that today span advanced technology, semiconductors, skyscraper and plant construction, petrochemicals, fashion, medicine, finance, hotels, and more. Our flagship company, Samsung Electronics, leads the global market in high-tech electronics manufacturing and digital media. Through innovative, reliable products and services; talented people; a responsible approach to business and global citizenship; and collaboration with our partners and customers, Samsung is taking the world in imaginative new directions.

ASUS

ASUS has ranked among the top 100 global IT companies on BusinessWeek's "InfoTech 100" for 11 consecutive years. Our product portfolio includes notebooks, Eee products, motherboards, graphics cards, and more. We are a top 5 notebook brand, and most of our notebooks incorporate dedicated graphics. We are the world leader in motherboards and graphics cards.

PNY

PNY Electronics, Inc. originated out of Brooklyn, N.Y. in 1985 as a company that bought and sold memory chips. To emphasize its expansion into manufacturing new forms of memory and complementary products, the company changed its name in 1997 to PNY Technologies®, Inc. The company continues to evolve by delivering a full spectrum of high-quality products to elevate the overall computing experience

Dell

Dell's mission is to deliver innovative and cost-effective solutions that meet today's real-life customer challenges and work seamlessly in existing environments and with other products. We gather requirements directly through customer interactions, organized events, social media venues, and customer panels, including partnerships with a wide variety of key industry software, hardware and component suppliers to give us a uniquely broad perspective on the computing landscape. Many innovations begin in-house, led by a global team of top engineers, product designers and technical experts. Others begin as a team effort with Dell's strategic partners like NVIDIA. We maintain strong internal development capabilities and we partner, rather than compete, with top industry technology suppliers and original development manufacturers to bring you products you can trust and value.

Christie Digital

Christie is celebrating its 80th anniversary as a leader in visual solutions for world-class organizations, offering diverse applications for business, entertainment, and industry. With expertise in film projection since 1929 and professional projection systems since 1979, Christie has established a global reputation as the world's single source manufacturer of a variety of display technologies and solutions for cinema, large audience environments, control rooms, business presentations, training facilities, 3D and virtual reality, simulation, education, media and government. Christie has installed over 75,000 projection solutions worldwide.

GOLD SPONSORS

Adobe

Adobe revolutionizes how the world engages with ideas and information. For 25 years, the company's award-winning software and technologies have redefined business, entertainment, and personal communications by setting new standards for producing and delivering content that engages people virtually anywhere at anytime. From rich images in print, video, and film to dynamic digital content for a variety of media, the impact of Adobe solutions is evident across industries and felt by anyone who creates, views, and interacts with information. With a reputation for excellence and a portfolio of many of the most respected and recognizable software brands, Adobe is one of the world's largest and most diversified software companies.

Supermicro

The SuperServer 6016GT-TF-TM2 integrates two NVIDIA Tesla™ M1060 GPU Processors via two Gen2 PCI-Express x16 interfaces along with dual Intel® Xeon® processors. With two double-width GPUs in PCI-E 2.0 x16 lanes, this 1U server delivers truly non-blocking GPU performance, or up to 2 Teraflops of processing power, which makes it the highest performing 1U server on the planet.

Citi

Citi is today's pre-eminent financial services company, with some 200 million customer accounts in more than 100 countries. Our history dates back to the founding of Citibank in 1812, Bank Handlowy in 1870, Smith Barney in 1873, Banamex in 1884, and Salomon Brothers in 1910.

Silicon Valley Bank

SVB Financial Group provides diversified financial services to emerging, growth and established technology companies and the life science, venture capital and premium wine markets. Through its focus on specialized markets and extensive knowledge of the people and business issues driving them, SVB Financial Group provides a level of service and partnership that measurably impacts its clients' success.

Deloitte

Deloitte LLP's Technology, Media and Telecommunications (TMT) industry group helps TMT companies evaluate complex issues, develop fresh approaches to problems, and implement practical solutions. Clients include some of the world's leading and fastest growing semiconductor, software and computer manufacturing companies, wireless operators, satellite broadcasters, advertising agencies—as well as leaders in digital media, gaming, clean tech, publishing and telecommunications.

SILVER SPONSORS

Acceleware

Acceleware provides software solutions to harness the parallel processing capabilities of multi-core GPUs/CPUs for the Electronic Design and Oil & Gas industries. Our Acceleration platform seamlessly integrates with applications from industry leading vendors in the Electromagnetic Simulation, Seismic Data Processing, Reservoir Modeling industries and Linear Algebra solvers

Sutter Hill Ventures

Sutter Hill Ventures has financed technology-based start-ups and assisted entrepreneurs in building market-leading companies since 1964. Through their decades of experience, they have developed strong industry networks, considerable operating and venture capital experience, and an understanding of the challenges that early-stage and high-growth companies face. Sutter Hill Ventures offers expertise in developing business strategy and partnerships, building management teams, and financing companies in both public and private markets.

Webcor

From Oracle World Headquarters to the W San Francisco and W San Diego, we are proud to say that we've built many of California's landmark buildings during the last thirty years. We believe that three key factors contribute to the success of every project.

Dassault Systemes

As a world leader in 3D and Product Lifecycle Management (PLM) solutions, Dassault Systèmes brings value to over 100,000 customers in 80 countries. A pioneer in the 3D software market since 1981, Dassault Systèmes develops and markets PLM software and services that support industrial processes and provide a 3D vision of the entire product lifecycle from conception to maintenance. Our portfolio consists of CATIA for designing virtual products – SolidWorks for 3D mechanical design – DELMIA for virtual production – SIMULIA for virtual testing – ENOVIA for global collaborative lifecycle management, and 3DVIA for online 3D lifelike experiences.
Visit http://www.3ds.com.

Advantest

Advantest is the global leader in ATE solutions to the semiconductor industry. Serving the IDM, Fabless and OSAT markets, Advantest's test cells and systems—SoC, RF and memory testers, handlers, interfaces and software—offer the industry's highest level of reliability and accuracy and are integrated into the world's most advanced semiconductor fabrication lines.

HP

HP is a technology company that operates in more than 170 countries around the world. We explore how technology and services can help people and companies address their problems and challenges, and realize their possibilities, aspirations and dreams. We apply new thinking and ideas to create more simple, valuable and trusted experiences with technology, continuously improving the way our customers live and work.

EXHIBITORS

3 Dream Team

A team of experienced professionals in the field of 3D-graphics, programming and project management, ready to implement the project of any complexity. Our projects – the whole worlds that can independtly explore from every angle, from any angle. We are proud of our projects, because they are unique, effective and always exceed the expectations of our customers. Because our team is one of the few in Russia, can offer as a solution to the serious business problems a real adventure.

911 Comp

TDVision and 911 comp will show the TDVReady PC platform equipped with NVIDIA chipset and Quadro GPU for use in the encode and decode of full 1920x1080 per-eye high definition stereoscopic 3D content for home distribution channels including IPTV, Blu-ray and download.

AccelerEyes

AccelerEyes delivers simple software for powerful visual computing. AccelerEyes' product, Jacket, is the GPU engine for MATLAB, enabling MATLAB code to run faster. Jacket is used by customers across all major technical industries. Jacket's focus is to make GPU programming easy. Within minutes of downloading Jacket, the scientist, engineer, or financial analyst is able to accelerate their MATLAB code.

AMAX Information Technologies

AMAX provides leading supercomputing solutions. AMAX specializes in GPU parallel computing, high performance computing systems, and customized cluster solutions. Our cutting-edge solutions use best-of-breed components, are highly customizable and work to accelerate breakthroughs across a wide range of industries.

Animated Media Inc

Animated Media Inc. (AMI) provides solutions to help customers build intuitive, stunning visuals for multimedia applications. Our VGK product platform is a collection of integrated software tools that extend Adobe's Flash product functionality by allowing the technology to operate in a wide variety of systems, from low-power embedded devices to high-end 3D computers. For more information, visit us at www.animatedmedia.ca.

Appro International

Appro is a leading developer of supercomputing solutions for the high performance computing markets. Appro accelerates technical applications and business results through balanced architecture, open standards and engineering expertise. Headquarters is in Milpitas, CA, sales and services in Houston, TX and global sales and service presence with strategic international partners.

ASA Computers

Since 1989, ASA Computers Inc has been assisting individuals, companies, and organizations with their computing and networking requirements. Through various iterations of hardware and software flavors, we have been partnering with our clients in keeping their computing environment current in a cost-effective manner. We specialize in custom server solutions.

Bell Technologies

Bell Technologies DBA Bell Computer is a Digital Content Creation tools company, providing a full line of affordable FX products, from render box rental to graphical workstations. Bell Computer is located in Los Angeles, center of Digital Content Creation industry. Company provides FX solution, with products such as Render Rack rental and high-end Intel workstations.

Boxx Technologies

BOXX Technologies offers high-performance solutions that empower digital artists to create at the speed of thought. Since 1998, our unrivaled expertise and innovative integration of cutting-edge technology has been represented by a line of record-setting solutions—desktop workstations, mobile workstations, and render farms that serve a host of industries, including media and entertainment, architecture, research, defense, and more.

C3 Technologies

C3 Technologies provides the most detailed and realistic 3D city models available on the market today. Using a unique algorithm from the Swedish defense industry C3 rapidly processes huge city areas automatically from aerial photos. The textured 3D models contain enormous amounts of 3D mesh data instantly streamed to online map applications.

Colfax International

Colfax International is a leading provider of innovative and expertly engineered workstations, servers, clusters, storage, and personal supercomputing solutions. Colfax International is uniquely positioned to offer the broadest spectrum of high performance computing solutions, all of them completely customizable to meet your needs. Founded in 1987, Colfax International is based in Sunnyvale, California and is privately held.

Confetti Special Effects

Confetti offers leading-edge special effects middleware and services that lets video game developers shorten their game development cycle time and cost. Our software fully exploits the capabilities of the GPU for a 1,000x increase in the number of dynamic particles on-screen compared to traditional products, by eliminating inefficient off-chip data transfer and memory buffering between CPU and GPU. The performance increase is especially large for next generation GPUs. At the same time, our unique architecture enables very efficient rendering on portable platforms, such as iPhone and SnapDragon. Confetti is located in Carlsbad, California.

Creative Consultants

Creative Consultants LLC is a woman owned, New Mexico based, small business producing efficient performance computing for scientists, engineers, and commercial business. Our mission is to design, build, install and support High Performance computing systems advancing leading edge, sensible technologies while safe-guarding our environment. We specialize in hybrid and stereoscopic custom computing solutions. Our company has successfully serviced America's National Laboratories for more than two decades.

Cubix Corporation

Cubix Corporation is a blade PC manufacturer that has developed an innovative remote graphics technology that delivers pixel-for-pixel perfect images over long distances with no video compression. This technology is used deliver highly secure computing with full graphics performance capabilities on Quadro FX Series adapters (FX 5800 on down including multiple adapters are supported) over a distance of hundreds of meters. In business for more 35 years, Cubix is located in Carson City, Nevada.

Dimensional Imaging

Founded in 2002 and based in Glasgow, Scotland, Dimensional Imaging Ltd is a world-leading supplier of human body 3D & 4D surface image capture and analysis solutions.The company's highly innovative Di3D™ passive stereo photogrammetry technology delivers ultra-high resolution 3D surface images from one or more stereo pairs of instantaneously captured digital stills images. Dimensional Imaging supplied systems are already in use in 13 countries over four continents worldwide for a range of applications such as oral and maxillofacial surgery, orthodontics, burns treatment, facial recognition and entertainment.

EXHIBITORS

Dux Soft

Dux Soft is uniquely positioned as the "SAP" for digital content creation and performance computing market. The company focuses on core processes in the industry and provides scalable and flexible products that optimize digital workflow, it infrastructure and data storage. Dux combines its offering with deep domain expertise, last mile handholding, affordable managed services and easy transition to cloud based services to provide clients unprecedented control over costs and performance.

EM Photonics

EM Photonics is a recognized leader in implementing computationally intense algorithms on commodity hardware platforms. We develop custom solutions using GPUs, FPGAs, and embedded systems, for clients seeking to optimize their scientific computing, image processing, and numerical analysis applications. We recently released CULAtools, our GPU-accelerated implementation of the LAPACK linear algebra library

ENODO

As our specialty is to provide 3D simulation counsel, and adapt cutting-edge video-games technology to heavy industrial projects, our focus is to deliver believable, immersive and interactive virtual environment. Whether it is industrial simulation, urban planning, Serious Games ,transportation, infrastructure design, ENODO's team will tirelessly imagine ways to stretch 3D graphics and artificial intelligence to new fields , in order to reach your project's objectives.

Exxact

Founded in 1992, Exxact Corporation is a value-added distributor and solution provider bringing over 20 years of experience and passion to the workstation and server markets. As a leading provider of computer hardware and peripherals, Exxact provides services beyond warehousing and transport to build strong and sustainable customer relationships. We pride ourselves on our extensive product knowledge and unparalleled customer support, allowing our customers the unique opportunity to increase and sustain business through personalized service, life cycle management, product allocation, and EOL product sourcing,

GE Fanuc Intelligent Platforms

GE Fanuc Intelligent Platforms, a leading manufacturer of rugged electronics, offers the latest NVIDIA GPUs with both Graphics and CUDA capability on specialized COTS boards for critical defense applications.

Graph Stream

GraphStream Incorporated (www.graphstream.com) is a leading provider of scalable GPU-accelerated computing systems that are based on cost-efficient Commercial Off-The-Shelf components. Beginning in 2003 with the world's first commercially integrated GPU-computing cluster, including InfiniBand interconnect, GraphStream has deployed a number of the world's most powerful GPU-computing and visualization systems, each with hundreds of NVIDIA GPUs working in parallel.

Hue AS

We design, develop, distribute and support high performance visualization software. Our technology is coined HueSpace, and is the only commercially available high performance volume rendering solution which runs on PC clusters. Our products are HueViewer, HueSpace Server and HueSpace SDK. To learn more about our technology, see our technology section.

Hynix

Hynix Semiconductor, based in Icheon, Korea is a leading supplier of advanced semiconductor memory solutions and Image sensor products. We design, develop, manufacture and market a wide variety of DRAM and NAND Flash memories and CMOS Image Sensors (CIS). These memory components are essential in today's leading-edge computing, consumer and wireless communications applications. Image Sensors are used in a wide range of portable consumer electronics products such as handsets and handheld games.

James River Technical Inc/ VelocityHPC

JRTI, a leading provider of (HPC) solutions to the marketplace and Velocity Micro, the premier high-performance personal computer provider in North America, are pleased to introduce VelocityHPC, our latest initiative focused on NVIDIA Tesla GPU Accelerated Computing solutions.

Leebug

Leebug is a free professional event networking website that provides the tools to make every connection a strategic one and empowers interactions that will enhance your business. It is a place to network before, during and after events.

Luminova

Luminova is a world leader in providing extremely large scale, real-time and cost effective project management systems, utilising a complete 3D database structure. Luminova utilizes core NVIDIA-mental images 64-bit product technologies in uniquely providing large scale construction/engineering project management systems through web based and real-time interactive 3D visual information portals.

mental images

mental images®, founded in 1986, is the recognized international leader in providing rendering and 3D modeling technology to the entertainment, computer-aided design, scientific visualization, architecture, and other industries that require sophisticated images.

Mirics

Mirics is a fabless semiconductor company developing silicon and software solutions for portable broadcast reception. Combining a poly-band RF front-end with proprietary algorithmic IP making use of existing processor resources, Mirics is bringing mass market converged broadcast reception to portable and PC platforms.

Motion DSP

MotionDSP develops software that dramatically improves video from a wide range of sources -- from mobile phones to surveillance cameras. Our patented, multi-frame video enhancement technology is available in Ikena®, a video forensic solution, and in vReveal™, a Windows application for consumers.

NextComputing

Founded in 1999, NextComputing is a privately held company led by CEO and founder Robert Labadini. It has approximately 30 employees and is headquarted in Nashua, NH. NextComputing created the flextop, a personal supercomputer offering an open-standards, modular architecture in a package the size of a small briefcase.

EXHIBITORS

NVIDIA

NVIDIA Corporation is the worldwide leader in programmable graphics processor technologies. The Company creates innovative, industry-changing products for computing, consumer electronics, and mobile devices. The NVIDIA® brands include NVIDIA GeForce®, NVIDIA Tesla,™ NVIDIA GoForce,® NVIDIA Quadro,® and NVIDIA nForce.® These product families are transforming visually-rich and computationally-intensive applications such as video games, film production, broadcasting, industrial design, financial modeling, space exploration, and medical imaging. Additionally, NVIDIA invents and delivers industry-shaping technologies, including NVIDIA SLI™ technology, a revolutionary approach to scalability and increased performance; and NVIDIA PureVideo™ high-definition video technology.

OptiTex

OptiTex is the leading developer of 2D & 3D CAD solutions for virtually all sewn-products industries. OptiTex presents a complete content creation solution, to create and visualize customized garments, to simulate fitting and draping of garments on fully parametric virtual models and to create movie clips in an immediate and direct manner.

Penguin Computing

Penguin Computing has specialized 100% on Linux since 1998. With our unmatched dedication to Linux, we concentrate all our resources on developing the best practices that deliver the world's best Linux-based user experience to our customers. The powerful combination of our Linux software, systems, and services dramatically simplifies the deployment and management of large pools of servers.

Prometech

Promotech represents selected manufacturers based on quality, reliability, and integrity. We supply customized digital picture frames, MP3/MP4 players and accessories, USB flash drives, cameras, bluetooth products, solar powered products, and computer peripherals. In order to best serve your needs, customer service satisfaction is our primary goal. We strive to ensure that each order is accurately fulfilled in a timely manner.

PSSC Labs

PSSC Labs is everything you expect from your technology provider, and more. With 20 years in business, PSSC Labs possesses the knowledge, expertise and procedures to deliver high performance computing solutions to the world's most demanding organizations. PSSC Labs computing solutions empower next generation science.

Scalable Display

Founded in 2004 by two MIT PhD's, Scalable Display Technologies, Inc. is a leading provider of software used to create large projection based displays with resolution far beyond HD. Scalable's patented software is the catalyst for an emerging class of displays. Its software simplifies the creation of super-resolution, multi-projector displays of the highest quality and scalable size. EasyBlend opens the door to widespread use of multi-projector edge-blended displays for a wide range of applications including simulators based on off-the-shelf components, as well as supporting new forms of digital signage and data visualization tools.

Scenecaster

SceneCaster is a 3D social media application and online community where anyone can visualize their ideas, share them, and make them real. It allows you to create your own 3D spaces and share them with your Facebook and other social networking friends. Choose from millions of 3D objects and textures to personalize your scenes. Add your own photos and link to your favorite YouTube videos and online music. Invite your friends to view or modify your scenes. SceneCaster is easy, fun and rewarding.

EXHIBITORS

Sea C02

Seac02 provides for the software market augmented reality software aimed at the improvement of the quality and efficiency of the processes of engineering, marketing, sales and communication and at the reinforcement of awareness in consumer purchasing.

Soft Kinetic

Softkinetic is the leading provider of natural interfaces that transform the way people interact with the digital world. Softkinetic enables a fully immersive, transparent and intuitive user experience by providing a 3D real-time gesture recognition middleware to Interactive Digital Entertainment, Consumer Electronics, Health Care and Serious Games companies. Softkinetic, iisu, and Interface Is You are trademarks or registered trademarks of Softkinetic S.A. For more information on Softkinetic please visit http://www.softkinetic.net

UsefulProgress

The development in computer graphics allows huge progress in the knowledge of Life and Matter. In Medical science, CT scanners allow to investigate the whole body with transparency. A very important step in data analysis consist to convert signals (X, MR, US) in digital data that could be treated by computers. UsefulProgress develops new software strategies based on computer graphics for high-performance visualisation.

Vertice

Founded in 2002, Vertice is located in France. Vertice was elected 1st Innovating Company in 2003. Vertice develops and distributes the Nova Software range dedicated to 3D real time and accompanies its customers in the customization of applications. Vertice product range allows to easily get high performance real time renderings for client computers and web servers.

Visioglobe

Visiogloble edit an innovative software for 3D real time and cover the main cities in the world with: – interactive navigation – geolocalization -local search. The technology is targeted towards the mobile phone, automotive and Internet, implementing cartographic, geographical and real facades data.

Wild Systems

Wild Systems, a brand of HPC Project, delivers cost-effective application-in-a-box solutions for demanding users who require intense computational power for modelling and simulation applications. To support its partners and customers, HPC Project provides services to help them master parallelism and new chipset architectures.

MARKETING PARTNERS

LEEBUG · Dr.Dobb's · Venture Beat · VC Experts · EUREKA PRIVATE EQUITY

HEDGE CONNECTION · insideHPC · DE Desktop Engineering · FINalternatives · CleanTech BRIEF

Over 200 Global Companies Rely On Acceleware to Deliver Faster Results

Discover Acceleware Solutions At Booth #17

"Acceleware's solution enables us to complete our computations and generate images in hours instead of days."

Dr. Susan Hagness
University of Wisconsin-Madison

ELECTROMAGNETICS | OIL & GAS | LINEAR ALGEBRA

acceleware.com

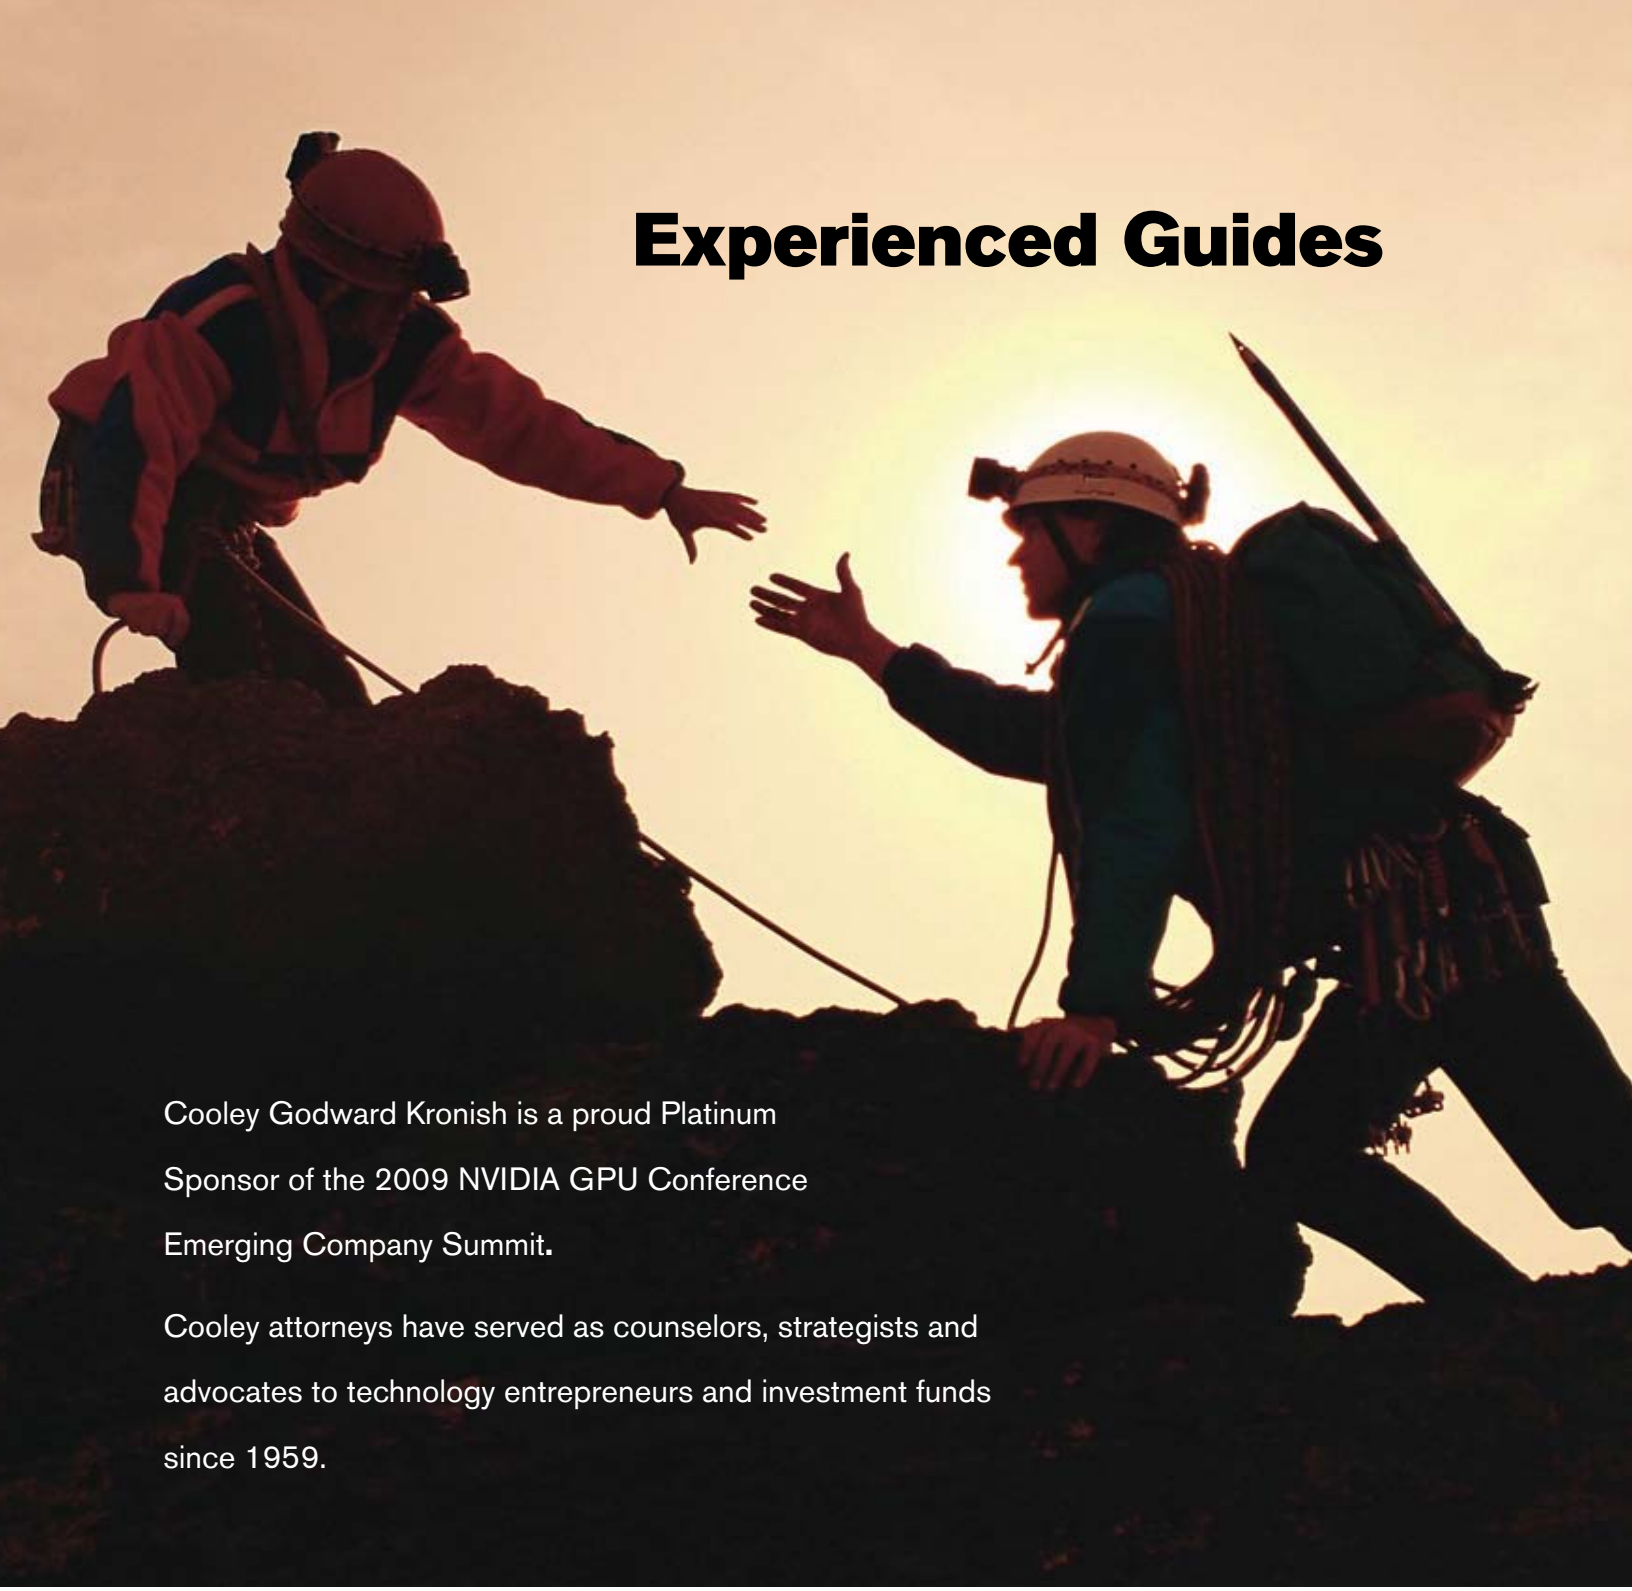

Experienced Guides

Cooley Godward Kronish is a proud Platinum Sponsor of the 2009 NVIDIA GPU Conference Emerging Company Summit.

Cooley attorneys have served as counselors, strategists and advocates to technology entrepreneurs and investment funds since 1959.

Cooley Godward Kronish, a national law firm for the converging worlds of high technology, high finance and high-stakes litigation. For more information, visit us at www.cooley.com

Cooley GODWARD KRONISH LLP

PALO ALTO | NEW YORK | SAN DIEGO | SAN FRANCISCO | RESTON, VA | BROOMFIELD, CO | WASHINGTON, DC | BOSTON | SEATTLE

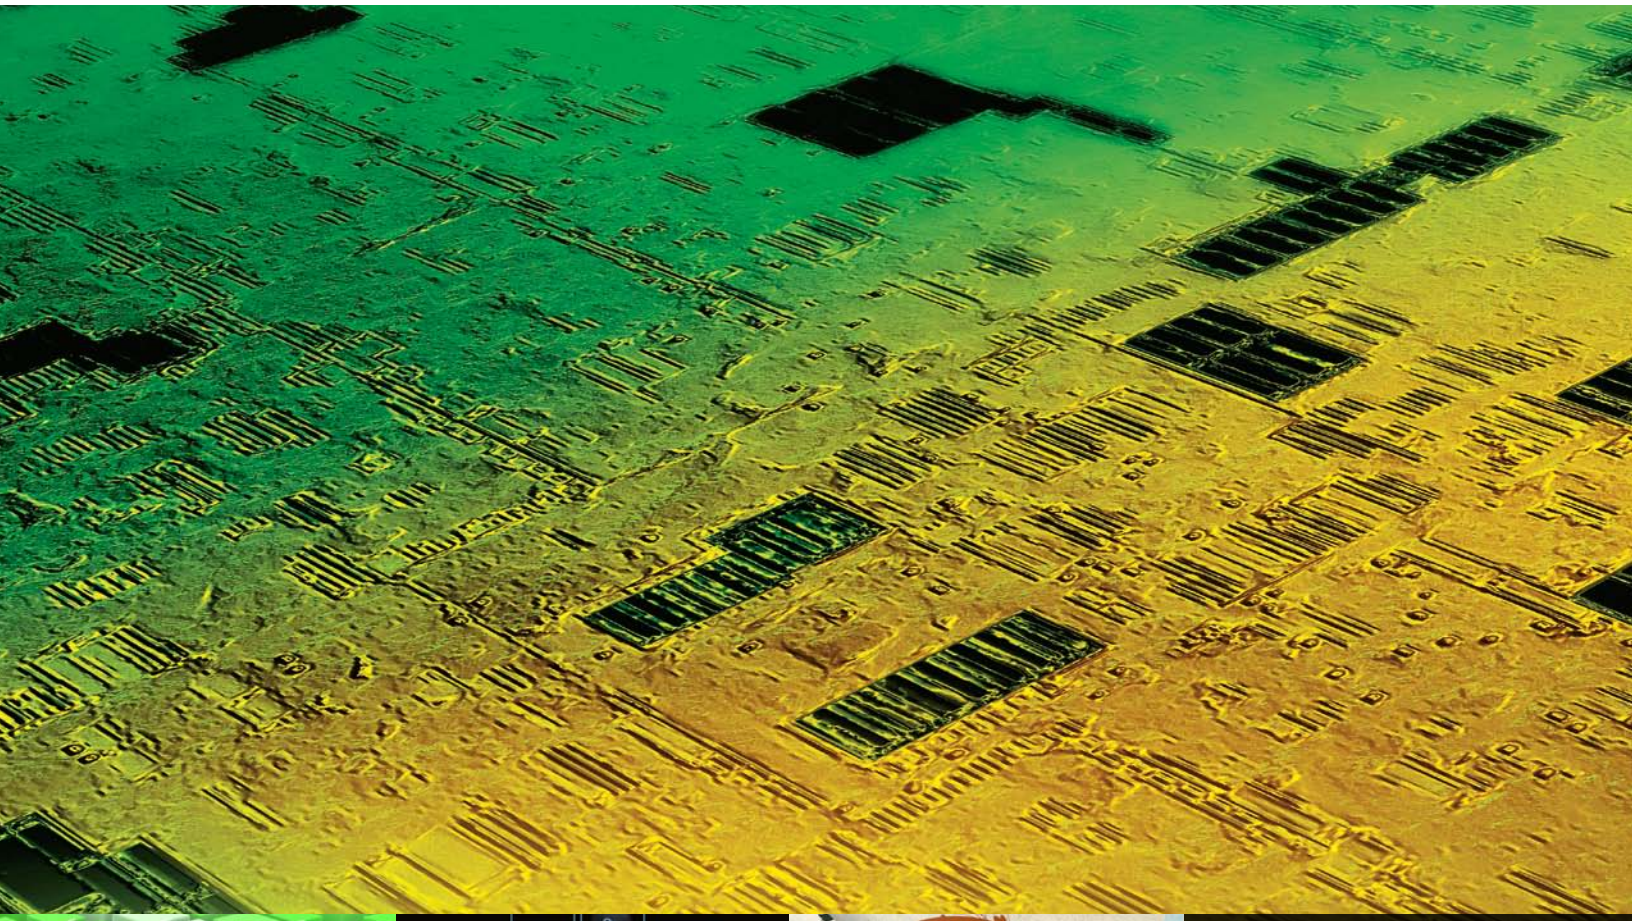

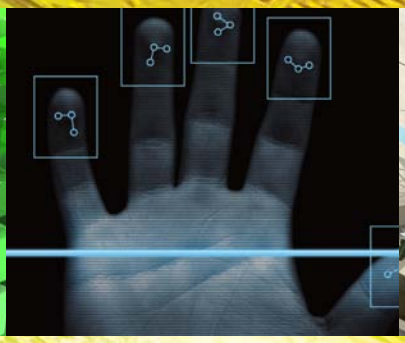

THE FAIRMONT HOTEL
SAN JOSE, CA
SEP 30 - OCT 2, 2009
www.nvidia.com/gtc

Images Courtesy of Heiko Klug / jesar.cgsociety.org and Chris Bobotis / mettle.com.
Entrants in CGSociety + NVIDIA Amazing Creations Challenge. CGSociety.org

Adobe
SUPERMICRO
citi
Silicon Valley Bank
A Member of SVB Financial Group
Deloitte.
acceleware
SUTTER HILL VENTURES
DASSAULT SYSTEMES
ADVANTEST
hp

FSC
Mixed Sources
Cert no. SCS-COC-001334
© 1996 FSC