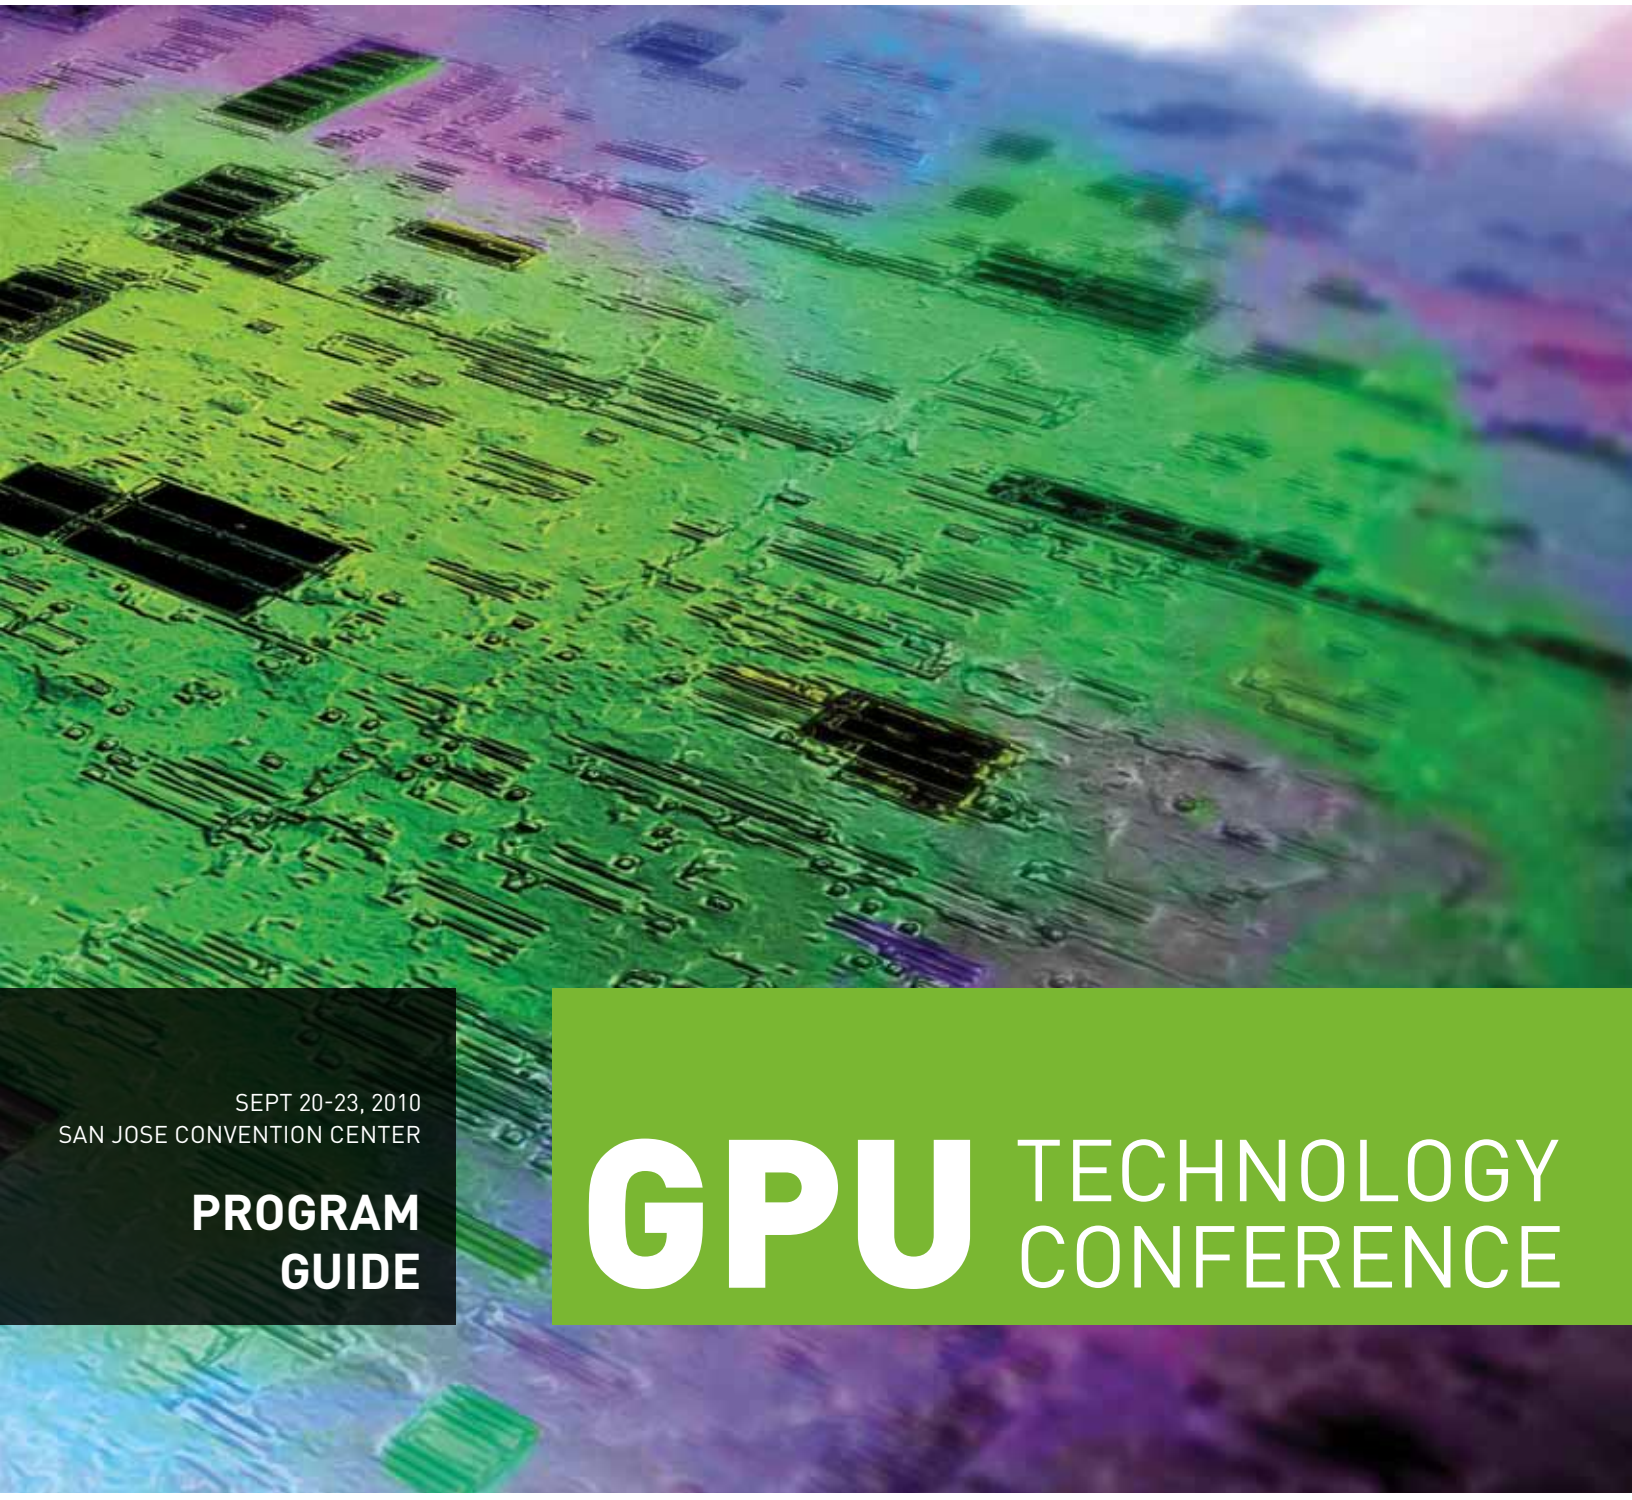

SEPT 20-23, 2010
SAN JOSE CONVENTION CENTER

PROGRAM GUIDE

GPU TECHNOLOGY CONFERENCE

PRESENTED BY

NVIDIA

SPONSORED BY

Adobe, HP, Microsoft, DELL, SUPERMICRO, PNY, Cooley GODWARD KRONISH LLP, SYNNEX CORPORATION, acer, tsmc, SAMSUNG, IBM

harnessing boundless capacity across client, cluster & cloud

Microsoft and Parallel Nsight – a powerful combination. Use Parallel Nsight to integrate into Microsoft Visual Studio, the world's most popular development environment. Run GPGPU accelerated applications on your desktop, on Windows Server or use Windows HPC Server 2008 for cluster applications.

Learn more at table 87 or the Parallel Nsight Lounge by Microsoft in the concourse area.

Microsoft®

Microsoft Visual Studio 2010 | Windows HPC Server 2008 R2 | Microsoft Excel 2010 | Windows Azure

THE WORLD, THESE DAYS, ISN'T FLAT. IT'S PARALLEL.

The parallel-processing power unleashed by the graphics processing unit, or GPU, is changing the face of computing. And, as it changes, our ability to address some of the world's most vexing challenges is improving.

Performing safer heart surgery. Making cars safer to drive. Drilling oil wells more accurately. Solutions to these and other complex computational problems have rapidly moved within reach, yielding results that have the potential to change lives and, ultimately, society as a whole.

Facilitating this work is the GPU, one of the most sophisticated processors ever manufactured. With up to three billion transistors in an area the size of a postage stamp, it can accelerate applications by several hundred times, shortening time to discovery from days to minutes.

At the same time, GPUs are far more power efficient than clusters designed exclusively with CPUs. And they are significantly less costly. In June, China's National Supercomputing Center, in Shenzhen, unveiled the world's second-most powerful super-computer, powered by Tesla GPUs, which was developed in a matter of months. Other enormously powerful GPU-powered supercomputers are on the way.

Another indicator of the triumph of parallel computing is the growing relevance of CUDA, the architecture developed by NVIDIA that enables GPUs to understand industry standard computing languages, as well as graphics APIs. Over the past year, NVIDIA has shipped more than 70 million CUDA-enabled GPUs; 10 textbooks about CUDA have been published in Chinese, English, Japanese and Russian; and two more universities each week, on average, are adopting CUDA into their curriculums.

And in the graphics space, NVIDIA's new Quadro professional graphics are ushering in a new era of computational visualization, bringing significant change to broadcast and film production, medical imaging and seismology, among other fields.

Perhaps the most immediately obvious sign of the importance of parallel computing, though, is the GPU Technology Conference itself. Last year's event far outstripped expectations in terms of attendance. And interest this year suggests that a revolution is at hand. Consider the following:

> The response to the call for talks at this year's conference rose more than fourfold from last year, while the number of sessions has more than doubled to some 300 hours.

> Representatives from more than 100 universities are registered.

> Attendees have arrived from some 50 countries.

For all these metrics, the ultimate success of the GPU Technology Conference, though, will be measured by the level of engagement it inspires, by the side conversations that it generates and the collaboration that it leads to.

Brace yourself for immersion in this brave new world!

The wait is over.

Experience breakthrough performance with the all-new Mercury Playback Engine in Adobe® Premiere® Pro CS5.

www.adobe.com/go/productionpremium

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

MONDAY, SEPTEMBER 20	8:00 AM to 6:00 PM
TUESDAY, SEPTEMBER 21	7:00 AM to 7:00 PM
WEDNESDAY, SEPTEMBER 22	8:00 AM to 6:00 PM
THURSDAY, SEPTEMBER 23	8:00 AM to 6:00 PM

EXHIBIT HALL AND MEAL HOURS

TUESDAY, SEPTEMBER 21	12:00 PM to 2:00 PM	Lunch / Exhibits Open
	6:00 PM to 8:00 PM	Reception / Exhibits Open
WEDNESDAY, SEPTEMBER 22	12:00 PM to 2:00 PM	Lunch / Exhibits Open
	6:00 PM to 8:00 PM	Reception / Exhibits Open
THURSDAY, SEPTEMBER 23	12:00 PM to 2:00 PM	Lunch / Exhibits Open

ENROLL INTO YOUR SESSIONS

Go to www.Nvidia.Com/GTC, click on "view schedule," and log in to start adding sessions into your personal schedule. Priority access into each session will be given to those who enroll. Enrolling into sessions also help us place the most popular sessions into the largest rooms.

WIRELESS INTERNET ACCESS

Free wireless internet access is available in most session rooms, the keynote hall, and also in the concourse outside of Ballroom A and the Exhibit Hall, under "GTC2010."

FIND OUT THE LATEST HAPPENINGS WITH THE CONFERENCE

Log on to www.nvidia.com/gtc to get the latest coverage on the event, along with any updates on room changes, access to session feedback survey, etc.

BUSINESS CENTER / SHIPPING

The Marriott Hotel and the Hilton Hotel both have business centers located on the first floor, near their front lobby. You can work out shipments with their respective front desks or bell desks. Alternatively, there is a FedEx Office Print & Ship Center at 93 E. San Carlos St, near 3rd St (3 blocks from the Convention Center, call 408-295-4336 for hours).

GO GREEN!

Take part in the shared goal of minimizing our collective impact on the environment. Please take only the conference materials you need and recycle your badges at the conclusion of the event. Also, we have provided GTC coffee mugs for those who opted in, so please use those to fill your hot and cold beverages to avoid contributing to more waste to the environment.

BAG AND COAT CHECK

Bag check is available at the bell desk of the Marriott and Hilton hotels, connected to the Convention Center. It is also available for a small fee on the ground floor of the San Jose Convention Center.

LOST AND FOUND

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

FIRST AID / EMERGENCY

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

Unparalleled PowerEdge flexibility with the M610x

The Dell™ PowerEdge™ M610x blade server allows you to creatively incorporate a vast array of expansion solutions, including the NVIDIA® Tesla™ GPGPU card.

Now, a single M610x, equipped with a NVIDIA Tesla GPGPU card and installed in a PowerEdge M1000e blade enclosure, can perform over 400 Gigaflops of double-precision computations for demanding, floating-point-intensive workloads. The x16 Gen2 PCIe slots in the PowerEdge M610x bring a new dimension of flexibility and performance to your data center.

Discover more at Dell.com/PowerEdge

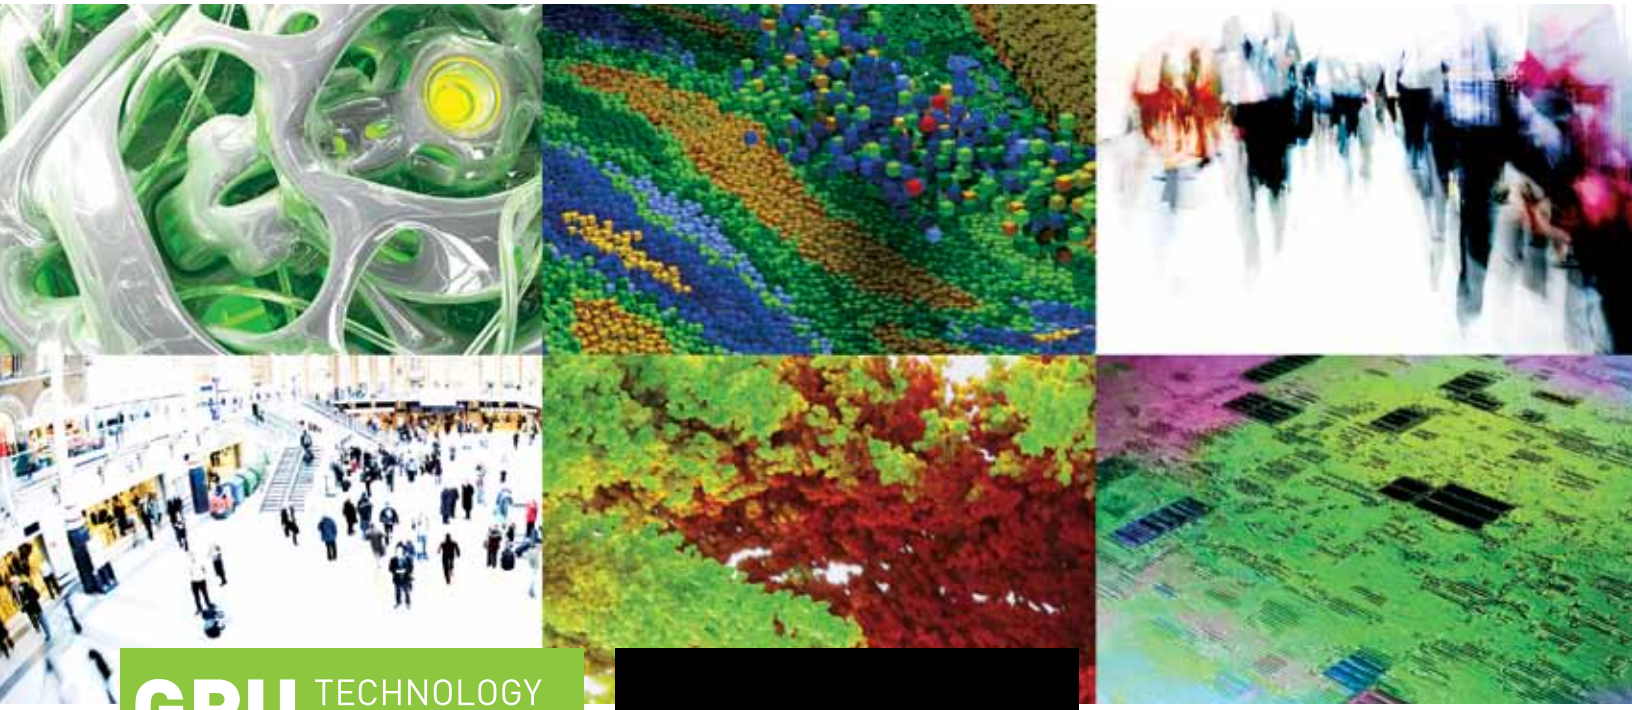

GPU TECHNOLOGY CONFERENCE

TABLE OF CONTENTS

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 Professional Graphics Advantage:

- 3 year warranty
- 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

Partner optimized applications and certified leading ISV's for:

- Large Scale Visualization
- HPC Environments
- CAD/CAE
- Digital Content Creation
- HD Broadcast and Film
- Video Editing and Effects

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

CONFERENCE HIGHLIGHTS DON'T MISS THESE EVENTS!

ALL WEEK LONG

Parallel Nsight Lounge, by Microsoft

While attending GTC, come learn from the experts at the Parallel Nsight™ Lounge by Microsoft, a casual environment for hands-on learning and instruction on Parallel Nsight, the industry's first development environment for GPU-accelerated applications. Experts from NVDIA and Microsoft will be available from 10am to 8pm each day to answer questions and provide instruction on Parallel Nsight, Visual Studio 2010, Windows HPC Server 2008 and CUDA C/C++ development.

CUDA Certification for GPU Computing Developers

The CUDA Certification Program is a response to the growing demand for qualified parallel programmers. To become CUDA certified, candidates must demonstrate an excellent working knowledge of the CUDA architecture and programming model, the ability to apply CUDA constructs to common algorithmic frameworks, and a strong understanding of optimization techniques related to CUDA C-based code.

> During GTC 2010, visit the "Ask the CUDA Experts" table to learn more!

> Additional info available at http://www.nvidia.com/certification

DigitalGuru: Where Smart People Get Smarter

DigitalGuru Technical Bookshop of Cupertino, California is pleased to participate in GTC 2010. Please visit our table during the conference for a wide and relevant selection of books on parallel programming, computer science, application tools and more. Books sold at GTC are available at 20% off list price. For more info, see www.digitalguru.com



SUPERMICRO®

Visit Us at Booth 71 and 73

The Evolution of IT

Optimized GPU Solutions

Scale up and Scale out GPUs using Supermicro's Twin® Technology

GPU+CPU
2 GPUs (2 TFLOPS)
The Fastest 1U Server in the World

2U Twin GPU
2 Hot-Swappable CPU+GPU
with Hot-Swappable Power Supplies

GPU+CPU SuperWorkstation
4 GPUs in 4U (4 TFLOPS)

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

ALL WEEK LONG

Birds of a Feather Gathering Places

Have a break in your schedule and want to network with some of the other brilliant minds here this week? Flock to one of the many "Birds of a Feather Gathering" tables out on the concourse. Each table is labeled with an interest group ranging from Computational Finance to Computer Vision. Table is open all day and evening.

TUESDAY

09:00 - 10:30	**Opening Keynote with Jen-Hsun Huang, NVIDIA CEO and Co-Founder** >Keynote Hall
12:00 - 14:00	**Exhibits Open / Networking Lunch** >Exhibit Hall
18:00 - 20:00	**Posters Showcase / Exhibits Open / Networking Reception** >Exhibit Hall and Concourse

WEDNESDAY

09:00 - 09:50	**Day 2 Keynote with Dr. Klaus Schulten, University of Illinois at Urbana-Champaign** >Keynote Hall
10:00 - 10:50	**Emerging Companies Summit Opening Address and Highlights** >Keynote Hall
12:00 - 14:00	**Exhibits Open / Networking Lunch** >Exhibit Hall
18:00 - 20:00	**Exhibits Open / Networking Reception** >Exhibit Hall

THURSDAY

09:00 - 09:50	**Emerging Companies Summit "Fireside Chat" featuring Quentin Hardy (Forbes Magazine) and Jen-Hsun Huang (NVIDIA)** >Keynote Hall
12:00 - 14:00	**Exhibits Open / Networking Lunch** >Exhibit Hall
17:00 - 18:30	**Closing Keynote with Sebastian Thrun, Professor / Distinguished Engineer, Stanford University and Google** >Keynote Hall

Closing Night Party for Charity

What happens when you throw together live music, raffle prizes, your GTC colleagues and a charity? A great party for a great cause. You can feel even better about letting loose as every dollar raised will be matched by the NVIDIA Foundation. For $10, you get a free drink ticket plus an entry into a raffle to win some stellar prizes. Learn more and buy your tickets at the NVIDIA Foundation table, the Gear Store booth, or buy at the door.
Time: **20:00 - ?**
Location: Voodoo Lounge, 14 S. Second Street, near Santa Clara Street. See your badge insert for a map to this location.



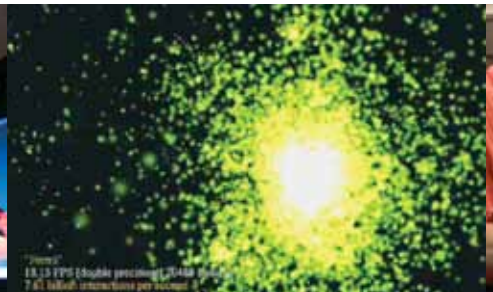


COMPUTE THE CURE

An initiative sponsored by the NVIDIA Foundation

There's a role for everyone in the search for a cure. Stop by the NVIDIA Foundation booth in the exhibit hall during GTC to learn about our efforts to help scientists speed the cure for cancer and to see how you can help.

NVIDIA
FOUNDATION

GPU TECHNOLOGY CONFERENCE

CLOSING NIGHT PARTY FOR CHARITY

What happens when you throw together live music, raffle prizes, your GTC colleagues and a charity? A great party for a great cause.

You can feel even better about letting loose as every dollar raised will be matched by the NVIDIA Foundation.

For $10, you get a free drink ticket plus an entry into a raffle to win some stellar prizes. Learn more and buy your tickets at the NVIDIA Foundation table on the concourse, the Gear Store booth, or buy at the door.

WHEN: **Thursday, September 23 at 8:00 PM**
WHERE: **Voodoo Lounge,** 14 S. Second Street, near Santa Clara Street

HopeLab

NVIDIA FOUNDATION

About the Cause

This year's recipient is Hope Lab, an organization that aims to enhance the physical health and psychological well-being of young people with cancer. Your donations will support the global distribution of Hope Lab's popular kids video game, Re-Mission, the first video game shown to induce positive behaviors that enhance the effectiveness of medical treatment. In Re-Mission, players pilot a nanobot named Roxxi as she travels through the bodies of fictional cancer patients destroying cancer cells, battling bacterial infections, and managing side effects associated with cancer and cancer treatment. Research shows that patients who played Re-Mission stuck to their prescribed treatments more consistently, a key component of successful cancer treatment, and showed increases in cancer knowledge and self-efficacy. Get your own copy at: www.re-mission.net.

About the NVIDIA Foundation

The NVIDIA Foundation is one of the only employee-led foundations in Silicon Valley. Each year, an employee board of directors receives feedback from our global employee population regarding the issues they care most about addressing, such as education and cancer research, and puts together an exciting set of programs that leverages NVIDIA's unique strengths as a company: having a transformative impact, engaging our employees, leveraging our ecosystem of customers, suppliers and vendors. The Foundation is developing a program called *Compute the Cure*, to drive forward a cure for cancer by supporting researchers working with gene sequencing technologies. Stop by the Foundation's booth at GTC, near the NVIDIA Store and Press Lounge to learn more.

RECOMMENDED FOR ACADEMICS

TUESDAY / SEPTEMBER 21

TIME	ID / SESSION TITLE
09:00 - 10:30	1001 – Opening Keynote with Jen-Hsun Huang
11:00 - 12:00	2223 – Academic Welcome Social and Poster Review
11:00 - 11:50	2112 – The Heisenberg Spin Glass Model on GPU: Myth versus Fact
14:00 - 14:50	2262 – CUDA Centers of Excellence Super-Session I
15:00 - 15:50	2263 – CUDA Centers of Excellence Super-Session II
16:00 - 16:50	2264 – CUDA Centers of Excellence Super-Session III
17:00 - 17:50	2265 – CUDA Centers of Excellence Super-Session IV
18:00 - 18:50	1005 – Research Poster Showcase / Exhibits Open / Networking Reception

WEDNESDAY / SEPTEMBER 22

TIME	ID / SESSION TITLE
09:00 - 09:50	1002 – Keynote with Dr. Klaus Schulten, University of Illinois at Urbana-Champaign
10:00 - 10:50	2280 – TSUBAME2.0 Experience
10:00 - 10:50	2082 – CU-LSP: GPU-based Spectral Analysis of Unevenly Sampled Data
10:00 - 10:20	2163 – Leveraging GPUs for Evolutionary Game Theory
10:00 - 10:50	2249 – New Programming Tools GPU Computing
10:00 - 10:50	2166 – The Triad of Extreme Computing-Fast Algorithms, Open Software and Heterogeneous Systems
10:00 - 10:50	2058 – A Practical Introduction to Computational Fluid Dynamics on GPUs
11:00 - 11:50	2078 – Shockingly fast and accurate CFD simulations
11:00 - 11:50	2286 – Towards Peta-Scale Green Computation- Applications of the GPU Supercomputers in the Chinese Academy of Sciences (CAS)
11:00 - 11:50	2177 – Simplifying Parallel Programming with Domain Specific Languages
14:00 - 14:50	2248 – Parallel Processing on GPUs at the University of Utah
14:00 - 14:50	2137 – CUDA for Real-Time Multigrid Finite Element Simulation of Soft Tissue Deformations
14:00 - 14:50	2000 – Gravitational N-body Simulations: How Massive Black Holes Interact with Stellar Systems
14:00 - 14:50	2164 – Analytical Performance Models to Improve the Efficiency of GPU Computing
14:00 - 14:50	2068 – Parallelizing FPGA Technology Mapping using GPUs
14:00 - 14:50	2204 – Bridging GPU Computing and Neuroscience to Build Large-Scale Face Recognition on Facebook.
15:00 - 15:50	2050 – Copperhead: Data-Parallel Python for the GPU
15:00 - 15:50	2029 – Computer Vision Algorithms for Automating HD Post-Production
15:00 - 15:50	2044 – GRASSY: Leveraging GPU Texture Units for Asteroseismic Data Analysis
15:00 - 15:50	2122 – Using GPUs for Real-Time Brain-Computer Interfaces
15:00 - 15:50	2281 – Domain-Specific Languages
16:00 - 16:50	2108 – Binary Black Holes Simulations using CUDA

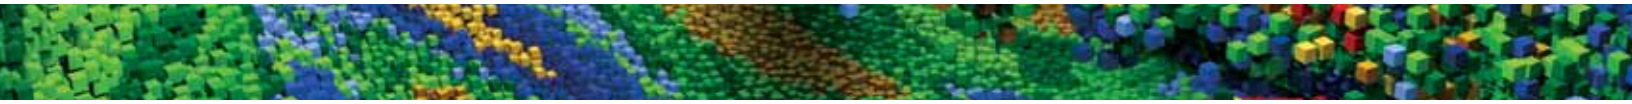

16:00 - 16:50	2118 – Large-scale Gas Turbine Simulations on GPU Clusters
16:00 - 16:50	2135 – Processing Petabytes per Second with the ATLAS experiment at the Large Hadron Collider at CERN
16:00 - 16:50	2226 – Reverse Time Migration with GMAC
17:00 - 17:50	2005 – Porting Large-Scale Legacy Fortran Codes
17:00 - 17:50	2242 – Swarming Bacteria and Diffusing Particles: High-Throughput Analysis of Microscopic 3D Motion
17:00 - 17:50	2167 – Designing a Geoscience Accelerator Library Accessible from High Level Languages

THURSDAY / SEPTEMBER 23

TIME	ID / SESSION TITLE
09:00 - 09:50	2030 – High-Throughput Cell Signaling Network Learning with GPUs
09:00 - 09:50	2236 – A Work-Efficient GPU Algorithm for Level Set Segmentation
10:00 - 10:50	2001 – Acceleration of the Freesurfer Suite for Neuroimaging Analysis
10:00 - 10:50	2269 – Bringing GPUs to Mainstream Molecular Dynamics Packages
10:00 - 10:50	2176 – Easy GPU Meta-programming: A Case Study in Biologically-Inspired Computer Vision
10:30 - 10:50	2292 – Implementation of High-Order Adaptive CFD Methods on GPUs
11:00 - 11:50	2007 – Folding@home: Petaflops on the Cheap Today; Exaflops Soon?
14:00 - 14:50	2054 – NAMD, CUDA, and Clusters: Taking GPU Molecular Dynamics Beyond the Desktop
14:00 - 14:50	2210 – GPU-Ocelot: An Open Source Debugging and Compilation Framework for CUDA
15:00 - 15:50	2062 – HOOMD-blue: Fast and Flexible Many-Particle Dynamics
17:00 - 18:30	1003 – Closing Keynote with Dr. Sebastian Thrun, Stanford University
20:00 - ???	1007 – Closing Night Party for Charity

EMERGING COMPANIES SUMMIT

GPU TECHNOLOGY CONFERENCE

Welcome back to NVIDIA's annual Emerging Companies summit. I'm delighted to report that 2010 marks the third consecutive and successful year for ECS, and our momentum continues to build!

The Emerging Companies Summit (an integral part of the GPU Technology Conference) is now the premier event for startups to share new applications, based on GPUs (graphics processing units) that are revolutionizing the computing industry. At the same time, these startups will have an opportunity to meet with hundreds of technologists, investors, analysts and executives who add additional "fuel" to the GPU computing ecosystem.

Familiar sectors such as media and entertainment have already been fundamentally altered by the GPU. Movie director James Cameron often says that Avatar could not have been created 10 years ago – it required powerful graphics processors to bring his vision to life. In addition to entertainment, new industries and applications are also being unleashed and significantly enhanced by GPU-based technologies.

At this year's ECS, I hope you will take full advantage of the opportunity to see and hear from 60 of the most promising companies in these fields. The companies, representing several countries from around the world, will showcase new technology in the fields of computer vision, robotics, video processing, cloud computing and mobile computing. What they all share in common is that they harness the massive power of GPUs to drive amazing performance for their applications.

In the spirit of innovation, this year we decided to introduce a new and exciting format for the startup presentations at ECS. While you will still be able to find booths in the exhibit hall for almost all 60 of the emerging companies, with the help of an advisory committee we have chosen a select group of 24 to participate in action-packed "CEO on Stage" sessions. The CEOs from these 24 emerging companies will have the special opportunity to both present and discuss their business strategies with panels comprised of some of the world's leading and most impressive venture capitalists, technology executives and industry analysts. I believe this will be one of the major highlights of this year's GPU Technology Conference, and I urge you to participate in as many of these sessions as possible.

In closing, I am extremely excited and honored to once again be part of the Emerging Companies Summit. The GPU computing ecosystem has now gathered a full head of steam, which will be clearly evident over the next few days. I would also like to add a special note of thanks to our sponsors which include Cooley Godward Kronish, Citi, Sutter Hill Ventures, Silicon Valley Bank, Deloitte, Mandel Communications, Churchill Club, and VentureBeat.

Thank you for attending, and welcome to the GPU computing revolution!

Cooley citi SUTTER HILL VENTURES Silicon Valley Bank

Deloitte. Mandel COMMUNICATIONS CHURCHILL CLUB VentureBeat

Jeff Herbst
Vice President of Business Development
NVIDIA

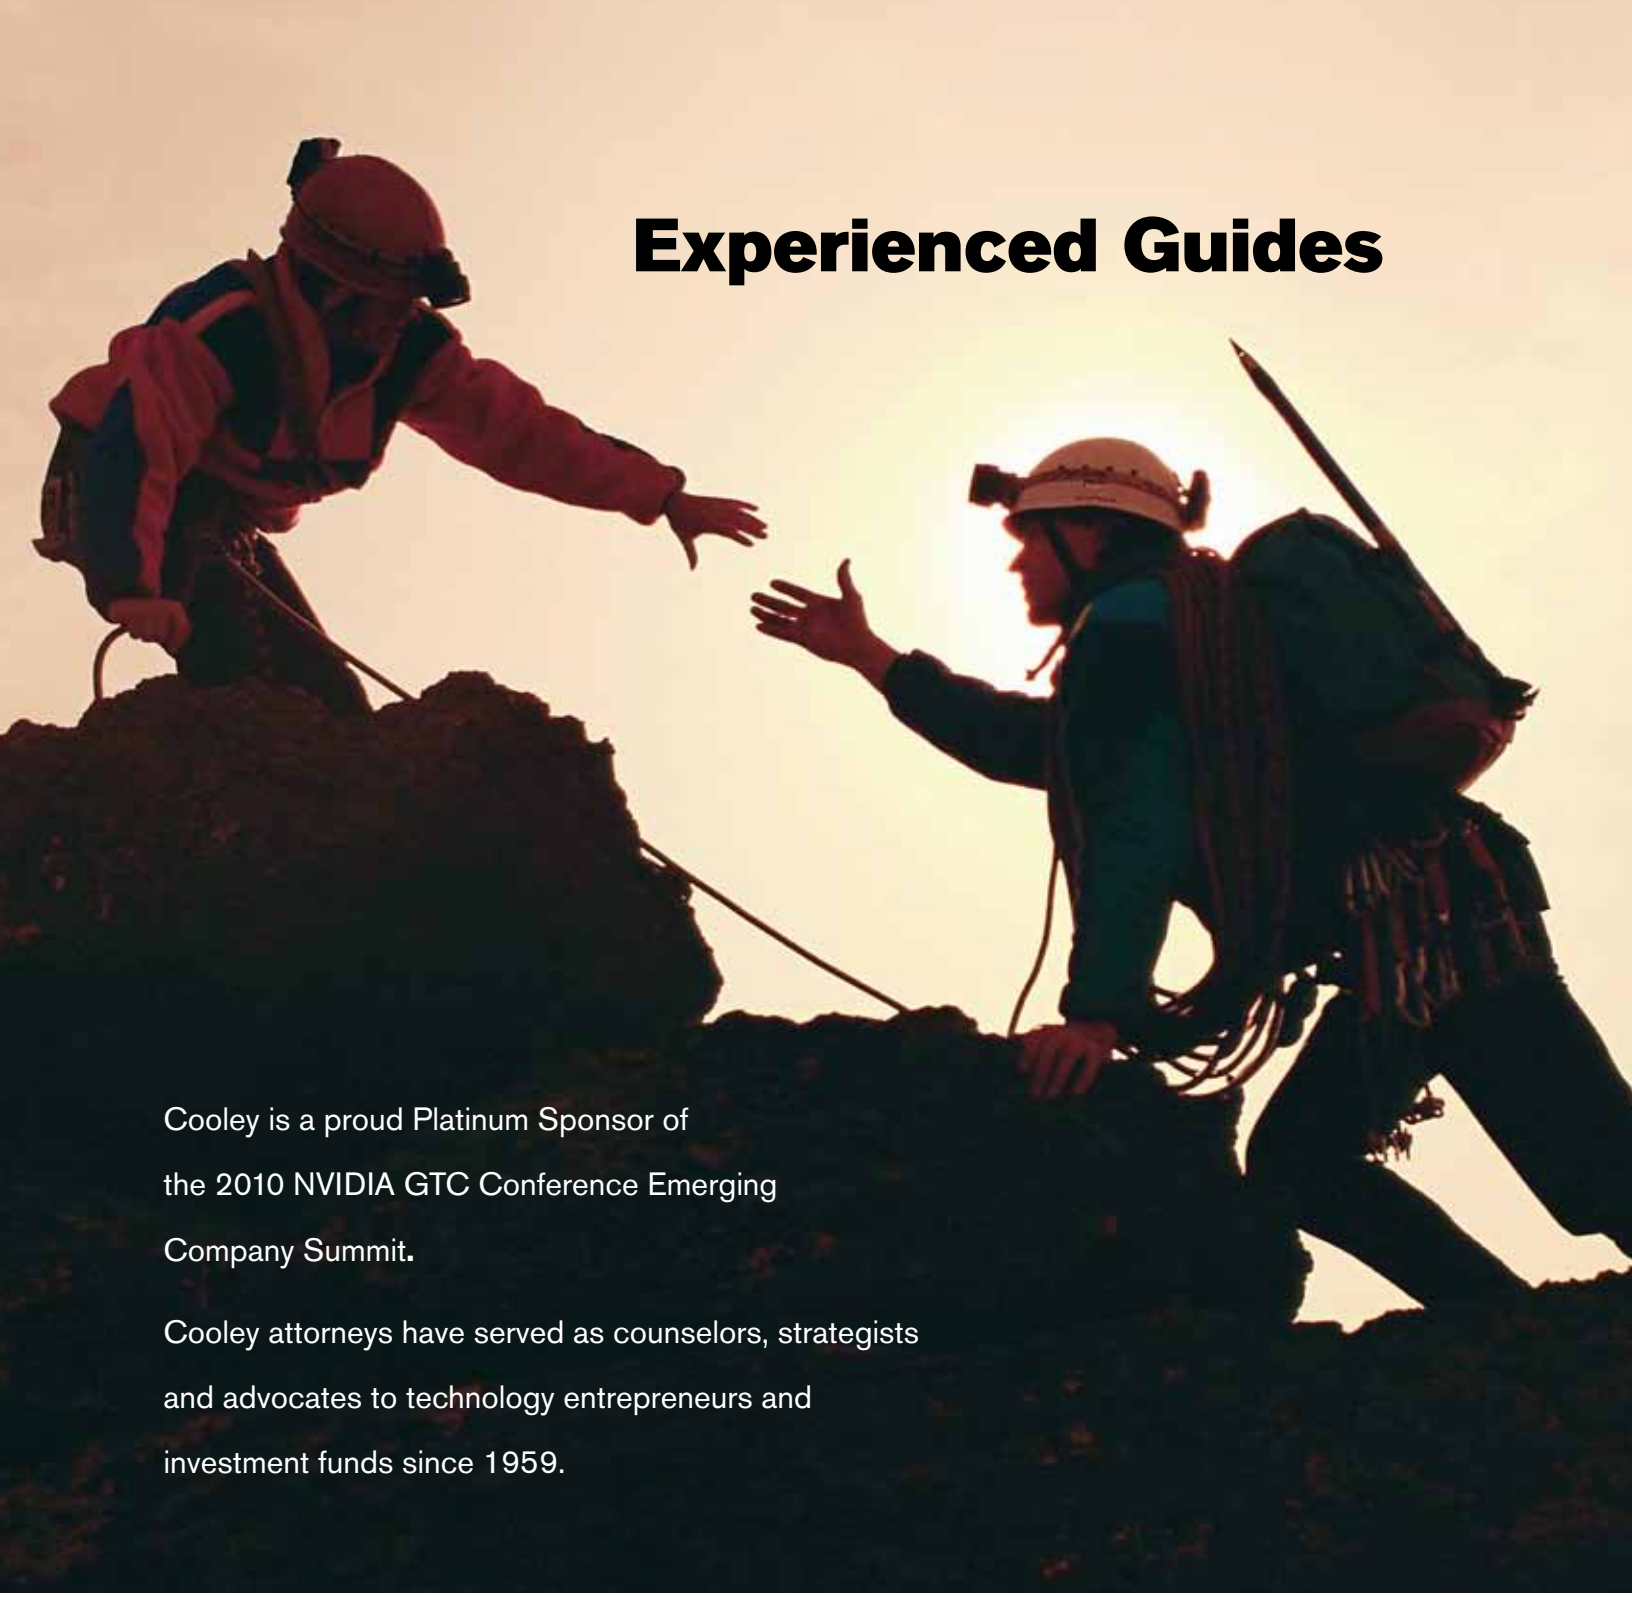

Experienced Guides

Cooley is a proud Platinum Sponsor of the 2010 NVIDIA GTC Conference Emerging Company Summit.

Cooley attorneys have served as counselors, strategists and advocates to technology entrepreneurs and investment funds since 1959.

Cooley, 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 LLP

PALO ALTO | NEW YORK | SAN DIEGO | SAN FRANCISCO | RESTON, VA | BROOMFIELD, CO | WASHINGTON, DC | BOSTON | SEATTLE

GPU TECHNOLOGY CONFERENCE

RECOMMENDED SESSIONS FEATURING EMERGING COMPANIES
EMERGING COMPANIES SUMMIT AGENDA

TUESDAY / SEPTEMBER 21

TIME	ID / SESSION TITLE
9:00 – 10:30	1001 – Opening Keynote with Jen-Hsun Huang
12:00 – 14:00	1004 – Exhibits Open / Networking Lunch
18:00 – 20:00	1005 – Exhibits Open / Networking Happy Hour/ Research Posters Showcase

WEDNESDAY / SEPTEMBER 22

TIME	ID / SESSION TITLE
9:00 – 9:50	1002 – Day 2 Keynote with Dr. Klaus Schulten, University of Illinois at Urbana-Champaign
10:00 – 10:50	4000 – Emerging Companies Summit Opening Address featuring Jeff Herbst (NVIDIA)
11:00 – 11:50	4001 – Emerging Companies Summit "CEO on Stage" featuring Sam Blackman (Elemental Technologies, Inc.), Sam Cox (Milabra), Chris Doran (Geomerics) and panelists Drew Lanza (Partner, Morgenthaler), Dan'l Lewin (Corporate VP and Strategic & Emerging Business Development, Microsoft), Jon Peddie (President, JPR), Jeff Herbst (Vice President of Business Development, NVIDIA)
12:00 – 14:00	1004 – Exhibits Open / Networking Lunch
14:00 – 14:50	4002 – Emerging Companies Summit "CEO on Stage" featuring Christopher Blewitt (miGenius), Sebastien Deguy (Allegorithmic), Philip Lunn (Bunkspeed) and panelists Drew Lanza (Partner, Morgenthaler), Dan'l Lewin (Corporate VP and Strategic & Emerging Business Development, Microsoft), Jon Peddie (President, JPR), Jeff Herbst (Vice President of Business Development, NVIDIA)
15:00 – 15:50	4003 – Emerging Companies Summit "GPUs for Computer Vision" moderated by Jon Peddie (Jon Peddie Research), featuring panelists Sam Cox (CEO, Milabra), Tom Dean (Research Scientist, Google) Janko Mrsic-Flogel (CTO, MirriAd), Joe Stam (Sr. Applications Engineer, NVIDIA), Yoram Yaacovi (CTO & General Manager, Technologies at Microsoft)
16:00 – 16:50	4004 - Emerging Companies Summit "CEO on Stage" featuring Michael Hummel (empulse GmbH), Natan Peterfreund (Playcast Media Systems), Austin Shoemaker (Cooliris) and panelists Nathan Brookwood (Research Fellow, Insight64), Charles Carmel (VP of Corporate Business Development, Cisco), Flip Gianos (General Partner, InterwestInterWest Partners), Jeff Herbst (Vice President of Business Development, NVIDIA)
17:00 – 17:50	4005 - Emerging Companies Summit "CEO on Stage" featuring Michel Tombroff (Softkinetic), Uri Tal (Rocketick), Kristian Raue (Jedox Business Intelligence) and panelists Nathan Brookwood (Research Fellow, Insight64), Charles Carmel (VP of Corporate Business Development, Cisco), Flip Gianos (General Partner, InterwestInterWest Partners), Jeff Herbst (Vice President of Business Development, NVIDIA)
18:00 – 20:00	1005 - Exhibits Open / Networking Happy Hour

THURSDAY / SEPTEMBER 23

TIME	ID / SESSION TITLE
9:00 – 9:50	4006 – Emerging Companies Summit "Fireside Chat" featuring Quentin Hardy (Forbes Magazine) and Jen-Hsun Huang (Co-founder & CEO, NVIDIA)
10:00 – 10:50	4007 – Emerging Companies Summit "CEO on Stage" featuring Andrew Jamison (Scalable Display Technologies), Jeroen Snepvangers (RTT), Michael Zeitlin (Aqumin) and panelists Rob Enderle (Analyst, Enderle Group), Jeff Herbst (Vice President of Business Development, NVIDIA), Savitha Srinivasan (Corporate Venture Partner, IBM), Norman Winarsky (VP of Ventures, Licensing and Strategic Programs, SRI)
11:00 – 11:50	4008 – Emerging Companies Summit "CEO on Stage" featuring David Peters (Universal Robotics), David Hayes (ICD) and panelists Rob Enderle (Analyst, Enderle Group), Jeff Herbst (Vice President of Business Development, NVIDIA), Savitha Srinivasan (Corporate Venture Partner, IBM), Norman Winarsky (VP of Ventures, Licensing and Strategic Programs, SRI)
12:00 – 14:00	1004 – Exhibits Open / Networking Lunch
14:00 – 14:50	4009 – Emerging Companies Summit "The `New Normal' For Building Emerging Companies Based On Disruptive Technologies" moderated by Jeff Herbst (NVIDIA), featuring panelists Gerald Brady (Silicon Valley Bank), Bill Frauenhofer (Managing Director, Citigroup Global Markets), Garrett Herbert (Partner, M&A Transaction Services, Deloitte & Touche LLP), Eric Jensen (Partner, Business Department Chair, Cooley LLP), Andrew T. Sheehan (Managing Director, Sutter Hill Ventures)
15:00 – 15:50	4010 - Emerging Companies Summit "CEO on Stage" featuring Yoram Burg (OptiTex), Sylvain Ordureau (Useful Progress), Torsten Reil (NaturalMotion) and panelists Tim Bajarin (Creative Strategies), Bill Tai (Charles River Ventures), Paul Weiskopf (Adobe)
16:00 – 16:50	4011 – Emerging Companies Summit "CEO on Stage" featuring Jeff Han (Perceptive Pixel), Lance Maurer (Cinnafilm, Inc.), Bruno Uzzan (Total Immersion) and panelists Tim Bajarin (President, Creative Strategies), Jeff Herbst (Vice President of Business Development, NVIDIA), Bill Tai (General Partner, CRV), Paul Weiskopf (Sr. VP of Corporate Development, Adobe)
17:00 – 18:30	1003 – Closing Ceremony / "Ones to Watch" Award Presentation and Closing Keynote with Dr. Sebastian Thrun, Stanford University
20:00 – ?	1007 – Closing Party for Charity

EMERGING COMPANIES IN THE EXHIBIT HALL

For full profile, please see "Sponsors and Exhibitors" section.

3DreamTeam
3DTV Solutions
AccelerEyes
Acceleware
Allegorithmic
Binatix, Inc.
Biodigital Systems
Bunkspeed
Cinnafilm
CodeSourcery
Cooliris
Cyberlink Corporation
Discretix Technologies
embodee
EM Photonics
Empulse GmbH
Filter Foundry
Geomerics
HPC Projects / Wild Systems
Israel Economic Mission (IEM)
Jedox Business Intelligence
Mersive Technologies
miGenius
Milabra
Milngleverse
NaturalMotion Ltd
OptiTex
Perceptive Pixel
PhaseSpace, Inc.
Phototour
Playcast Media Systems
Prometech Software, Inc.
RealityFrontier
Reservoir Labs
RTT
Scalable Display Technologies
ScaleForm Corporation
SEACO2
Softkinetic
Stonetrip
Tide Powered Ltd.
Trinity Racing Concepts, LLC
Universal Robotics
Useful Progress
VisiSonics Corporation

Allegorithmic

Allegorithmic is the company behind *Substance*, the first professional middleware for the authoring and on-the-fly rendering of smart textures. Substance allows content developers to produce textures twice as fast as usual, while Substance description files are typically 500-1000 times smaller than regular bitmaps: Creating and distributing textures for online and mobile 3D games has never been that efficient.

Speaker Sébastien Deguy, Founder & CEO

Speaker Session
4002 - Emerging Companies: CEO on Stage featuring Allegorithmic, Bunkspeed, and miGenius (Wednesday, Sept 22, 14:00)

CEO Sébastien Deguy

Investors Undisclosed

Capital Raised Undisclosed

Aqumin

Aqumin LLC invented AlphaVision™ to give financial professionals the ability to see relationships between data across entire markets at once. More than just another heat map or static table, AlphaVision™ converts financial data into interactive three-dimensional financial landscapes that enable real-time, multi-variate comparative analysis. Coupled with data services from premier providers such as Bloomberg, ThompsonReuters, ActivFinancial and others, Aqumin is developing a rich view library that simplifies the process of gathering, organizing, and presenting relevant information. At the click of a button, AlphaVision™ facilitates the seamless transition from current table-based views of global securities markets to an interactive, multi-dimensional workspace.

Speaker Michael Zeitlin, CEO

Speaker Session
4007 - Emerging Companies: CEO on Stage featuring Aqumin, RTT, and Scalable Display Technologies (Thursday, Sept 23, 10:00)

CEO Michael Zeitlin

Investors Private Individuals

Capital Raised $2 Million

Bunkspeed

Bunkspeed is a 3D rendering and animation software developer based in Carlsbad, California. Our philosophy has been to create easy to use use software for creative people with no prior 3D modeling or rendering experience, thus expanding the marketing beyond traditional bounderies established by complex rendering software. Founded in 2003, Bunkspeed software has become the standard in the industrial design community and is spreading rapidly to the engineering design and marketing communities. Recently Bunkspeed has introduced it's new generation of 3D rendering software based on mental images iray accelerated by the NVIDIA CUDA GPU's.

Speaker Philip Lunn, CEO

Speaker Session
Speaking Session 4002 - Emerging Companies: CEO on Stage featuring Allegorithmic, Bunkspeed, and miGenius (Wednesday, Sept 22, 14:00)

CEO Philip Lunn

Investors Undisclosed

Capital Raised Undisclosed

Cinnafilm

Cinnafilm™ was founded to address the absence of quality, software-based tools to visually optimize, convert and repurpose video images in the post production market. Recognizing the power of modern Graphics Processing Units (GPUs) and the accelerating migration to file-based workflows, Cinnafilm disregarded prevailing industry methods in 2005 and started from scratch to create what has become the world's fastest, most accurate GPU-based image processing engine, Pixel Strings™. Cinnafilm is a growing company that has successfully partnered with some of the strongest names in the motion picture and television industry: ARRI, Quantel, Rhozet, and NVIDIA. Our partners have recognized the power of Pixel Strings and the superlative image quality which can result when this power is properly harnessed. Cinnafilm is a privately held company, headquartered in Albuquerque, NM, amongst powerful resources such as the nation's defense laboratories and New Mexico's highly competitive film tax incentives.

Speaker Lance Maurer, CEO

Speaker Session
4011 - Emerging Companies: CEO on Stage featuring Cinnafilm, Perceptive Pixel and Total Immersion (Thursday, Sept 23, 16:00)

CEO Lance Maurer

Investors Undisclosed

Capital Raised Undisclosed

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®

Cooliris

Cooliris was founded with a simple mantra: "Think beyond the browser". The company creates products that make discovering and enjoying the Web more exciting, efficient,and personal. 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 over 12 million installs of the product, with thousands more downloads everyday

Speaker Austin Shoemaker, CTO

Speaker Session
4004 - Emerging Companies: CEO on Stage featuring Cooliris, empulse GmbH, and Playcast Media Systems (Wednesday, Sept 22, 16:00)

CEO Soujanya Bhumkar

Investors Kleiner Perkins Caufield & Byers, DAG Ventures, The Westly Group, and T-Venture

Capital Raised $20+ Million

Elemental Technologies

Elemental Technologies is the leading provider of massively parallel video processing solutions for broadcast and online video customers. Elemental's products use off-the-shelf, programmable graphics processing units (GPUs) for compute-intensive video processing tasks. The product line is ideal for digital media workflows that require video encoding for Internet and mobile distribution.

Speaker Sam Blackman, CEO & Co-Founder

Speaker Session
4001 - Emerging Companies: CEO on Stage featuring Elemental Technologies, Geomerics, and Milabra (Wednesday, Sept 22, 11:00)

CEO Sam Blackman

Investors Elemental's venture capital investors are Steamboat Ventures, headquartered in Burbank, California; General Catalyst, headquartered in Cambridge, Massachusetts; and Voyager Capital, headquartered in Seattle, Washington. Each venture capital investor has a board seat, alongside with other industry experts on the Elemental board who hail from Adobe and Pixelworks.

Capital Raised $14 Million

empulse GmbH

empulse was founded in 2007 by two ex-Accenture managers to help clients realize complex interactive web projects. Today, empulse employs 20 IT-professionals and realizes projects for leading international corporations.
To deliver outstanding search and analysis performance empulse developed ParStream - an innovative database for handling billions of structured data sets.

Speaker Michael Hummel, Managing Director

Speaker Session
4004 - Emerging Companies: CEO on Stage featuring Cooliris, empulse GmbH, and Playcast Media Systems (Wednesday, Sept 22, 16:00)

CEO Joerg Bienert

Investors Privately held by three founders

Capital Raised Undisclosed

Geomerics

Geomerics is an innovation-led company built on advanced in-house IP, a world-class research team, and strong management experience. Geomerics' first product is Enlighten. Enlighten redefines the way lighting is handled in computer games. Instead of pre-baking the effects of global illumination into the scene they are computed at run time, allowing for fully dynamic lighting that dramatically enhances quality. Enlighten gives artists total control over lighting, driving a new generation of games that rival film for their manipulation of mood and atmosphere. Licensees of Enlighten include AAA titles in production at EA DICE, CCP and FunCom.

Speaker Chris Doran, Founder & Chief Operating Officer

Speaker Session
4001 - Emerging Companies: CEO on Stage featuring Elemental Technologies, Geomerics, and Milabra (Wednesday, Sept 22, 11:00)

CEO Chris Doran

Investors Undisclosed

Capital Raised Undisclosed

ICD

ICD designs technology. We create purposeful products driven by user centric design principles using the latest hardware and software technologies available. Our products are optimized for usability, aesthetics, technology and purpose. We specialize in NVIDIA chipset technologies running Google and Windows OS. We have basic goals: Simplify. Innovate. Impress.

Speaker David Hayes, CEO

Speaker Session
4008 - Emerging Companies: CEO on Stage featuring ICD and Universal Robotics (Thursday, Sept 23, 11:00)

CEO David Hayes

Investors Undisclosed

Capital Raised Undisclosed

Jedox Business Intelligence

Jedox developed a centralized, in-memory, GPU-based calculation engine that controls and stores the spreadsheet-based business Intelligence data contained in every Excel, Open Office and Google spreadsheet in an organization. This technology stops "Spreadsheet Spreadmart Chaos" (hundreds of spreadsheets with uncoupled, non-verifiable data "running amok" in an organization).

Speaker Kristian Raue, CEO & Founder

Speaker Session
4005 - Emerging Companies: CEO on Stage featuring Jedox Business Intelligence, Rocketick, and Softkinetic (Wednesday, Sept 22, 17:00)

CEO Kristian Raue

Investors Klaus Wecken, eCapital, KfW

Capital Raised 7 Million Euros

From rule breaker

to rising star

to game changer.

SVB *Accelerator*

SVB *Growth*

SVB *Corporate Finance*

We're with you.

Silicon Valley Bank. For nearly three decades, we've led the way in recognizing the vast potential in technology and life science companies. From jump starting start ups to global cash management to providing debt and asset management for industry leaders, SVB understands what you do and has the resources to provide what you need. Beyond commercial banking, we offer venture capital services and funds, valuations and analytics, private banking and more. Silicon Valley Bank.

Named one of Forbes' Top 5 Best Banks.

We'll help you find a way, all the way. svb.com

Gerald Brady, Managing Director
2400 Hanover Street Palo Alto, California 94304
Phone 650.855.3072 E-mail gbrady@svb.com

SVB› *Find a way*

miGenius

Using a powerful combination of mental images Reality Server®, iray® renderer and NVIDIA's CUDA based GPU hardware systems; both businesses and consumers alike can now rapidly and simply upload any 3D content to individually customised websites that can be immediately shared and explored with friends and colleagues, in both accurate photorealistic detail and real-time. miGenius is developing a toolset to enable predominantly 'non-technical' users to easily create customised User Interfaces with a wide range of viewing and management controls and upload their detailed 3D scenes onto either a dedicated GPU server or onto the rapidly emerging 'GPU cloud computing' networks.

Speaker Christopher Blewitt, CEO

Speaker Session
4002 - Emerging Companies: CEO on Stage featuring Allegorithmic, Bunkspeed, and miGenius (Wednesday, Sept 22, 14:00)

CEO Chris Blewitt

Investors No current investor, but we are looking for prospect investors

Capital Raised Undisclosed

Milabra

Milabra uses proprietary, patent pending machine vision software to create visual data about the web. We bridge the visual relationship between the Ad, the Page, and the Audience for increased advertiser performance. We analyze the visual attributes and content of webpages in order to target online display advertising, optimizing creative choice based on the visual environment within which it will appear. We analyze page elements as well as the photo and video content, providing real time data for targeting and optimization as well as defending advertisers against negative content associations.

Speaker Sam Cox, CEO

Speaker Session
4001 - Emerging Companies: CEO on Stage featuring Elemental Technologies, Geomerics, and Milabra (Wednesday, Sept 22, 11:00)

CEO Sam Cox

Investors Undisclosed

Capital Raised Undisclosed

NaturalMotion Ltd

NaturalMotion Ltd is a leading entertainment software company with offices in Oxford (England), San Francisco (California) and Seoul (Korea). 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, Ubisoft, LucasArts, Disney and Bioware. NaturalMotion recently established its own games division, NaturalMotion Games. Its first title, Backbreaker, has since become one of the most successful sports games on mobile devices, with more than 2.2 million downloads.

Speaker Torsten Reil, CEO

Speaker Session
4010 - Emerging Companies: CEO on Stage featuring NaturalMotion, OptiTex, and Useful Progress (Thursday, Sept 23, 15:00)

CEO Torsten Reil

Investors Undisclosed

Capital Raised Undisclosed

OptiTex

OptiTex is the premiere 2D and 3D CAD software 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". OptiTex system consists of three main components: cloth content creation system with our PDS software, 3D Runway Designer, a virtual try-onsystem, 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.

Speaker Yoram Burg, President

Speaker Session
4010 - Emerging Companies: CEO on Stage featuring NaturalMotion, OptiTex, and Useful Progress (Thursday, Sept 23, 15:00)

CEO Yoram Burg

Investors Undisclosed

Capital Raised Undisclosed

Perceptive Pixel

PERCEPTIVE PIXEL

Perceptive Pixel is dedicated to the research, development and deployment of multi-touch interfaces for the knowledge worker. The company's hardware and software solutions enable both novice and expert users to manipulate complex datasets through a new class of intuitive yet powerful and visually rich interface techniques.

Speaker Jeff Han, Founder, Chief Scientist

Speaker Session
4011 - Emerging Companies: CEO on Stage featuring Cinnafilm, Perceptive Pixel and Total Immersion (Thursday, Sept 23, 16:00)

CEO Jeff Han

Investors Undisclosed

Capital Raised Series B

Playcast Media Systems

Playcast

Playcast Media Systems brings video games to the world's largest media distribution platform – Pay TV networks. The company's solution includes a head-end based system, which streams a game's audiovisual content as a standard MPEG stream, as well as the provisioning of the content and programming itself. Playcast's media streaming systems, located in operators' headends, host the games and stream them over the existing video network to an already distributed base of set-top boxes. Playcast is a privately owned, venture capital backed company, based in Israel and the UK.

Speaker Natan Peterfreund, CTO

Speaker Session
4004 - Emerging Companies: CEO on Stage featuring Cooliris, empulse GmbH, and Playcast Media Systems(Wednesday, Sept 22, 16:00)

CEO Guy de Beer

Investors Xenia Ventures and Private Investors

Capital Raised Undisclosed

Ellison & Zander 2009

Hastings & Eisner 2010

Top Ten Tech Trends 2010

Join us **Sept. 28th** for **Churchill Club's** *Annual Dinner 2010* with: Juniper Network's CEO, Kevin Johnson

CHURCHILL CLUB

Business and Technology Insiders

Igniting conversations to encourage innovation and economic growth since 1985.

www.churchillclub.org

Churchill Club is a member-supported non-profit organization. For information about membership and upcoming programs, please visit www.churchillclub.org or contact us at 408.265.0130.

Ignite your own conversations

RTT

Realtime Technology (RTT) AG stands for creative and fascinating 3D visualization solutions, which bring products to life in realtime and portray them in a natural and realistic environment. The company provides its clients with assistance during each stage of the life cycle of their products – from the initial product design stage through to development and subsequent marketing and sales. The 3D data model from the product development stage serves as the basis for all the following steps in the product lifecycle. It can be used, for example, to rapidly create computer generated, photorealistic product illustrations for the marketing department or to develop a 3D online product configurator on a website. In this way, RTT doesn't just speed up decision making and development processes for its clients, but it also opens up new opportunities with regard to marketing and sales. The company was founded in 1999 and its head office is in Munich, Germany. RTT AG has over 400 employees and is represented in 14 locations worldwide. Many leading businesses have put their trust in RTT and its portfolio of clients includes names such as Adidas, Audi, BASF, BMW, Bosch, Daimler, EADS, Harley-Davidson, Miele, Porsche, Samsung, Thyssen-Krupp, Toyota and Volkswagen. RTT AG is a stock market listed company (Xetra:R1T; WKN: 701220; ISIN: DE0007012205). For more information visitwww.rtt.ag.

Speaker Jeroen Snepvangers, President and CEO

Speaker Session
4007 - Emerging Companies: CEO on Stage featuring Aqumin, RTT, and Scalable Display Technologies (Thursday, Sept 23, 10:00)

CEOs Ludwig A. Fuchs and Jeroen Snepvangers

Investors Balderton Capital, Siemens VC, Heliad

Capital Raised 15 Million Euros

Rocketick

Rocketick is revolutionizing the chip verification world by bringing acceleration to the fingertips of every engineer. Rocketsim is the world's first GPU-based logic simulation accelerator .Rocketsim is a highly cost-effective solution. It is a pure software-product that runs on commercial, low cost GPUs that are widely available from NVIDIA.

Speaker Uri Tal, CEO

Speaker Session
4005 - Emerging Companies: CEO on Stage featuring Jedox Business Intelligence, Rocketick, and Softkinetic (Wednesday, Sept 22, 17:00)

CEO Uri Tal

Investors Peregrine Ventures

Capital Raised Undisclosed

Scalable Display Technologies

Scalable is a leading provider of auto-calibration software used to create edge-blended displays. Its patented EasyBlend™ software simplifies the creation of super-resolution, multi-projector displays of the highest quality and scalable size. EasyBlend enables widespread use of multi-projector displays for a new class of simulators, and supports new forms of digital signage/data visualization tools.

Speaker Andrew Jamison, CEO

Speaker Session
4007 - Emerging Companies: CEO on Stage featuring Aqumin, RTT, and Scalable Display Technologies (Thursday, Sept 23, 10:00)

CEO Andrew Jamison

Investors Undisclosed

Capital Raised Undisclosed

Softkinetic

Softkinetic is the leading provider of 3D gesture recognition technologies that transform the way people interact with the digital world. We provide the most advanced software platform for building immersive, transparent and intuitive user experiences within the fields of Interactive Digital Entertainment, Consumer Electronics, Interactive Marketing, Digital Signage and other markets. Our patented 3D gesture recognition middleware, iisu™, supports all 3D cameras and supports all platforms, from PCs to embedded systems and set-top-boxes.

Speaker Michel Tombroff, CEO

Speaker Session
4005 - Emerging Companies: CEO on Stage featuring Jedox Business Intelligence, Rocketick, and Softkinetic (Wednesday, Sept 22, 17:00)

CEO Michel Tombroff

Investors Privately held company

Capital Raised 5 Million Euros

Total Immersion

Total Immersion is the global leader in augmented reality. Through its patented D'Fusion™ technology, Total Immersion blurs the line between the virtual world and the real world by integrating real time interactive 3D graphics into a live video stream. 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.

Speaker Bruno Uzzan, Founder & CEO

Speaker Session
4011 - Emerging Companies: CEO on Stage featuring Cinnafilm, Perceptive Pixel and Total Immersion (Thursday, Sept 23, 16:00)

CEO Bruno Uzzan

Investors Partech International, I Source Gestion, Elaia Partners

Capital Raised $8 Million

Universal Robotics

Universal Robotics creates software that enables machines to learn from their experiences, react and adapt to their surroundings, and perform tasks that are costly, dangerous or difficult for humans to undertake. The company's signature technology, Neocortex, which was developed over seven years at NASA and Vanderbilt University, will increase efficiency and worker safety across industries in applications including warehousing, mining, handling hazardous waste and automating vehicles such as forklifts.

Speaker David Peters, Founder & CEO

Speaker Session
4008 - Emerging Companies: CEO on Stage featuring ICD and Universal Robotics (Thursday, Sept 23, 11:00)

CEO David Peters

Investors Private Equity

Capital Raised Undisclosed

Useful Progress

UsefulProgress, was created in 2003 develops new software strategies and computers intended for 3D imaging with high-definition. UsefulProgress receives the agreement of «R&D Company» since 2007. The company is based in the University Paris Descartes. The applications affect sectors as diverse as medical, pharmaceutical, industrial, mining (diamonds), archeology, higher education.

Speaker Sylvain Ordureau, CEO

Speaker Session
4010 - Emerging Companies: CEO on Stage featuring NaturalMotion, OptiTex, and Useful Progress (Thursday, Sept 23, 15:00)

CEO Sylvain Ordureau

Investors Undisclosed

Capital Raised Undisclosed

SPONSORED BY:

NVIDIA

SYNNEX CORPORATION

GTC SYNNEX NETWORK

Please visit these Tesla Preferred Partner exhibits and be entered in a drawing to win a free **NVIDIA Tesla C2050**!

Ace Computers	Appro	CIARA Technologies	Colfax International
Creative Consultants	Exxact Corporation	GraphStream Inc.	JRTI
Koi Computers	Microway	PSSC Labs	Penguin Computing
SGI			

MONDAY, SEPT 20, 13:00 (80 MINUTES)
MARRIOTT SAN JOSE BALLROOM

2004 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 90-minute tutorial session will cover the key features and differences between the major programming languages, APIs and development tools available today. Attendees will also learn several high level design patterns for consumer, professional and HPC applications, with practical programming considerations for each.

Speaker(s):	Phillip Miller (Director, Workstation Software Product Management, NVIDIA), Holger Kunz (Director, Workstation Software Development, NVIDIA), Brian Harrison (NVIDIA), Thomas Ruge (Software Manager, NVIDIA)
Topic(s):	Programming Languages & Techniques, Tools & Libraries

MONDAY, SEPT 20, 13:00 (80 MINUTES)
ROOM B

2024 NVIDIA Acceleration Engines Overview (Pre-Conference Tutorial)

Come learn of the software engines NVIDIA freely provides to application developers to rapidly leverage new GPU capabilities and dramatically reduce the time it takes to bring compelling features to end users.

Speaker(s):	Phillip Miller (Director, Workstation Software Product Management, NVIDIA), Holger Kunz (Director, Workstation Software Development, NVIDIA), Brian Harrison (NVIDIA)
Topic(s):	Programming Languages & Techniques, Computer Vision, Ray Tracing

MONDAY, SEPT 20, 13:00 (80 MINUTES)
ROOM A5

2157 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.

Speaker(s):	Cem Cebenoyan (Senior Manager, Developer Technology, NVIDIA)
Topic(s):	Computer Graphics, General Interest

MONDAY, SEPT 20, 13:00 (80 MINUTES)
ROOM C

2158 Pipeline with OpenGL (Pre-Conference Tutorial)

This tutorial session teaches attendees how to program the NVIDIA Quadro Digital Video Pipeline with OpenGL. It will go in-depth into the techniques and recommended practices.

Speaker(s):	Thomas True (Applied Engineer, NVIDIA)
Topic(s):	Programming Languages & Techniques, Video Processing, Computer Graphics

MONDAY, SEPT 20, 14:30 (80 MINUTES)
MARRIOTT SAN JOSE BALLROOM

2131 Introduction to CUDA C (Pre-Conference Tutorial)

Starting with a background in C or C++, learn everything you need to know in order to start programming in CUDA C. Beginning with a "Hello, World" CUDA C program, explore parallel programming with CUDA through a number of hands-on code examples. Examine more deeply the various APIs available to CUDA applications and learn the best (and worst) ways in which to employ them in applications. Master the first half of the book "CUDA by Example" as taught by the author, pointing you on a trajectory to complete the second half on your own after course completion.

Speaker(s):	Jason Sanders (Senior Software Engineer, NVIDIA)
Topic(s):	Programming Languages & Techniques,

MONDAY, SEPT 20, 14:30 (80 MINUTES)
ROOM C

2159 Programming the NVIDIA Digital Video Pipeline with Direct3D (Pre-Conference Tutorial)

Learn how to program the NVIDIA Quadro Digital Video pipeline using Direct3D. This session will provide an overview of the SDK, discuss device control, data transfers, performance measuring and tuning, ancillary data and application design considerations.

Speaker(s):	Thomas True (Applied Engineer, NVIDIA)
Topic(s):	Programming Languages & Techniques, Video Processing, Computer Graphics

MONDAY, SEPT 20, 14:30 (80 MINUTES)
ROOM A5

2260 DirectCompute (Pre-Conference Tutorial)

Learn how to to use the DirectCompute API to solve GPU computing problems. This tutorial will introduce the DirectCompute API, cover the recommended best practices for GPU programming, and go over examples of how to use this API efficiently and effectively to solve compute-intensive problems.

Speaker(s):	Eric Young (Manager of Developer Technology Professional and Consumer Applications, NVIDIA)
Topic(s):	Programming Languages & Techniques

MONDAY, SEPT 20, 14:30 (80 MINUTES)
ROOM B

2261 Introduction to GPU Ray Tracing with NVIDIA OptiX (Pre-Conference Tutorial)

Learn how to use NVIDIA OptiX to quickly develop high performance ray tracing applications for interactive rendering, offline rendering, or scientific visualization. This session will explore the latest available OptiX version.

Speaker(s):	Dave McAllister (Development Lead, OptiX Development, NVIDIA), Phillip Miller (Director, Workstation Software Product Management, NVIDIA)
Topic(s):	Ray Tracing, High Performance Computing, Computer Graphics

MONDAY, SEPT 20, 16:00 (80 MINUTES)
ROOM C

2010 Implementing Stereoscopic 3D in Your Applications (Pre-Conference Tutorial)

Let's dive into the 3rd dimension. This talk presents a comprehensive technical overview of NVIDIA's stereo technology and tools. After a complete introduction to NVIDIA's stereo technology, we will then explore in more detail production techniques for the new artistic space of effects and creativity offered by 3D stereo. The take away of this session will be a solid understanding of NVIDIA's stereo technology and how to take best advantage of it.

Speaker(s):	Samuel Gateau (Developer Technology Engineer, NVIDIA), Steve Nash (Applied Engineer, NVIDIA)
Topic(s):	Programming Languages & Techniques, Stereoscopic 3D

MONDAY, SEPT 20, 16:00 (80 MINUTES)
MARRIOTT SAN JOSE BALLROOM

2018 OpenCL on the GPU (Pre-Conference Tutorial)

OpenCL is Khronos' new open standard for parallel programming of heterogeneous systems. This tutorial session will introduce the main concepts behind the standard and illustrate them with some simple code walkthrough. Attendees will also learn how to make efficient use of the API to achieve good performance on the GPU.

Speaker(s):	Cliff Woolley (CUDA Developer Technology Engineer, NVIDIA)
Topic(s):	Tools & Libraries

MONDAY, SEPT 20, 16:00 (80 MINUTES)
ROOM A3

2127 OpenGL (Pre-Conference Tutorial)

This session will discuss the latest OpenGL features offered by NVIDIA for both Quadro and Geforce line of products. Learn more about OpenGL 4 as well as NVIDIA specific OpenGL extensions.

Speaker(s):	Mark Kilgard (Principal System Software Engineer, NVIDIA)
Topic(s):	Programming Languages & Techniques,

MONDAY, SEPT 20, 16:00 (80 MINUTES)
ROOM B

2245 Parallel Nsight for Microsoft Visual Studio (Pre-Conference Tutorial)

NVIDIA Parallel Nsight provides access to the power of the GPU from within the familiar environment of Microsoft Visual Studio. In this session, you will learn how to use Parallel Nsight to develop GPU computing and graphics applications.

Learn how to use the powerful Parallel Nsight debugger to identify errors in CUDA C/C++ kernels and HLSL shaders using GPU breakpoints and direct memory and variable inspection. See how Parallel Nsight displays system-wide performance characteristics, allowing you to create efficient GPU algorithms.

Speaker(s):	Kumar Iyer (Product Manager, NVIDIA)
Topic(s):	Tools & Libraries

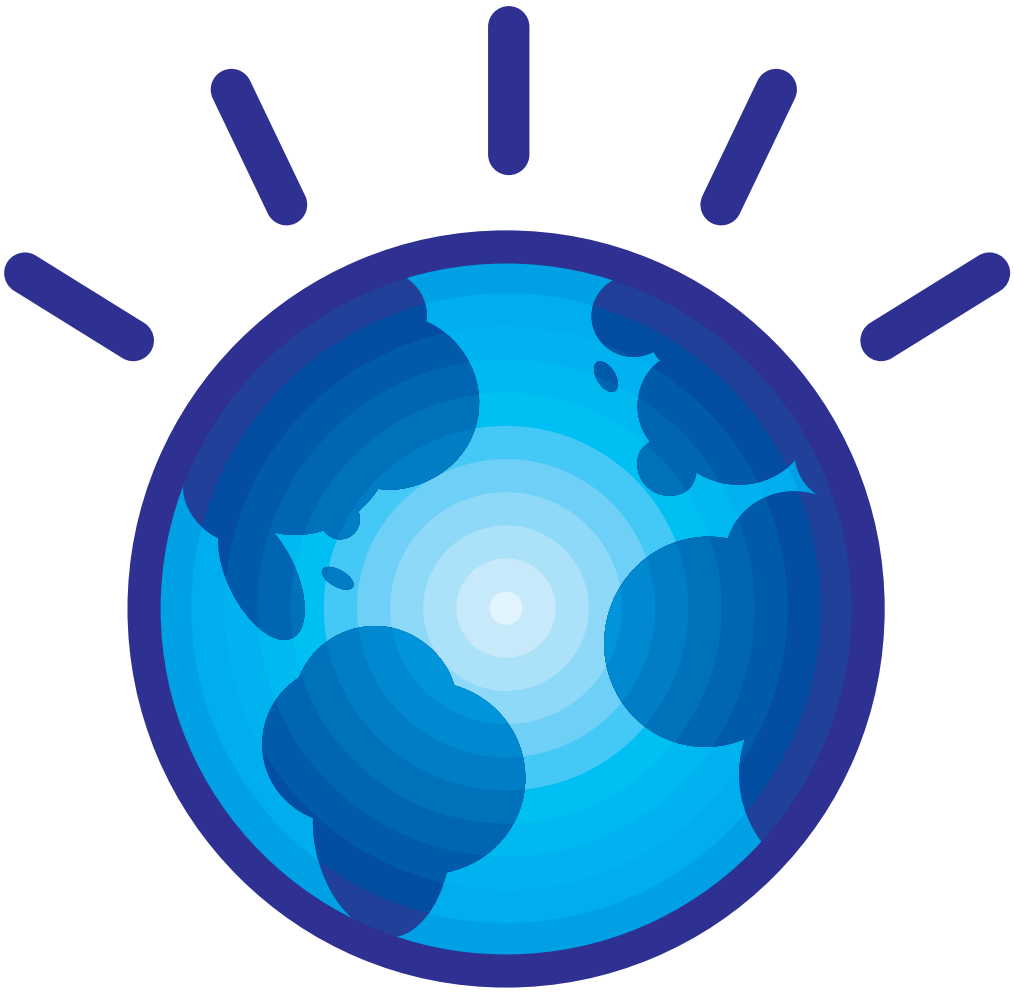

Welcome to the decade of smart.

A year ago, we began a global conversation about how our planet can become smarter.

One year into this new era, the signs of a smarter planet are all around us. Smarter systems are creating value in every major industry and across every region of both the developed and developing worlds.

Intelligence is being infused into the systems and processes that make the world work—into things no one would recognize as computers: cars, appliances, roadways, power grids, clothes, even natural systems such as agriculture and waterways.

Trillions of digital devices, connected through the Internet, are producing a vast ocean of data. And that information can now be turned into knowledge because we have the computational power and advanced analytics to make sense of it all.

In a study of 439 cities, those with transportation congestion systems reduced average travel delays by more than 700,000 hours annually.

Eight hospitals and 470 primary care clinics improved clinical results and operational efficiency by up to 10% through information access at the point of care.

Leading retailers reduced supply chain costs by up to 30% and increased sales by up to 10%.

With sophisticated mathematical models, we can actually begin to predict and react to changes in our systems. New York has smart crime fighting. Galway has smart water. A smart grid in Copenhagen keeps energy flowing.

We've learned a lot over the past year about what it takes to build a smarter planet. We've also learned about the issues it raises—like protecting personal information and securing critical infrastructures.

The good news is that business leaders, policymakers and government officials around the world are stepping up to these challenges. Above all, they realize that we cannot let this moment pass. The time to act is now, and the way to act is together. The decade of smart is under way.

Let's build a smarter planet. Join us and see what others are doing at **ibm.com**/smarterplanet

TUESDAY, SEPT 21, 09:00 (90 MINUTES)
KEYNOTE HALL

1001 Opening Keynote with Jen-Hsun Huang, NVIDIA

Do not miss this opening keynote, featuring Jen-Hsun Huang, CEO and Co-Founder of NVIDIA and special guests. Hear about what's next in computing and graphics, and preview disruptive technologies and exciting demonstrations from across industries.

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.

Speaker(s):	Jen-Hsun Huang (CEO & Co-Founder, NVIDIA)
Topic(s):	General Interest

TUESDAY, SEPT 21, 11:00 (50 MINUTES)
MARRIOTT SAN JOSE BALLROOM

2223 Academic Welcome Social and Poster Review

This session is open to academic attendees only. We invite you to join your fellow academics to preview this year's NVIDIA Research Summit Posters and mingle with your colleagues. Included will be a special presentation from our 2010-2011 Graduate Fellowship recipients to showcase the research that earned them this prestigious award. These students were selected from 268 applications in 28 countries. Their research confronts a variety of challenges of immense technical and strategic importance, including light-transport simulation, computer vision, programmability and optimization for heterogeneous systems, and much more. We believe that these minds lead the future in our industry. Light snacks and drinks will be served.

Speaker(s):	Bill Dally (NVIDIA), David Luebke (NVIDIA),NVIDIA Graduate Fellows
Topic(s):	General Interest

TUESDAY, SEPT 21, 11:00 (50 MINUTES)
ROOM K

2047 Bridging Ray and Raster Processing on GPUs

Explore new techniques in real time rendering. We will discuss a system for ray traced global illumination (GI) carefully integrated with a traditional raster renderer using an incremental irradiance cache. Covers novel GPU methods for spawning secondary GI rays on only visible cells, smoothly sampling the visible 3D cache into 2D, and incrementally ray traced spherical harmonics basis. Details applying a range of optimizations to achieve real-time frame rates with the OptiX ray tracing engine.

Speaker(s):	Kenny Mitchell (Research Lead, Black Rock Studio)
Topic(s):	Ray Tracing

TUESDAY, SEPT 21, 11:00 (50 MINUTES)
ROOM C

2079 A Fast, Scalable High-Order Unstructured Compressible Flow Solver

We will describe a scalable and efficient high-order unstructured compressible flow solver for GPUs. The solver allows the achievement of arbitrary order of accuracy for flows over complex geometries. High-order solvers require more operations per degree of freedom, thus making them highly suitable for massively parallel processors. Preliminary results indicate speed-ups up to 70x with the Tesla C1060 compared to the Intel i7 CPU. Memory access was optimized using shared and texture memory.

Speaker(s):	David M. Williams (Ph.D. Candidate, Stanford University), Patrice Castonguay (Ph.D. Candidate, Stanford University)
Topic(s):	Computational Fluid Dynamics, Algorithms & Numerical Techniques, Physics Simulation

TUESDAY, SEPT 21, 11:00 (50 MINUTES)
ROOM A2

2096 High-Speed CT Reconstruction in Medical Diagnosis & Industrial NDT Applications

We present the software platform CERA developed by Siemens, which utilizes (multiple) graphics processing units (GPUs) in order to deliver high-speed CT reconstructions, and describe its implementation challenges using CUDA and OpenCL. We further show how GPU acceleration enables the utilization of reconstruction approaches which provide highly improved reconstruction quality in NDT applications.

Speaker(s):	Holger Scherl (Computer Scientist, Siemens AG)
Topic(s):	Medical Imaging & Visualization, Imaging

TUESDAY, SEPT 21, 11:00 (50 MINUTES)
ROOM N

2112 The Heisenberg Spin Glass Model on GPU: Myth versus Fact

Dive into implementations of the 3D Heisenberg spin glass model for GPUs. We will discuss results showing that fast shared memory gives better performance with respect to slow global memory only under certain conditions. Covers careful kernel tuning to achieve significant speedup with respect to a state-of-art high end multicore processor.

Speaker(s):	Massimo Bernaschi (Professor, Istituto Applicazioni del Calcolo - C.N.R.)
Topic(s):	Physics Simulation

TUESDAY, SEPT 21, 11:00 (50 MINUTES)
ROOM A3

2119 Supercomputing for the Masses: Killer-Apps, Parallel Mappings, Scalability and Application Lifespan

Hear the latest on how supercomputing for the masses is changing the world. We will look at some of the one- to three-orders of magnitude faster killer apps and see how they do it. We will discuss specific mapping to GPGPU hardware and techniques for high performance and near-linear scalability both within and across multiple GPGPUs. We will also consider software investment and the decades long longevity of some successful massively parallel Investments in multithreaded software, scalability, balance metrics, lack of consensus on programming models, and lifecycle considerations.

Speaker(s):	Robert Farber (Senior Scientist, PNNL)
Topic(s):	High Performance Computing, Algorithms & Numerical Techniques, Machine Learning & Artificial Intelligence, Physics Simulation

TUESDAY, SEPT 21, 11:00 (50 MINUTES)
ROOM A7

2130 GPU Computing and a Revolution in Design Engineering

Join design enginneering experts for a discussion of the technology needs limiting their use of GPUs and how vendors are addressing those needs. We will cover performance in analysis, simulation, and rendering, as well as how to use GPUs to accelerate and improve the engineering design process.

Speaker(s):	Peter Varhol (HPC Editor, Desktop Engineering Magazine)
Topic(s):	General Interest

TUESDAY, SEPT 21, 11:00 (50 MINUTES)
ROOM A8

2132 Accelerating Biologically Inspired Computer Vision Models

Join us for a discussion on applying commodity-server-based clusters and GPU-based clusters to simulating computer vision algorithms at a scale that approaches that of biological vision. We consider the limitations of each technology, survey approaches taken thus far, and suggest new hybrid models and programming frameworks to overcome current limitations and substantially improve performance.

Speaker(s):	Tom Dean (Research Scientist, Google Inc.)
Topic(s):	Computer Vision, Machine Learning & Artificial Intelligence

TUESDAY, SEPT 21, 11:00 (50 MINUTES)
ROOM A1

2165 Rendering Revolution

Learn how GPU technologies are transforming the making of pixels. This talk will cover GPU-centric rendering techniques that leverage both the raw computational capabilities of NVIDIA's GPUs and advanced pixel-shading techniques for interactive visualization and rendering.

Speaker(s):	Ken Pimentel (Director, Media & Entertainment, Autodesk)
Topic(s):	Computer Graphics, Film

TUESDAY, SEPT 21, 11:00 (50 MINUTES)
ROOM D

2172 Unveiling Cellular & Molecular Events of Cardiac Arrhythmias

George Mason University is using CUDA technology to get a 20x speed-up in simulations of intracellular calcium dynamics, thought to play a major role in the generation of cardiac arrhythmias. We will discuss the novel algorithms we have developed for Markov Chain Monte Carlo Simulation and their use in investigating elementary events of calcium release in the cardiac myocyte. The resulting extremely fast simulation time has generated new insights into how defects in the control of intracellular calcium may lead to cardiac arrhythmia.

Speaker(s):	Tuan Hoang-Trong (PhD student, George Mason University)
Topic(s):	Life Sciences, Algorithms & Numerical Techniques, Physics Simulation

TUESDAY, SEPT 21, 11:00 (50 MINUTES)
ROOM L

2214 Faster Simulations of the National Airspace System

Learn about twenty-four hour, fast-time simulations of traffic in the National Airspace System, which use GPU technology to help perform key steps in the trajectory prediction of flights. GPUs enabled us to improve the runtime by up to two orders of magnitude versus the previously required tens of minutes per execution. We will present a brief overview of the problem domain and a description of how the GPU has opened doors to uncharted research areas.

Speaker(s):	Joseph Rios (Research Aerospace Engineer, NASA)
Topic(s):	General Interest

TUESDAY, SEPT 21, 11:00 (50 MINUTES)
ROOM A5

2267 GPU Computing with MATLAB®

MATLAB is a widely used tool for scientific, engineering and financial applications. As the popularity of GPUs has grown, there is strong interest from engineers and scientists who solve computationally intensive problems to be able to leverage GPUs within MATLAB and other products from MathWorks. This talk will discuss how MathWorks tools can help engineers and scientist to take advantage of GPU resources while continuing to work in the familiar MATLAB environment. A range of capabilities will be discussed and demonstrated.

Speaker(s):	Loren Dean (Director of Engineering, MATLAB Products, MathWorks)
Topic(s):	Tools & Libraries

TUESDAY, SEPT 21, 11:30 (20 MINUTES)
ROOM B

2149 Overview of Parallel Nsight for Visual Studio

NVIDIA Parallel Nsight provides access to the power of the GPU from within the familiar environment of Microsoft Visual Studio. This session is an entry level overview of the GPU computing and graphics development features of Parallel Nsight as well as a glimpse into the future of this powerful tool.

Speaker(s):	Kumar Iyer (Product Manager, NVIDIA)
Topic(s):	Tools & Libraries

TUESDAY, SEPT 21, 12:00 (120 MINUTES)
EXHIBIT HALL

1004 Exhibits Open / Networking Lunch

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.

Topic(s):	General Interest

TUESDAY, SEPT 21, 14:00 (50 MINUTES)
ROOM A8

2013 iray - GPUs and the Photorealistic Rendering Revolution

Hear about the ongoing revolution in the production of photorealistic imagery being powered by GPUs. We will explore the algorithms and concepts behind iray – a CUDA accelerated software library from mental images/NVIDIA that provides an interactive, push-button, fast synthetic digital camera in software to a variety of OEM applications and platforms. We will demonstrate iray embedded in commercial CAD and Digital Content Creation applications as well as in 3D cloud computing platforms.

Speaker(s):	Michael Kaplan (Vice President of Strategic Development, mental images/NVIDIA)
Topic(s):	Digital Content Creation (DCC), Cloud Computing, Ray Tracing

TUESDAY, SEPT 21, 14:00 (50 MINUTES)
ROOM C

2015 Efficient Tridiagonal Solvers for ADI methods and Fluid Simulation

Learn about new techniques to efficiently implement the Alternating Direction Implicit method on GPU for large 2D and 3D domains with complex boundaries.

A novel tridiagonal solver for systems with variable sizes and a new hybrid approach will be covered in detail. Comprehensive performance analysis and key Fermi optimizations will be explored.

Various applications of tridiagonal solvers such as 3D direct numerical fluid simulation and a 2D depth-of-field effect for games will be briefly discussed.

Speaker(s):	Nikolai Sakharnykh (Developer Technology Engineer, NVIDIA)
Topic(s):	Algorithms & Numerical Techniques, Computational Fluid Dynamics

TUESDAY, SEPT 21, 14:00 (50 MINUTES)
ROOM A5

2019 GPU-Accelerated Internet Technologies & Trends

Join us for a whirlwind demo-punctuated tour of up-and-coming technologies that promise to bring GPU acceleration to the Worldwide Web. We'll cover 2D graphics, 3D graphics and video. In addition to summarizing the emerging standards and technologies, performance test results showing how they scale on various GPUs will be presented, along with recommendations for how to design for best performance. Finally, adoption trends and ecosystem dynamics will be summarized. Attendees should leave with a richer understanding of the possibilities enabled by the GPU-Accelerated Web, and new insights into when and how it will matter.

Speaker(s):	Chris Pedersen (Market Development Manager, NVIDIA)
Topic(s):	GPU Accelerated Internet, Stereoscopic 3D, Video Processing

TUESDAY, SEPT 21, 14:00 (50 MINUTES)
ROOM K

2028 Mathematica for GPU Programming

Mathematica is widely used in scientific, engineering, mathematical fields and education. In this session, new tools for general GPU programming in the next release of Mathematica are presented. These tools build on top of Mathematica's technology which provides a simple, yet powerful, interface to the large base of compiling tools. Applications of CUDA and OpenCL from within Mathematica will be presented. These examples will provide a general overview of the powerful development environment for GPU programming that Mathematica can offer not just for researchers but for anybody with basic knowledge of Mathematica and GPU programming.

Speaker(s):	Ulises Cervantes-Pimentel (Senior Kernel Developer, Wolfram Research)
Topic(s):	Programming Languages & Techniques, Algorithms & Numerical Techniques, Imaging, Tools & Libraries

TUESDAY, SEPT 21, 14:00 (50 MINUTES)
ROOM A3

2057 CUDA-Accelerated LINPACK on Clusters

This talk will illustrate the use of GPUs to accelerate the LINPACK benchmark on clusters with GPUs, where both the CPUs and the GPUs are used in synergy. The acceleration is obtained executing DGEMM (matrix multiply) and DTRSM (for the solution of triangular systems) calls simultaneously on both GPU and CPU cores. Details of the implementation will be presented together with results that shows how effective the solution is, both for performance and power efficiency.

Speaker(s):	Everett Phillips (Applied Engineer - GPU Computing, NVIDIA), Massimiliano Fatica (Manager, NVIDIA)
Topic(s):	High Performance Computing, Algorithms & Numerical Techniques

TUESDAY, SEPT 21, 14:00 (50 MINUTES)
ROOM A2

2094 Nearly Instantaneous Reconstruction for MRIs

GE's Autocalibrating Reconstruction for Cartesian Imaging (ARC) is a computationally intensive, widely used algorithm in MRI Reconstruction using Parallel Imaging. We demonstrate that an optimized CUDA implementation of ARC on a GPU can enable nearly instantaneous reconstruction and speedups of up to 10x over an optimized dual socket QuadCore CPU implementation. We will discuss challenges both with computational intensity and data read/write efficiency. We will also compare the Fermi C2050 with the C1060.

Speaker(s):	Babu Narayanan (Lab Manager, GE Global Research)
Topic(s):	Medical Imaging & Visualization, High Performance Computing

TUESDAY, SEPT 21, 14:00 (50 MINUTES)
ROOM B

2150 Parallel Nsight: Debugging Massively Parallel Applications [Advanced]

Data parallel algorithms that provide real-time financial options pricing or identification of hidden oil reserves are utilizing the massively parallel nature of the GPU for industry changing performance gains. Developers require industry standard development tools to create the software that accomplishes these parallel tasks.

NVIDIA Parallel Nsight delivers the power of the GPU within the familiar environment of Microsoft Visual Studio. In this session, you will learn advanced techniques for debugging CUDA C/C++ and DirectCompute code using Parallel Nsight, including conditional and data breakpoints as well as out of bound GPU memory access detection.

Speaker(s):	Sebastien Domine (Sr. Dir. Developer Tools, NVIDIA)
Topic(s):	Tools & Libraries, Programming Languages & Techniques

TUESDAY, SEPT 21, 14:00 (50 MINUTES)
ROOM A1

2152 Using Virtual Texturing to Handle Massive Texture Data

A virtual texture implementation allows applications the ability to manage gigantic amounts of texture data for rendering complex data sets. However, practical utilization involves feeding it adequate data. The GPU offers a powerful engine capable of accelerating the transcoding of efficient storage formats into formats useful for rendering. This session will demonstrate a virtual texturing implementation and the steps needed to GPU accelerate the non-rendering portions of managing and loading the virtual texture data.

Speaker(s):	Evan Hart (Software Engineer, NVIDIA)
Topic(s):	Computer Graphics

TUESDAY, SEPT 21, 14:00 (50 MINUTES)
ROOM A7

2222 Working Man's Guide to 3D Video Editing

Video editing is currently at two simultaneous inflections points: use of GPUs for video processing and the beginning of wide spread adoption of 3D. At this time however, identifying and navigating through the necessary tools and equipment to create compelling 3D video content is challenging.

This session is intended to provide a pragmatic guide to creating prosumer 3D video content and how the GPU greatly

assists and speeds up this process.

The intended audience is anyone interested in how to create compelling 3D movies at a prosumer level.

Speaker(s):	Ian Williams (Director PSG Applied Engineering, NVIDIA), Kevan O'Brien (NVIDIA)
Topic(s):	Digital Content Creation (DCC), Stereoscopic 3D

TUESDAY, SEPT 21, 14:00 (50 MINUTES)
ROOM M

2233 Solving Your GPU Computing Needs (Sponsored by HP)

In this session we will go into detail and you will learn about HP's GPU enabled systems, from Workstations to our GPU enabled servers and clusters. You will get the latest information on configurations, options, GPU management and use cases.

Speaker(s):	Dave Korf (Marketing, HP), Will Wade (Business Alliance Manager, HP)
Topic(s):	High Performance Computing

TUESDAY, SEPT 21, 14:00 (50 MINUTES)
MARRIOTT SAN JOSE BALLROOM

2262 CUDA Centers of Excellence Super-Session I

Come hear about the groundbreaking research taking place at the CUDA Centers of Excellence, an elite group of world-renown research universities that are pushing the frontier of massively parallel computing using CUDA. Researchers from these top institutions will survey cutting-edge research that is advancing the state of the art in GPU computing and dozens of application fields across science and engineering.

In this session we will hear from Professor Hanspeter Pfister of Harvard University and Professor Jeff Vetter of Georgia Tech and Oak Ridge National Laboratory.

Speaker(s):	Hanspeter Pfister (Professor, Harvard University), Jeffrey Vetter (Professor, Georgia Tech / Oak Ridge National Laboratory)
Topic(s):	General Interest

TUESDAY, SEPT 21, 14:00 (50 MINUTES)
ROOM L

2276 Using GPUs to Run Next-Generation Weather Models

We are using GPUs to run a new weather model being developed at NOAA's Earth System Research Laboratory (ESRL) called the Non-hydrostatic Icosahedral Model (NIM). NIM is slated to run at high resolution (4km global scale) within two years. This presentation will highlight work required to parallelize and run the NIM. We will describe progress running on multiple GPUs, report on our evaluation of two FORTRAN GPU compilers, and give performance updates of NIM using Fermi. We will also discuss special challenges developing and running operational weather models on GPUs.

Speaker(s):	Mark Govett (Chief, Advanced Computing Section, NOAA Earth System Research Laboratory)
Topic(s):	General Interest

TUESDAY, SEPT 21, 14:00 (20 MINUTES)
ROOM N

2299 Integrating CUDA BLAS with IMSL Fortran

As GPU hardware becomes more prevalent in both research and commercial institutions, software that takes advantage of this specialized hardware is growing in demand. In many cases, it is infeasible or impossible to rewrite an existing program to run entirely on the GPU, so the goal is often to offload as much work as possible. As the IMSL Library team at Rogue Wave Software considers how best to tackle the GPU realm with a general mathematical library, the IMSL Fortran Library takes an initial step where the CUDA BLAS library is utilized to offload CPU work to GPU hardware. This presentation will discuss the approach and architecture of the solution. Benchmark results will show where success has been found. Plans for future products will also be covered.

Speaker(s):	Chris Gottbrath (Principal Product Manager, TotalView Technologies, Inc., a Rogue Wave Software company)
Topic(s):	Tools & Libraries

TUESDAY, SEPT 21, 14:00 (50 MINUTES)
ROOM D

2303 Using Tegra to Solve The Electric Car Power Dilemma

Explore how advanced SoC technologies are transforming the world of automotive industry. Learn on how using nVidia Tegra increased the available range while pushing the envelope on next-gen driver experience. Sharing the lessons learned in the world of electric cars and challenges in constructing a mass production electric vehicle.

Speaker(s):	Theo Valich (President, Bright Side Network Inc.)
Topic(s):	Embedded & Automotive, Computer Vision, Video Processing, Computer Graphics

TUESDAY, SEPT 21, 15:00 (50 MINUTES)
ROOM A3

2017 Lessons Learned Deploying the World's First GPU-Based Petaflop System

Learn what to expect when deploying PetaFLOP or larger systems. The June 2010 list of the Top 500 computer systems featured the first GPU based cluster to exceed 1 PetaFLOP of foating point power -- a system that was built in a fraction of the time and the cost a CPU-only system of that performance would have required. An overview of how system builders and administrators should prepare for large-scale HPC deployments.

Speaker(s):	Dale Southard (Senior Solution Architect, NVIDIA)
Topic(s):	High Performance Computing

TUESDAY, SEPT 21, 15:00 (50 MINUTES)
ROOM A2

2074 Driving a Product from Rasterization to Ray Tracing: The Developer Experience

Learn from the challenges encountered while using DirectX to update the Bunkspeed Move rasterization engine to work with Mental Images' iRay. This work was part of the creation of Bunkspeed Shot, which allows the user to leverage both the high quality image generation of iRay and a highly interactive, good

CUDA

Open GL

ActionScript Accelerator

Multi Parallel Animations

3D Apps

The VGK platform is a collection of software tools that enable Flash content to run on a variety of systems, from low-power electronics and mobile devices to high-end 3D desktop and game systems. Flash can be integrated into your 2D and 3D applications with ease.

PUT SOME FLASH IN YOUR APPS!

Go to www.animatedmedia.ca for our demos

Contact Us: sales@animatedmedia.ca or 416.977.7187

Follow us on Twitter @animatedmedia

ANIMATED MEDIA INC

quality rasterization engine (used for quick setup of a scene). Covers major differences between a ray tracing based interactive system, including GPU based ray tracing, and a traditional GPU rasterization engine.

Speaker(s):	Nick Gebbie, (Senior Graphics Programmer, Bunkspeed)
Topic(s):	Ray Tracing

TUESDAY, SEPT 21, 15:00 (50 MINUTES)
ROOM C

2085 Tridiagonal Solvers: Auto-Tuning and Optimizations

In this presentation, we will discuss and analyze the performance of three optimization techniques for tridiagonal solvers. We first present a hybrid Parallel Cyclic Reduction(PCR)-Gaussian Elimination(GE) tridiagonal solver, which combines work-efficient and step-efficient algorithms for high performance. We further discuss an auto-tuned variant of this technique which selects the optimal switching point between algorithms on a per-machine basis. Next, we present a technique to handle large systems, where shared memory constraints prohibit previous work to solve these systems directly. Finally, we will discuss optimizations on a cyclic reduction technique that avoid bank conflicts on current hardware.

Speaker(s):	Andrew Davidson (Graduate Student, University of California, Davis), Yao Zhang (Graduate Student, University of California, Davis)
Topic(s):	Algorithms & Numerical Techniques, Computational Fluid Dynamics

TUESDAY, SEPT 21, 15:00 (50 MINUTES)
ROOM N

2090 Developing Highly Scalable Particle-Mesh Codes for GPUs: A Generic Approach

Dive deep into a multi-parallel Particle in Cell code that utilizes MPI, pthreads, and CUDA. Around this specific application a general C++ framework for transparent data transfers between GPUs has been developed and will be presented. Further techniques employed include interleaving of communication and computation, particle tiling and a study of how well CUDA performance can be transferred to OpenCL.

Speaker(s):	Guido Juckeland (Senior System Engineer (HPC), Leader Hardware Accelerator Group, TU Dresden - ZIH), Michael Bussmann (Junior Group Leader Computational Radiation Physics, Forschungszentrum Dresden-Rossendorf)
Topic(s):	Physics Simulation, Astronomy & Astrophysics, High Performance Computing

TUESDAY, SEPT 21, 15:00 (50 MINUTES)
ROOM L

2103 Development of an Efficient GPU-Accelerated Model for Fully Nonlinear Water Waves

This work is concerned with the development of an efficient high-throughput scalable model for simulation of fully nonlinear water waves (OceanWave3D) applicable to solve and analyze large-scale problems in coastal engineering. The goal can be achieved through algorithm redesign and parallelization of an optimized sequential single-CPU algorithm based on a flexible-order Finite Difference Method. High performance is pursued by utilizing many-core processing in the model focusing on GPUs for acceleration of code execution. This involves combining analytical methods with an algorithm redesign of the current numerical model.

Speaker(s):	Allan Peter Engsig-Karup (Assistant Professor, Scientific Computing, Technical University of Denmark)
Topic(s):	Computational Fluid Dynamics, Algorithms & Numerical Techniques, Physics Simulation

TUESDAY, SEPT 21, 15:00 (50 MINUTES)
ROOM A8

2113 WebGL: Bringing 3D to the Web

WebGL is a newly-emerging standard for 3D graphics and visual computing on the web. Supported and developed by major web browser vendors, WebGL enables rich interactive 3D graphics delivered through a web browser, on both desktop and mobile platforms. This session will contain an introduction to WebGL, and will focus application development issues unique to the web platform, optimization concerns, and how web technologies such as offline app support, HTML5 video and audio, File and WebSockets integrate with WebGL. Experienced OpenGL developers will learn how to transition their knowledge to WebGL development.

Speaker(s):	Vladimir Vukicevic (Principal Engineer, Mozilla Corporation)
Topic(s):	GPU Accelerated Internet, Tools & Libraries, Computer Graphics

TUESDAY, SEPT 21, 15:00 (50 MINUTES)
ROOM B

2147 GPGPU Development for Windows HPC Server

Attend this demo-driven session to see how to schedule jobs to a Windows compute cluster that includes GPUs. We will also demonstrate GPU-enhanced versions of some commonly used HPC open-source codes, and show how NVIDIA Parallel Nsight™ can be used to debug GPU applications on a cluster. Provides a brief introduction to performance profiling tools that allow developers to analyze system, CPU and GPU events.

Speaker(s):	Calvin Clark (Senior Consultant, Microsoft)
Topic(s):	High Performance Computing, Tools & Libraries

TUESDAY, SEPT 21, 15:00 (20 MINUTES)
ROOM K

2148 Rapid Prototyping and Visualization with OpenCL Studio

Learn about OpenCL Studio, an integrated OpenCL and OpenGL development environment for parallel programming and visualization. We will discuss building end user applications and using its integrated visualization capabilities to better understand the output and internal structure of parallel algorithms. We will also demonstrate its capabilities using several sample applications including particle systems, volumetric rendering, and image processing.

Speaker(s):	Jochen Stier (Founder, Geist Software Labs)
Topic(s):	Tools & Libraries

TUESDAY, SEPT 21, 15:00 (50 MINUTES)
ROOM A7

2224 GPU Acceleration in Adobe Creative Tools

Hear experts explain how Adobe Creative Suite 5 harnesses the power of CUDA technology in several of its core software applications. We will focus on the complete redesign of the core video playback and rendering engine in Adobe Premiere Pro CS5 and how it uses the power of GPUs to deliver superior

performance and change the game for Adobe in professional video production.

Speaker(s):	Paul Young (Adobe), Steve Hoeg (Adobe), Al Mooney (Adobe)
Topic(s):	Video Processing, Imaging

TUESDAY, SEPT 21, 15:00 (50 MINUTES)
ROOM A1

2227 OpenGL 4.0 Tessellation for Professional Applications

The new generation of accelerated graphics is elevating visual computing to new heights. Tessellation, one of its most anticipated features, is already used in many scenarios to bring 3D graphics to an unprecedented level of realism.

This talk will introduce tessellation using OpenGL 4.0. We will also describe how an existing application can be adapted to efficiently take advantage of this new feature and also how to overcome some of the challenges.

Speaker(s):	Philippe Rollin (Applied Engineer, NVIDIA)
Topic(s):	Computer Graphics, Tools & Libraries

TUESDAY, SEPT 21, 15:00 (50 MINUTES)
ROOM A5

2235 Advanced Medical Volume Rendering and Segmentation on the GPU

Learn how to speed up your interactive medical visualization pipeline by an order of magnitude and dramatically improve rendering quality at the same time. Leading researchers in medical imaging informatics describe recent advances in volume visualization and interactive segmentation.

Emphasis is on the underlying parallel GPU algorithms and acceleration data structures.

Speaker(s):	Mike Roberts (Research Assistant, Hotchkiss Brain Institute, University of Calgary, Canada), Eric Penner (Research Associate, Hotchkiss Brain Institute, University of Calgary, Canada)
Topic(s):	Medical Imaging & Visualization, Algorithms & Numerical Techniques, Computer Vision, Computer Graphics

TUESDAY, SEPT 21, 15:00 (50 MINUTES)
MARRIOTT SAN JOSE BALLROOM

2263 CUDA Centers of Excellence Super-Session II

Come hear about the groundbreaking research taking place at the CUDA Centers of Excellence, an elite group of world-renown research universities that are pushing the frontier of massively parallel computing using CUDA. Researchers from these top institutions will survey cutting-edge research that is advancing the state of the art in GPU computing and dozens of application fields across science and engineering.

In this session we will hear from Dr. Wei Ge at the Chinese Academy of Science, Professor Amitabh Varshney at the University of Maryland, and Adjunct Assistant Professor Stan Tomov at the University of Tennessee – Knoxville.

Speaker(s):	Stan Tomov (University of Tennessee), Amitabh Varshney (Professor, University of Maryland), Wei Ge (Professor, Institute of Process Engineering, Chinese Academy of Sciences)
Topic(s):	General Interest

TUESDAY, SEPT 21, 15:00 (50 MINUTES)
ROOM M

2270 Appro's GPU Computing Solutions

Learn how GPU's are changing the High Performance Computing landscape to deliver price/performance levels that were previously considered unachievable. Join Appro (http://www.appro.com), a leading provider of supercomputing solutions; to discuss the introduction of the Appro Tetra server, the most powerful GPU server available today in a 1U form factor and the availability of a new modular GPU expansion blade, both based on NVIDIA Tesla 20-series GPUs. The availability of these two products is a confirmation of Appro's commitment in providing the most innovative and powerful computing platforms at very attractive prices to the High Performance Computing markets.

Speaker(s):	John Lee (Vice President, Appro International, Inc.)
Topic(s):	High Performance Computing

TUESDAY, SEPT 21, 15:30 (20 MINUTES)
ROOM K

2268 Think Data-Parallel! Building Data-Parallel Code with M

Discover and leverage parallelism inherent in pre-existing codes. Often times, parallelism is hidden in seemingly serial programs. This is due obfuscation via indexing or looping wherein the parallelism is seemingly non-existent. Several real-world examples of seemingly serial code demonstrate simple, yet surprisingly effective rules for detecting potential parallelism.

For each example, learn how to express the code at a higher, more concise level in M by vectorizing computations. We give several canned techniques of vectorization for many common, and sometimes very difficult, use cases.

Learn how such vectorization concisely brings the parallelism of code to the forefront and transforms programs that might have been originally difficult to run on a SIMT device very suitable for execution on the GPU. GPU speedups will be shown utilizing Jacket.

Speaker(s):	Gallagher Pryor (VP of Research, AccelerEyes)
Topic(s):	General Interest

TUESDAY, SEPT 21, 16:00 (50 MINUTES)
ROOM A1

2022 Solving PDEs on Regular Grids with OpenCurrent

OpenCurrent is an open source library with support for structured 3D grids and various PDE solvers that operate on them, including a multigrid Poisson solver and an incompressible Navier-Stokes solver. It also includes extensions for splitting grids across multiple GPUs. This talk will provide a basic introduction to the code base and its design principles.

Speaker(s):	Jonathan Cohen (Senior Research Scientist, NVIDIA Research)
Topic(s):	Computational Fluid Dynamics

TUESDAY, SEPT 21, 16:00 (50 MINUTES)
ROOM A2

2036 Algorithms for Automated Segmentation of Medical Imaging Studies Utilizing CUDA

Discover how GPU computing can help doctors make sense of modern imaging studies. This session is intended for a general audience as well as medical informatics specialists. The focus will be on algorithmic approaches to segmentation as it pertains

to CTA (computed tomography angiography) studies. Topics covered will include specialized optimization algorithms and novel lumen tracking methodologies.

Speaker(s):	Supratik Moulik (University of Pennsylvania)
Topic(s):	Medical Imaging & Visualization, Computer Vision

TUESDAY, SEPT 21, 16:00 (50 MINUTES)
ROOM A3

2052 Power Management Techniques for Heterogeneous Exascale Computing

Power consumption has become the leading design constraint for large scale computing systems. In order to achieve exascale computing, system energy efficiency must be improved significantly. Our approach will focus on investigating software methodologies to achieve energy efficient computing on heterogeneous systems accelerated with GPUs.

Speaker(s):	Xiaohui Cui (Research Scientist, Oak Ridge National Laboratory)
Topic(s):	High Performance Computing

TUESDAY, SEPT 21, 16:00 (50 MINUTES)
ROOM C

2056 Next-Generation Rendering with CgFX

Dive into the details of using CgFX – Cg's effect framework – to combine ray-tracing with real-time rendering and enable the next generation of complex high-quality rendering. You will learn how to use CgFX to create complex rendering effects in a concise and elegant fashion by:

• Blending material-level and scene-level effects in a consistent way,

• Seamlessly integrating CUDA-based data processing within the CgFX rendering pipeline,

• Mixing OptiX-based rendering with CgFX and OpenGL.

Speaker(s):	Tristan Lorach (Computer Graphics Engineer, NVIDIA)
Topic(s):	Computer Graphics

TUESDAY, SEPT 21, 16:00 (50 MINUTES)
ROOM A5

2067 Experiences with Code Optimizations for High Performance GPGPU Programs

Attend this session to learn and share code optimizations to achieve high performance GPU computing. We will cover code transformations for memory coalesing, workload management at both thread and thread-block levels, and different ways to handle memory partition conflicts. We will also discuss Integration of code optimizations into a compiler.

Speaker(s)	Huiyang Zhou (Associate Professor, North Carolina State University), Yi Yang (Ph.D. student, North Carolina State University)
Topic(s):	Programming Languages & Techniques

TUESDAY, SEPT 21, 16:00 (50 MINUTES)
ROOM N

2129 Hardware Subdivision and Tessellation of Catmull-Clark Surfaces

See how the new DirectX 11 Hardware Tessellation and Compute Shader can be used to implement an adaptive Catmull-Clark subdivision surface renderer. We use a table driven approach to performing Catmull-Clark subdivision in parallel utilizing one thread per output mesh vertex.

Speaker(s):	Charles Loop (Senior Researcher, Microsoft Research)
Topic(s):	Computer Graphics

TUESDAY, SEPT 21, 16:00 (50 MINUTES)
ROOM B

2151 Parallel Nsight: Analyzing and Optimizing Massively Parallel Applications [Advanced]

Life altering products that provide early detection of breast cancer or simulate molecular behavior, accelerating drug discovery, are becoming reality thanks to the power of the GPU. As these technologies become mainstream, mainstream tools are required to support these development efforts.

NVIDIA Parallel Nsight delivers the power of the GPU within the familiar environment of Microsoft Visual Studio. In this session, you will learn advanced techniques for visualizing your application's workloads and performance characteristics across the CPU, GPU, and operating system, and explore the depths of Parallel Nsight profilers, including GPU performance counters and how to use them.

Speaker(s):	Sebastien Domine (Sr. Dir. Developer Tools, NVIDIA)
Topic(s):	Tools & Libraries, Programming Languages & Techniques

TUESDAY, SEPT 21, 16:00 (50 MINUTES)
ROOM A7

2161 NVIDIA Quadro Digital Video Pipeline Overview

This session will provide an overview of the Quadro Digital Video Pipeline. It will cover a description of the DVP components, application architectures software architectures, and programming resources available.

Speaker(s):	Thomas True (Applied Engineer, NVIDIA)
Topic(s):	Computer Graphics, Video Processing, Programming Languages & Techniques

TUESDAY, SEPT 21, 16:00 (20 MINUTES)
ROOM K

2179 GPU - An R Library for Native GPU Objects

Come learn about the GPU R package. R is the widely popular open source statistical programming language. The GPU package extends R by providing GPU-based types, classes and methods implementing GPU versions of R vectors, matrices, lists and data frames. Subsequent operations with these are executed on the GPU. Users are not required to create special bindings or implement special syntax, nor do they need copy objects between CPU and GPU. The GPU packages allows programmers access to the computational power of GPUs with little modification to existing code.

Speaker(s):	Christopher Brown (Partner, Open Data)
Topic(s):	Tools & Libraries, Algorithms & Numerical Techniques, High Performance Computing

TUESDAY, SEPT 21, 16:00 (50 MINUTES)
ROOM M

2247 Reconfiguring a Pool of GPUs on The Fly (Sponsored by NextIO)

Today's HPC applications break down large data set problems into smaller, independent elements solved by massively parallel processor systems. GPU's as co processing devices are optimized for this task and their popularity in technical computing is rapidly

advancing. Like many rapidly advancing technologies, they leave in their wake new and challenging problems. In the effort to cut costs while increasing performance, damaging ripple-effects can occur; resources can be over or under provisioned, inventory difficult to manage, lots of single points of failure mean constant job interruptions, manual reconfiguration of resources are required for each job, servicing and lifecycle management require outages. Most of these problems can be addressed and overcome by combining GPU resources into managed, structured pools. NextIO will present and demonstrate a new and innovative approach to consolidating and managing pools of NVIDIA GPU resources along with the cost and operational savings benefits associated with top of rack GPU consolidation appliances.

Companies that consolidate GPU resources in a top of rack appliance can reduce GPU operational costs by $1000 per GPU per year which in most cases pays for the GPU itself. NextIO will present the TCO models for a 500 GPU installation using traditional upgrade methodology vs. a top of rack appliance solution which reduces installation time to 5 minutes per GPU without incurring application downtime on server nodes.

Speaker(s):	K.C. Murphy (Vice President of Marketing, NextIO)
Topic(s):	High Performance Computing, Cloud Computing

TUESDAY, SEPT 21, 16:00 (50 MINUTES)
MARRIOTT SAN JOSE BALLROOM

2264 CUDA Centers of Excellence Super-Session III

Come hear about the groundbreaking research taking place at the CUDA Centers of Excellence, an elite group of world-renown research universities that are pushing the frontier of massively parallel computing using CUDA. Researchers from these top institutions will survey cutting-edge research that is advancing the state of the art in GPU computing and dozens of application fields across science and engineering.

In this session we will hear from Dr. Wen-mei Hwu at the University of Illinois at Urbana – Champaign, Professor Yangdong Deng at Tsinghua University and Dr. Charles D. Hansen at the University of Utah.

Speaker(s):	Wen-mei Hwu (Professor, University of Illinois, Urbana-Champaign), Yangdong Deng (Associate Professor, Tsinghua University), Charles Hansen (Professor, University of Utah)
Topic(s):	General Interest

TUESDAY, SEPT 21, 16:00 (50 MINUTES)
ROOM A8

2274 Harnessing the Power of the GPU in Internet Explorer 9

Internet Explorer 9 is bringing the power of modern GPUs to Web. Thanks to hardware accelerated graphics, the websites that you use every day become faster and developers can create new classes of web applications which were previously not possible. This session will provide an inside look into how Internet Explorer was redesigned to leverage the GPU. We'll show detailed performance results, discuss our architectural approach, and look at the impact of the GPU on HTML5. A session by engineers for engineers with lots of fun demos.

Speaker(s):	Jason Weber (Internet Explorer Performance Lead, Microsoft)
Topic(s):	GPU Accelerated Internet

TUESDAY, SEPT 21, 16:00 (50 MINUTES)
ROOM L

2295 Large-Scale CFD Applications and a Full GPU Implementation of a Weather Prediction Code on the TSUBAME Supercomputer

Many CFD applications have been successfully accelerated on GPUs, but for large-scale simulations that require memory beyond a single GPU, communication is required between GPUs over cluster nodes through PCI-Express and interconnects. To overcome performance bottlenecks and preserve parallel scalability, an overlapping technique between computation and communication is essential. This work presents results of an LBM for incompressible flow, and a Tsunami simulation solving the shallow water equation for simulations on the NVIDIA Tesla-based TSUBAME supercomputer of Tokyo Tech. In addition results will be presented on a complete GPU implementation of a production-level weather prediction code developed by the JMA that achieves 15 TFLOPS for an 80-fold speedup.

Speaker(s):	Takayuki Aoki (Professor, Tokyo Institute of Technology)
Topic(s):	Computational Fluid Dynamics

TUESDAY, SEPT 21, 16:30 (20 MINUTES)
ROOM K

2111 Using R for High-Performance Data Analysis

Data analysis is the art and the science of getting the correct quantitative models and their numerical parameters from the observed data. In this talk, we report on a project to integrate CUDA into the open source data analysis environment R. The combined use of the CPU and GPU resources can efficiently exploit the significant amount of data parallelism inherent in most data analysis problems and methods. This makes interactive analysis possible even for large, compute-intensive problems. The implementation and the achievable performance gains will be demonstrated on a concrete example from quantitative finance.

Speaker(s):	Domokos Vermes (Associate Professor, Worcester Polytechnic Insitute)
Topic(s):	Tools & Libraries, Databases & Data Mining, Finance, Life Sciences

TUESDAY, SEPT 21, 17:00 (50 MINUTES)
ROOM A2

2009 4D Visualization and Analysis of Flow

4D flow or vector data is now common in CFD simulations as well as acquisition techniques like 4D flow MRI to study abnormal blood flow patterns. We show how by mixing compute and graphics combined with stereo we are now able to interactively analyze and visualize the resulting data to understand abnormal flow patterns. Topics include flow field rendering, computing derived quantities, merging volumetric rendering with computed geometry such as particles and surfaces, and integration 3d vision stereo.

Speaker(s):	Shalini Venkataraman (Applied Engineer, NVIDIA)
Topic(s):	Medical Imaging & Visualization, Computational Fluid Dynamics, Stereoscopic 3D

TUESDAY, SEPT 21, 17:00 (20 MINUTES)
ROOM M

2026 MatCloud: Accelerating Matrix Math GPU Operations with SaaS

We present MatCloud (www.mat-cloud.com), a cloud infrastructure and service for scientific computing using state-of-the-art GPU clusters. MatCloud is a service infrastructure exposed by a simple web terminal interface to run Matlab-like commands/scripts. Join us to see how GPU technology can not only be applied to cloud computing community, but also boost the adoption of cloud computing for its dramatic performance gains over traditional cloud infrastructures. MatCloud is an in-progress academic project and is under active development.

Speaker(s):	Xing Wu (Research Assistant, North Carolina State University), Frank Mueller (Associate Professor, North Carolina State University)
Topic(s):	Cloud Computing, Tools & Libraries

TUESDAY, SEPT 21, 17:00 (50 MINUTES)
ROOM A8

2060 GPUs in a Flash: Mapping the Flash Animated Software Vector Rendering Model to the GPU

Explore the Flash rendering architecture including the challenges of mapping from an animated software vector rendering model to a GPU. We will also discuss how the landscape of mobile, desktop, devices, drivers, and APIs impacts the design and deployment of a GPU based Flash Player.

Speaker(s):	Lee Thomason (Principal Scientist, Adobe Systems)
Topic(s):	GPU Accelerated Internet

TUESDAY, SEPT 21, 17:00 (50 MINUTES)
ROOM A5

2084 State of the Art in GPU Data-Parallel Algorithm Primitives

Learn about the importance of optimized data-parallel algorithm primitives as building blocks for efficient real-world applications. Fundamental parallel algorithms like sorting, parallel reduction, and parallel scan are key components in a wide range of applications from video games to serious science. This session will cover the state of the art in data-parallel primitive algorithms for GPUs. Starting with an explanation of the purpose and applications of the algorithms, we will discuss key algorithm design principles, demonstrate current open source algorithm libraries for GPUs (CUDPP and Thrust), describe optimizations using new features in the Fermi architecture, and explore future directions.

Speaker(s):	Mark Harris (Senior Developer Technology Engineer, NVIDIA)
Topic(s):	Algorithms & Numerical Techniques, High Performance Computing, Tools & Libraries

TUESDAY, SEPT 21, 17:00 (50 MINUTES)
ROOM N

2102 Evacuate Now? Faster-than-real-time Shallow Water Simulation on GPUs

Learn how to simulate a half an hour dam break in 27 seconds! We present how shallow water simulation with interactive visualization is successfully mapped to modern graphics hardware. Featuring a live demo, we will present interactive shallow water simulations running on a standard laptop. The implementation has been verified against analytical and experimental data, supports multi-gpu simulation, and can run up-to 6300x6300 domain sizes at 320 million cells per second on the GTX 480.

Speaker(s):	André Rigland Brodtkorb (Research Scientist, SINTEF ICT)
Topic(s):	Physics Simulation, Computational Fluid Dynamics

TUESDAY, SEPT 21, 17:00 (50 MINUTES)
ROOM A7

2205 A Highly Reliable RAID System Based on GPUs

While RAID is the prevailing method of creating reliable secondary storage infrastructure, many users desire more flexibility than offered by current implementations. To attain needed performance, customers have often sought after hardware-based RAID solutions. This talk describes a RAID system that offloads erasure correction coding calculations to GPUs, allowing increased reliability by supporting new RAID levels while maintaining high performance.

Speaker(s):	Matthew Curry , Sandia National Laboratories and the University of Alabama at Birmingham
Topic(s):	High Performance Computing

TUESDAY, SEPT 21, 17:00 (50 MINUTES)
ROOM B

2212 Parallel Nsight for Accelerated DirectX 11 Development [Advanced]

Parallel Nsight is NVIDIA's new development environment for graphics and GPU computing. In this advanced session, you will learn how Parallel Nsight can accelerate debugging and profiling of Direct3D 11 applications. Attendees will learn how to debug Direct3D frames and HLSL shaders using Parallel Nsight's powerful Graphics Inspector and Debugger which allows developers to inspect Direct3D resources and state, set breakpoints in HLSL shaders, examine shader variables, and see which graphics primitives are live on the GPU. Attendees will also learn how to use the Frame Profiler to capture and mine performance information, and easily pinpoint bottlenecked GPU units.

Speaker(s):	Simon Barrett (Senior Software Engineer, NVIDIA)
Topic(s):	Programming Languages & Techniques, Computer Graphics

TUESDAY, SEPT 21, 17:00 (50 MINUTES)
ROOM A3

2225 Tools for Managing Clusters of NVIDIA GPUs

Learn about the suite of tools NVIDIA provides to manage large installations of GPUs from the NVIDIA Tesla Series. The presentation will cover cluster management – tool and library–, as well as the GPUDirect technology that enables GPUs to communicate faster across the network.

Speaker(s):	Peter Buckingham (Tesla Software Manager, NVIDIA), Andrew Iles (Software Engineer, NVIDIA)
Topic(s):	Tools & Libraries

TUESDAY, SEPT 21, 17:00 (50 MINUTES)
ROOM L

2239 Fast GPU Preconditioning for Fluid Simulations in Film Production

Explore how a less efficient, but highly parallel algorithm can still be a superior alternative to a sequential CPU method. This talk will present a simple CUDA-based Poisson solver to the conjugate gradient method designed for solving well-conditioned matrices such as those that arise from the pressure projection stage of a Navier-Stokes fluid solver. In contrast to other active areas of research in this field, we show that a more brute force approach can still significantly out-perform the best CPU alternatives by sacrificing a high convergence rate in place of achieving much faster iterations.

Speaker(s):	Dan Bailey (R&D, Double Negative)
Topic(s):	Computational Fluid Dynamics, Algorithms & Numerical Techniques, Film

TUESDAY, SEPT 21, 17:00 (50 MINUTES)
ROOM C

2250 GPU Ray Tracing Exposed: Under the Hood of the NVIDIA OptiX Ray Tracing Engine

Take a deep dive into many of the design choices and implementation details of the NVIDIA OptiX ray tracing engine. Learn how domain specific compilation, a unique execution model and a general object model, are combined into a flexible and powerful API.

Speaker(s):	Austin Robison (Lead Developer, OptiX Integration, NVIDIA)
Topic(s):	Ray Tracing

TUESDAY, SEPT 21, 17:00 (50 MINUTES)
MARRIOTT SAN JOSE BALLROOM

2265 CUDA Centers of Excellence Super-Session IV

Come hear about the groundbreaking research taking place at the CUDA Centers of Excellence, an elite group of world-renown research universities that are pushing the frontier of massively parallel computing using CUDA. Researchers from these top institutions will survey cutting-edge research that is advancing the state of the art in GPU computing and dozens of application fields across science and engineering.

In this session we will hear from Professor Ting-wai Chiu at National Taiwan University, Dr. Satoshi Matsuoka at Tokyo Tech and Dr. Paul Calleja at the University of Cambridge.

Speaker(s):	Paul Calleja (University of Cambridge), Ting-Wai Chiu (Professor, National Taiwan University), Satoshi Matsuoka (Professor, Tokyo Institute of Technology)
Topic(s):	General Interest

TUESDAY, SEPT 21, 17:00 (50 MINUTES)
ROOM K

2297 Developing CUDA Accelerated .NET Plugins for Microsoft Excel

Quantifi will demo its xLDevelopment environment, which provide developers with an easy to use development environment which allows cuda functionality to be in Microsoft Excel. With as little as four lines, one will also select the position of the function in the menu bar, xml markup language will display in the excel help functionality, and objects can be easily added to the object cache. These objects can then be inspected by the end user or developer. Performance information can also be displayed in the object cache. The environment provides the developer an environment where he can focus on developing high performance functionality, and all intermediate layers of interface are taking care of by the environment.

Speaker(s):	Peter Decrem (Director, Rates Products, Quantifi)
Topic(s):	Tools & Libraries, Finance

TUESDAY, SEPT 21, 17:00 (50 MINUTES)
ROOM D

2304 Harnessing the GPU to Accelerate Automotive Development

Learn how GPU technologies broke speed limits in automotive development. By using GPU-accelerated tools, small team of engineers created a complete certifiable vehicle in only two years, using fraction of the budget used in conventional industry. Talk will cover tools and techniques used in creation of XD concept, as well as how to overcome challenges moving a product from concept to mass production stage.

Speaker(s):	Theo Valich (President, Bright Side Network Inc.)
Topic(s):	Embedded & Automotive, Computer Vision, Computer Graphics

TUESDAY, SEPT 21, 18:00 (120 MINUTES)
KEYNOTE HALL

1005 Research Posters Showcase / Exhibits Open / Networking Reception

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 research posters to see "what's next" out of the world of research and academia, and meet the poster presenters, who will be stationed near their posters during this time. Appetizers and drinks will be served.

Topic(s):	General Interest

GE
Intelligent Platforms

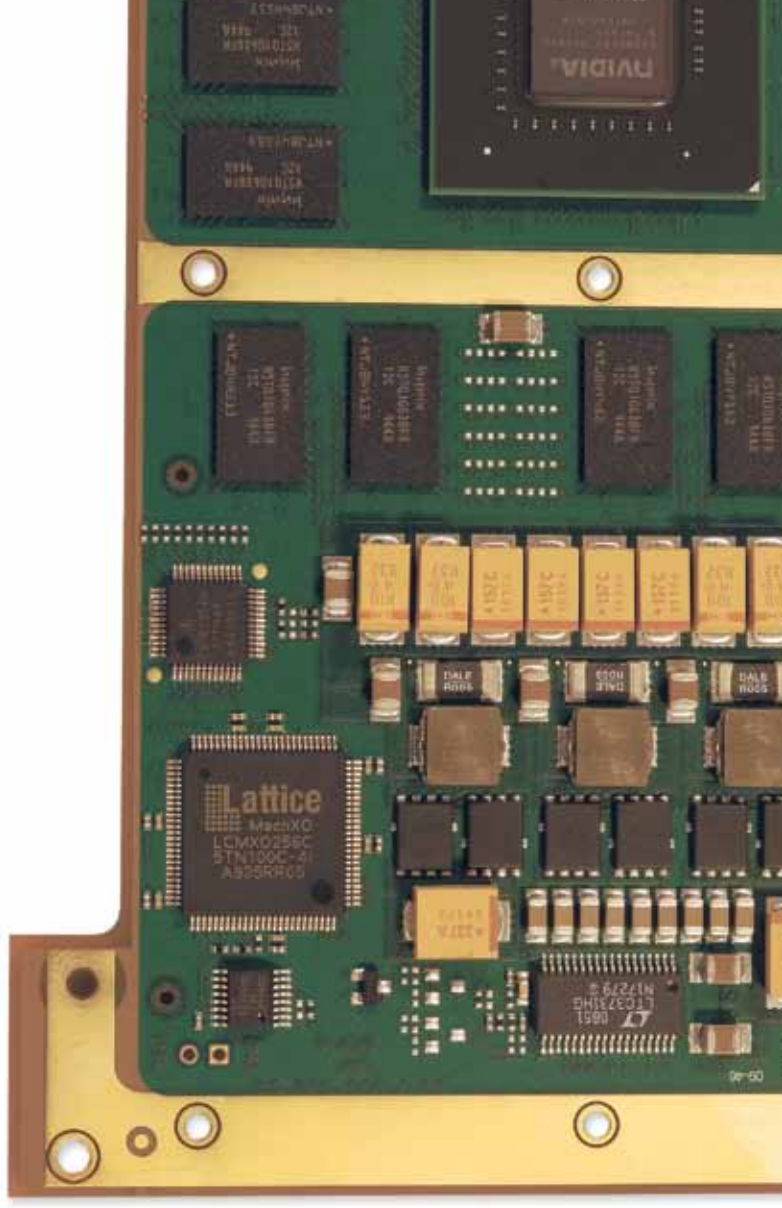

1 Rugged Cluster. 290 Processors.

We've redefined the size, weight and power equation for rugged applications.

Thanks to our partnership with NVIDIA®, you can now take advantage of massively parallel GPGPU processing in rugged applications such as radar, sonar, signal intelligence, signal processing, image processing and software defined radio. Indeed, any rugged application that involves parallel execution of multiple threads will benefit. Best of all, these performance improvements are accessible to almost everyone because the CUDA™ architecture allows you to program in C.

We offer a number of NVIDIA-based GPGPU products in various rugged form factors, such as the 6U VPX IPN250, which is CUDA-capable and supports OpenCL™. And our AXISLib VSIPL DSP libraries will help you develop high performance multi-core and multi-processor applications without needing to understand the complexities of the underlying hardware. You can also easily migrate AXIS applications to new platforms. For more information on our growing family of CUDA-capable GPGPU boards, please visit **www.ge-ip.com/gpgpu**

imagination at work

WEDNESDAY, SEPT 22, 09:00 (50 MINUTES)
KEYNOTE HALL

1002 Day 2 Keynote with Dr. Klaus Schulten, University of Illinois, Urbana-Champaign

How does the H1N1 "Swine Flu" virus avoid drugs while attacking our cells? What can we learn about solar energy by studying biological photosynthesis? How do our cells read the genetic code? What comes next in computational biology?

Computational biology is approaching a new and exciting frontier: the ability to simulate structures and processes in living cells. Come learn about the "computational microscope," a new research instrument that scientists can use to simulate biomolecules at nearly infinite resolution. The computational microscope complements the most advanced physical microscopes to guide today's biomedical research. In this keynote address, computational biology pioneer Dr. Klaus Schulten of the University of Illinois, Urbana-Champaign, will introduce the computational microscope, showcase the widely used software underlying it, and highlight major discoveries made with the aid of the computational microscope ranging from viewing protein folding, translating the genetic code in cells, and harvesting solar energy in photosynthesis. He will also look towards a future when cell tomography and computing will establish atom-by-atom views of entire life forms.

Klaus Schulten received his Ph.D. from Harvard University in 1974. He is Swanlund Professor of Physics and is also affiliated with the Department of Chemistry as well as with the Center for Biophysics and Computational Biology. Professor Schulten is a full-time faculty member in the Beckman Institute and directs the Theoretical and Computational Biophysics Group at the University of Illinois Urbana-Champaign, IL. Honors and awards: Award in Computational Biology 2008; Humboldt Award of the German Humboldt Foundation (2004); University of Illinois Scholar (1996); Fellow of the American Physical Society (1993); Nernst Prize of the Physical Chemistry Society of Germany (1981).

Topic(s):	General Interest, Life Sciences
Speaker(s):	Klaus Schulten (Professor, University of Illinois, Urbana-Champaign)

WEDNESDAY, SEPT 22, 10:00 (50 MINUTES)
ROOM C

2058 A Practical Introduction to Computational Fluid Dynamics on GPUs

Learn step-by-step procedures to write an explicit CFD solver based on final difference methods with staggered grid allocations and boundary fitted coordinates. We will discuss the derivation of the mathematical model, discretization of the model equations, development of the algorithms, and parallelization and visualization of the computed data using OpenCL and OpenGL. Compares case studies of natural convection, driven cavity, scaling analysis, and magneto-thermal convection computed using CSIRO's CPU/GPU supercomputer cluster to known analytical and experimental solutions.

Speaker(s):	Tomasz Bednarz (3d Visualisation Software Engineer, CSIRO)
Topic(s):	Computational Fluid Dynamics, Algorithms & Numerical Techniques, High Performance Computing, Physics Simulation

WEDNESDAY, SEPT 22, 10:00 (50 MINUTES)
MARRIOTT GUADALUPE ROOM

2082 CU-LSP: GPU-based Spectral Analysis of Unevenly Sampled Data

Standard FFT algorithms cannot be applied to spectral analysis of unevenly sampled data. Alternative approaches scale as O(N^2), making them an ideal target for harnessing the raw computing power of GPUs. To this end, I have developed CU-LSP, a CUDA spectral analysis code based on the Lomb-Scargle periodogram. Preliminary benchmarking indicates impressive speed-ups, on the order of 400 relative to a single core of a modern CPU. An initial application of CU-LSP will be the analysis of time-series data from planet-search and asteroseismology satellites.

Speaker(s):	Richard Townsend (Assistant Professor, University of Wisconsin-Madison)
Topic(s):	Astronomy & Astrophysics, Algorithms & Numerical Techniques, Signal processing

WEDNESDAY, SEPT 22, 10:00 (50 MINUTES)
ROOM A1

2134 Ultra High Resolution Displays and Interactive Eyepoint Using CUDA

We'll go over the challenges we have overcome in building 100 million pixel seamless displays. One customer requirement involves interactive changes of the eyepoint as a person moves, relative to the screen, yet the distortions computed are quite non-linear. We discuss our use of a gpu to implement this procedure.

Speaker(s):	Rajeev Surati (President, Scalable Display Technologies)
Topic(s):	Computer Graphics, High Performance Computing, Medical Imaging & Visualization

WEDNESDAY, SEPT 22, 10:00 (50 MINUTES)
ROOM A2

2141 Moving the Frontier of Oil and Gas Exploration and Production with GPUs

Learn how the Oil and Gas Industry is embracing GPUs in order to tackle new and complex oil and gas plays around the world. The first part of this talk gives an overview of the business and geopolitical drivers of the industry, followed with the critical contribution of computation in the quest for secure supply of energy.

Speaker(s):	Maurice Nessim (Schlumberger), Shashi Menon (R&D Manager, Schlumberger)
Topic(s):	Energy Exploration, High Performance Computing

WEDNESDAY, SEPT 22, 10:00 (50 MINUTES)
ROOM L

2160 StarPU: a Runtime System for Scheduling Tasks

See how StarPU provides task scheduling facilities for a hybrid platform and a powerful data management library that transparently takes care of data across the entire machine. We will discuss the significant performance improvements resulting from its flexible scheduler as well as its ability to mix parallel CPU kernels (eg. written in OpenMP or TBB) with CUDA/OpenCL and MPI.

Speaker(s):	Cedric Augonnet (INRIA)
Topic(s):	Tools & Libraries, High Performance Computing

WEDNESDAY, SEPT 22, 10:00 (20 MINUTES)
ROOM D

2163 Leveraging GPUs for Evolutionary Game Theory

Learn how GPUs are being used to accelerate the study of the emergence of cooperative behavior in biology, from the interactions of humans to viruses to bacteria. The work presented here achieves a speedup of 209x on a cluster of 4 Tesla GPUs.

Speaker(s):	Amanda Peters (PhD Candidate, Harvard University)
Topic(s):	Algorithms & Numerical Techniques, Life Sciences

WEDNESDAY, SEPT 22, 10:00 (50 MINUTES)
ROOM A3

2166 The Triad of Extreme Computing-Fast Algorithms, Open Software and Heterogeneous Systems

The first wave of successful GPU accelerations has been crowded with highly-parallel methods that adapted well to the hardware. But the easy-pickings are now running out. The truly challenging applications require "going back to the algorithmic drawing board." To develop new versions of the most effective fast algorithms, such that our science can most benefit, an ideal environment is created by the open software model, where efforts can be shared. We will describe one area of application --electrostatics of biomolecules in solution-- where we see at work the triad of extreme computing: fast algorithms, open software, and heterogeneous computing.

Speaker(s):	Lorena Barba (Assistant Professor, Boston University)
Topic(s):	Algorithms & Numerical Techniques, Physics Simulation

WEDNESDAY, SEPT 22, 10:00 (50 MINUTES)
ROOM B

2168 Interactive Molecular Dynamics for Nanomechanical and Nanochemical Experiments

Hear how the combination of GPU accelerated molecular dynamics simulation software, 3D TV displays, affordable haptic game controllers, and high performance molecular visualization is leading to new ways to study materials and objects on the nanoscale. We will present the concept of an appliance for integrated virtual nanoscale experiments and challenges related to software and hardware.

Speaker(s):	Axel Kohlmeyer (Associate Director, Institute for Computational Molecular Science, Temple University)
Topic(s):	Molecular Dynamics

WEDNESDAY, SEPT 22, 10:00 (50 MINUTES)
ROOM A7

2169 Real-time Volumetric Medical Ultrasound Applications for GPU Computing

Real-time volumetric medical ultrasound requires computationaly intensive rapid processing of data for visualization of aquired acoustic data. Clinical applications of GPU-based technologies in obstetrics and cardiology will be discussed.

Speaker(s):	Roee Lazebnik (Director of Product Development, Siemens Healthcare)
Topic(s):	Medical Imaging & Visualization, Imaging, Stereoscopic 3D, Computer Graphics

WEDNESDAY, SEPT 22, 10:00 (50 MINUTES)
ROOM E

2231 Driving on Mars, Redux: System Level Simulation of Dynamic Systems

Learn how GPU and HPC computing are used to predict through simulation the dynamics of large complex mechanical systems such as tracked vehicles including the Mars Rover. The presentation outlines the physics based approach and numerical solution methods that enabled the simulation of dynamic systems with millions of bodies on the GPU. The presentation will also explain how a HPC cluster is used to effectively render scenes with tens of thousands of bodies for generating animations that can be used by Engineers in the design process.

Speaker(s):	Dan Negrut (Assistant Professor, University of Wisconsin)
Topic(s):	Physics Simulation, Algorithms & Numerical Techniques, High Performance Computing

WEDNESDAY, SEPT 22, 10:00 (50 MINUTES)
ROOM A5

2249 New Programming Tools GPU Computing

This session will focus on new parallel programming tools for GPU computing. The type of tools that fit into the session include (1) Planning tools for porting legacy applications to use GPU computing, (2) High-level programming and scripting tools for GPU computing, (3) Automation of common performance optimizations for GPU computing, (4) Performance analysis and diagnosis tools for GPU computing, (5) Tools that simplify heterogeneous parallel computing.

Speaker(s):	Wen-mei Hwu (Professor, University of Illinois, Urbana-Champaign), Andrew Schuh (Project Manager, University of Illinois)
Topic(s):	Tools & Libraries

WEDNESDAY, SEPT 22, 10:00 (50 MINUTES)
MARRIOTT SAN JOSE BALLROOM

2280 TSUBAME2.0 Experience

Tsubame2.0 is the next-generation multi-petaflops supercomputer that been designed and built at Tokyo Tech, with more than 4000 NVIDIA Fermi GPUs. as a successor to the highly successful Tsubame1. Deep design considerations were made based on experiences on Tsubame1 retrofitted with the previous generation Tesla to maximize the versatility and the competitiveness of the system across considerable number of application domains, as well as accommodating as much strong scaling as possible. This resulted in a totally new custom system design in collaboration with HP and NEC, rather than a machine with a retrofitted GPUs. The resulting supercomputer hopefully will become a design template of future large-scale GPU systems to come.

Speaker(s):	Satoshi Matsuoka (Professor, Tokyo Institute of Technology)
Topic(s):	High Performance Computing

WEDNESDAY, SEPT 22, 10:00 (50 MINUTES)
ROOM A8

2305 PantaRay: Accelerating Out-Of-Core Ray Tracing of Sparsely Sampled Occlusion

Modern VFX rendering pipelines are faced with major complexity challenges: a film like Avatar requires rendering hundreds of thousands of frames, each containing hundreds of millions

or billions of polygons. Furthermore, the process of lighting requires many rendering iterations across all shots. In this talk, we present the architecture of an efficient out-of-core ray tracing system designed to make rendering precomputations of gigantic assets practical on GPUs. The system we describe, dubbed PantaRay, leverages the development of modern ray tracing algorithms for massively parallel GPU architectures and combines them with new out-of-core streaming and level of detail rendering techniques.

Speaker(s):	Luca Fascione (Senior Research and Development Engineer, Weta Digital)
Topic(s):	Digital Content Creation (DCC)

WEDNESDAY, SEPT 22, 10:00 (20 MINUTES)
ROOM K

2306 Gate-Level Simulation with GP-GPUs

Logic simulation is a critical component of the digital design tool flow. It is used from high-level descriptions down to gate-level to validate several aspects of the design, particularly functional correctness. Despite development houses investing vast resources in the simulation task it is still far from achieving the performance demands of validating complex modern designs at gate-level. We developed a GP-GPU accelerated gate-level simulator using NVIDIA CUDA.

We leverage novel algorithms for circuit netlist partitioning and found that our experimental prototype could handle large, industrial scale designs comprised of millions of gates while delivering 13x speedup on average over a typical commercial simulator.

Speaker(s):	Debapriya Chatterjee (Graduate Student, University of Michigan)
Topic(s):	General Interest

WEDNESDAY, SEPT 22, 10:00 (50 MINUTES)
ROOM N

2308 Building Cutting-Edge Realtime 3D Applications with NVIDIA SceniX

Learn how NVIDIA SceniX is a rapid start to building state of the art, realtime 3D applications, and how raytracing can be combined with raster graphics for new levels of interactive realism.

Speaker(s):	Brian Harrison (NVIDIA), Michael Morrison (NVIDIA)
Topic(s):	Computer Graphics, Computer Vision, Ray Tracing, Stereoscopic 3D

WEDNESDAY, SEPT 22, 10:00 (50 MINUTES)
KEYNOTE HALL

4000 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 compute-intensive problems. The Opening Address includes an overview of NVIDIA's GPU ecosystem development activities and an interaction on stage with selected companies building groundbreaking applications on top of the GPU platform.

The ECS is a great opportunity to discover new players in the GPU ecosystem, find great investments, explore partnership and customer/vendor opportunities, network/build relationships, and discuss the future of an industry that is reshaping computing.

Speaker(s):	Jeff Herbst (Vice President of Business Development, NVIDIA)
Topic(s):	General Interest

WEDNESDAY, SEPT 22, 10:30 (20 MINUTES)
ROOM D

2109 Migration of a Complete 3D Poisson Solver from Legacy Fortran to CUDA

We describe our journey of migrating a legacy direct solver library for Poisson equations written in Fortran77 to CUDA in order to harness the computational power provided by the Tesla device ("Fermi"). This legacy library is still widely used today as it is the most complete library that can deal with three different boundary conditions (Dirchlet, Neumann and Cyclic) and two grid configurations (staggered and centered) independently in any of the three dimensions (x, y, z); giving a total of over 200 configurations.

Speaker(s):	Huynh Phung (Research Engineer, A*STAR Institute of High Performance Computing)
Topic(s):	Tools & Libraries, Computational Fluid Dynamics

WEDNESDAY, SEPT 22, 10:30 (20 MINUTES)
ROOM K

2300 High-Performance Compressive Sensing using Jacket

This talk will present the ongoing work that I am doing in the L1-optimization group at Rice Universtiy. The purpose of the work is to merge both compressive sensing, for image/signal reconstructions and GPU computation, using NVIDIA's GPUs to enhance the technology of CS.

This talk will cover basic concepts in compressive sensing and the easy adaptation of operating on the GPU, in particular working with Jacket (by AccelerEyes). We willthen cover some of our numerical experiments that encompass the use of different flavors of algorithms.

Speaker(s):	Nabor Reyna Jr. (Graduate Student, Rice University)
Topic(s):	Imaging, Tools & Libraries

WEDNESDAY, SEPT 22, 11:00 (20 MINUTES)
ROOM B

2034 Reformulating Algorithms for the GPU

Important applications in signal, data processing and bioinformatics that use dynamic programming are difficult to parallelize due to intrinsic data dependencies. We demonstrate a novel technique to extract parallelism out of data dependent algorithms and reformulate the same for GPUs.

This simple technique breaks the dependencies and resolves them at an optimal point later in time, thus obtaining remarkable speedup on GPUs. We present a case study from computational biology i.e., protein motif-finding. We also present how the same technique can be extended and applied to other relevant problems such as gene-prediction and phylogenetics.

Speaker(s):	Narayan Ganesan (Research Scientist, University of Delaware), Michela Taufer (Assistant Professor, University of Delaware)
Topic(s):	Life Sciences, Algorithms & Numerical Techniques, High Performance Computing

WEDNESDAY, SEPT 22, 11:00 (20 MINUTES)
ROOM K

2039 GPU Debugging with Allinea DDT

Discover how a debugger can help you fix those hard to find bugs in your GPU software, with this introduction to the special CUDA features in Allinea DDT.

Speaker(s):	David Lecomber (CTO, Allinea Software)
Topic(s):	Tools & Libraries

WEDNESDAY, SEPT 22, 11:00 (50 MINUTES)
ROOM A2

2059 Industrial Seismic Imaging on GPUs

At Hess Corporation, we have moved the most computationally intensive parts of our seismic imaging codes from CPUs to GPUs over the past few years. In this talk I will give an overview of seismic imaging, highlighting the physical and computational algorithms of these codes. I will discuss our software approach and the programming effort to port them to GPUs, concluding with a summary of our progress in adopting GPUs in production.

Speaker(s):	Scott Morton (Geophysical Advisor, Hess Corporation)
Topic(s):	Energy Exploration, High Performance Computing

WEDNESDAY, SEPT 22, 11:00 (50 MINUTES)
ROOM E

2065 Massively Accelerating Iterative Gauss-Newton Fitting

To measure three-dimensional shape data of objects, we build up a measurement system that assigns three-dimensional coordinates to the position of projected measurement labels in a camera image. To achieve high measurement accuracy across high amounts of measurement points, we need a very quick routine to localize measurement labels with high precision. To speed up the computation, we evaluate the fits using the CUDA architecture. The final implementation speeds up the fitting of 104 two-dimensional Gauss functions by a factor of 90.

Speaker(s):	Daniel Härter (University of Freiburg, IMTEK, Laboratory for Process Technology)
Topic(s):	Computer Vision, Stereoscopic 3D

WEDNESDAY, SEPT 22, 11:00 (50 MINUTES)
ROOM A1

2071 Large Scale Visualization Soup

The unprecedented realism that is possible today allows for visualization at an ever larger scale. This talk will walk through several case studies from high resolution single displays to completely immersive environments. Details will be shared on how to architect and implement these installations, with attention to the typical issues encountered. It will cover how to implement stereo 3D in OpenGL, Direct3D, as well as how that relates to the different display technologies (projectors, multi-display, CAVEs, etc.).

Speaker(s):	Steve Nash (Applied Engineer, NVIDIA)
Topic(s):	Computer Graphics, Stereoscopic 3D

WEDNESDAY, SEPT 22, 11:00 (50 MINUTES)
MARRIOTT SAN JOSE BALLROOM

2078 Shockingly fast and accurate CFD simulations

In the last three years we have demonstrated how GPU accelerated discontinuous Galerkin methods have enabled simulation of time-dependent, electromagnetic scattering from airplanes and helicopters.

In this talk we will discuss how we have extended these techniques to enable GPU accelerated simulation of supersonic airflow as well.

Speaker(s):	Timothy Warburton (Associate Professor, Rice University)
Topic(s):	Computational Fluid Dynamics, Algorithms & Numerical Techniques, High Performance Computing

WEDNESDAY, SEPT 22, 11:00 (50 MINUTES)
ROOM L

2092 Integrating CUDA into a Large-Scale Commercial Database Management System

In a large-scale database installation where data tables are distributed across multiple servers, computational throughput can be optimized by using GPUs on each server and integrating database management with GPU resources.

In the Department of Physics and Astronomy at The Johns Hopkins University, we are experimenting with a set of software tools that closely couple SQL statements with GPU functionality. While still under development, the new framework is now routinely used in our research projects, e.g., to study the spatial clustering of galaxies as well as genomics.

Speaker(s):	Richard Wilton (Research Scientist, The Johns Hopkins University), Tamas Budavari (Research Scientist, Johns Hopkins University)
Topic(s):	Databases & Data Mining, Astronomy & Astrophysics, High Performance Computing, Tools & Libraries

WEDNESDAY, SEPT 22, 11:00 (50 MINUTES)
MARRIOTT GUADALUPE ROOM

2099 Cosmology Powered by GPUs Redux

Cosmological simulations aim at reproducing the physical processes which occur on the largest scales of the Universe since the Big-Bang by means of numerical calculations on supercomputers. Using CUDA, I have implemented standard cosmological techniques on GPU architecture (PM N-Body solver, Hydrodynamics & moment-based radiative transfer) and designed them to run on supercomputing facilities by means of MPI+CUDA mixed programming. These applications are able to run on 100 or more graphics devices with typical scalar x50 accelerations and with a communication overhead limited to 15%. It allow to explore physical regimes which were out of reach of current simulations.

Speaker(s):	Dominique Aubert (Lecturer, Strasbourg University)
Topic(s):	Astronomy & Astrophysics

WEDNESDAY, SEPT 22, 11:00 (50 MINUTES)
ROOM N

2104 Rapid Prototyping Using Thrust: Saving Lives with High Performance Dosimetry

Radiation poisoning is an everpresent danger for intervention teams that must visit nuclear sites. Virtual reality can help teams prepare for intervention, but efficient computation of radiation dosage is critical to study complex scenarios. Radiation protection research often uses codes based on the straight line attenuation method. As with other approaches, geometrical computations (finding all the interactions radiation rays/objects intersection) remain the simulation bottleneck. This talk will describe how we have used the Thrust high-level library for CUDA C/C++ to quickly prototype innovative algorithms and achieve a significant speed up.

Speaker(s):	Guillaume Saupin (CEA)
Topic(s):	High Performance Computing, Algorithms & Numerical Techniques, Physics Simulation, Ray Tracing

WEDNESDAY, SEPT 22, 11:00 (50 MINUTES)
ROOM A7

2146 Virtual Surgery

Come see how 3D Vision technology is used in Virtual Surgery Training for Medical Education. BioDigital Systems in conjuncture with University of California San Francisco (UCSF), has developed a dental injection simulator to teach students of dentistry the mechanics of nerve block injection. 3D Vision Technology has added a new dimension of realism by providing users with a unique immersive experience.

Speaker(s):	Aaron Oliker (Managing Partner/Director 3D Technology, BioDigital)
Topic(s):	Medical Imaging & Visualization, Stereoscopic 3D

WEDNESDAY, SEPT 22, 11:00 (50 MINUTES)
ROOM C

2177 Simplifying Parallel Programming with Domain Specific Languages

Explore a new approach in parallel programming which leverages Domain Specific Languages (DSLs) to simplify programming heterogeneous systems (multi-core processors and GPUs). This approach allows DSL users to take advantage of the power of GPUs without having working knowledge of lower level programming models such as CUDA. Topics will cover the advantages of the DSL approach in parallel programming, and the runtime implementation details with optimizations to have the performance benefits of using GPUs.

Speaker(s):	HyoukJoong Lee (PhD Student, Stanford University), Hassan Chafi (PhD Candidate, Stanford University)
Topic(s):	Tools & Libraries, High Performance Computing

WEDNESDAY, SEPT 22, 11:00 (50 MINUTES)
ROOM D

2207 Playing Zero-Sum Games on the GPU

A Zero-Sum game is a match for which the gain of one results in loss of the other. Tic-Tac-Toe, Checkers and Chess are Zero-Sum board game examples. For realizing the best player move, the game is abstracted as a tree, often quite deep, consisting of all possible configurations. We present an efficient GPU implementation of the Mini-Max search algorithm, enhanced with Alpha-Beta pruning. We highlight challenges for deploying non-tail recursion of a highly irregular algorithm on GPUs, proposing a hybrid of compiler and user managed stack. We demonstrate superior performance for running many thousands of 3D Tic-Tac-Toe matches, simultaneously.

Speaker(s):	Avi Bleiweiss (Principal Architect, NVIDIA Corporation)
Topic(s):	Machine Learning & Artificial Intelligence

WEDNESDAY, SEPT 22, 11:00 (50 MINUTES)
ROOM A3

2216 CUDA Libraries Open House

Learn about NVIDIA's CUDA libraries and meet the engineers that develop them. Lead developers will cover the capabilities, performance and future directions for NVIDIA's CUFFT, CUBLAS, CURAND, and NPP libraries (other libraries such as CUSPARSE and open source Thrust are covered in other talks). After the presentation, NVIDIA developers will remain in the room to chat and answer questions during the lunch break.

Speaker(s):	Ujval Kapasi (CUDA Platform SW, NVIDIA), Philippe Vandermersch (Senior Software Engineer, NVIDIA), Elif Albuz (NVIDIA), Nathan Whitehead (CUDA Software Engineer, NVIDIA), Frank Jargstorff (Software Engineer, NVIDIA)
Topic(s):	Tools & Libraries

WEDNESDAY, SEPT 22, 11:00 (50 MINUTES)
ROOM A5

2275 The Evolution of GPUs for General Purpose Computing

Learn how the GPU evolved from its humble beginning as a "VGA Accelerator" to become a massively parallel general purpose accelerator for heterogeneous computing systems. This talk will focus on significant milestones in GPU hardware architecture and software programming models, covering several key concepts that demonstrate why advances in GPU parallel processing performance and power efficiency will continue to outpace CPUs.

Speaker(s):	Ian Buck (Software Director of GPU Computing, NVIDIA)
Topic(s):	General Interest

WEDNESDAY, SEPT 22, 11:00 (50 MINUTES)
ROOM A8

2286 Towards Peta-Scale Green Computation - Applications of the GPU Supercomputers in the Chinese Academy of Sciences (CAS)

China now holds three spots in the June 2010 Top500 list of GPU-based supercomputers, and two of them, using NVIDIA GPUs, are related to CAS. Efficient use of these systems is more important than peak or Linpack performance. This session will cover some of the large-scale multi-GPU applications in CAS, ranging from molecular dynamics below nano-scale to complex flows on meter-scale and porous media on geological scales, as well as fundamental linear algebra and data/image analysis. The idea of keeping high-efficiency and generality of the computation platform by maintaining a consistency among the target physical system, the computational model and algorithm, and the computer hardware will be explained in detail and demonstrated through a number of super-computing applications in the chemical, oil, mining, metallurgical and biological industries.

Speaker(s):	Wei Ge (Professor, Institute of Process Engineering, Chinese Academy of Sciences), Xiaowei Wang (Dr., Institute of Process Engineering), Yunquan Zhang (Professor, Institute of Software, CAS), Long Wang (Associate Professor, Super Computing Center, Institute)
Topic(s):	High Performance Computing

WEDNESDAY, SEPT 22, 11:00 (50 MINUTES)
ROOM M

2293 Scaling Up and Scaling Out GPUs with Supermicro's Twin™ Architecture (Sponsored by Supermicro)

Find out how Supermicro scales up and scales out GPU performance by using Twin™ architecture. In this session, we outline Supermicro's Twin™ architecture advantages across 1U/2U GPU servers and the design of personal supercomputer, and how we are able to scale and optimize GPU technology for datacenter environment and for professional workstation.

Speaker(s):	Don Clegg (Supermicro)
Topic(s):	High Performance Computing, Computer Vision

WEDNESDAY, SEPT 22, 11:00 (50 MINUTES)
KEYNOTE HALL

4001 Emerging Companies: CEO on Stage featuring Elemental Technologies, Geomerics, and Milabra

See the hottest new technologies from startups that are transforming computing.

In a lively and fast-paced exchange, the "Emerging Companies Summit - CEO on Stage" sessions will feature CEOs from three startups who will each have 15 minutes to introduce their companies and interact with a panel of leading venture capitalists, technology executives, and industry analysts.

Panelist(s):	Drew Lanza (Partner, Morgenthaler), Dan'l Lewin (Corporate VP and Strategic & Emerging Business Development, Microsoft), Jon Peddie (President, JPR), Jeff Herbst (Vice President of Business Development, NVIDIA)
Speaker(s):	Sam Cox (CEO, Milabra), Sam Blackman (CEO and Co-Founder, Elemental Technologies, Inc.), Chris Doran (Founder and Chief Operating Officer, Geomerics)
Topic(s):	General Interest, Video Processing, Computer Graphics

WEDNESDAY, SEPT 22, 11:30 (20 MINUTES)
ROOM B

2035 Simulations of Large Membrane Regions

Learn how to study membrane-bound protein receptors by moving beyond the current state-of-the-art simulations that only consider small patches of physiological membranes. Towards this end, this session presents how to apply large-scale GPU-enabled computations of extended phospholipid bilayer membranes using a GPU code based on the CHARMM force field for MD simulations. Our code enables fast simulations of large membrane regions in NVT and NVE ensembles and includes different methods for the representation of the electrostatic interactions, i.e., reaction force field and Ewald summation (PME) methods. Performance and scientific results for dimyristoylphosphatidylcholine (PC) based lipid bilayers are presented.

Speaker(s):	Michela Taufer (Assistant Professor, University of Delaware), Narayan Ganesan (Research Scientist, University of Delaware), Sandeep Patel (Assistant Professor, University of Delaware)
Topic(s):	Molecular Dynamics, High Performance Computing, Physics Simulation

WEDNESDAY, SEPT 22, 11:30 (20 MINUTES)
ROOM K

2117 Migration of C and Fortran Apps to GPGPU using HMPP

GPGPU is a tremendous opportunity to many application fields. Migrating legacy software to GPGPU is a complex process that requiresmastering the technological risks (e.g. loss of code portabilit, extensive code restructuration, debugging complexity) as well as costs. In this talk, we present a methodology based on HMPP (Heterogeneous Multicore Parallel Programming), allowing incremental processes that reduce the cost and risks of porting codes to GPGPU.

Speaker(s):	Francois Bodin (CTO, CAPS Entreprise)
Topic(s):	High Performance Computing, Tools & Libraries

WEDNESDAY, SEPT 22, 12:00 (120 MINUTES)
EXHIBIT HALL

1004 Exhibits Open / Networking Lunch

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.

Topic(s):	General Interest

WEDNESDAY, SEPT 22, 14:00 (50 MINUTES)
MARRIOTT GUADALUPE ROOM

2000 Gravitational N-body Simulations: How Massive Black Holes Interact with Stellar Systems

Astrophysics is a field where super computing is a must to obtain new scientific results. in particular, the study of the interaction among massive black holes and surrounding stars is a hot topic, which requires heavy computations to have good representation of what happens in the inner regions of galaxies. We present the results obtained with our high precisioned N-body code, NBSymple, which exploits the joint power of a multi core CPU system together with the high performance NVIDIA Tesla C1060 GPUs.

The code is available at the website:
astrowww.phys.uniroma1.it/dolcetta/nbsymple.html

Speaker(s):	Roberto Capuzzo-Dolcetta (Professor, Sapienza Univ. of Roma), Alessandra Mastrobuono Battisti (PhD Student, Sapienza- University of Rome)
Topic(s):	Astronomy & Astrophysics, Algorithms & Numerical Techniques

WEDNESDAY, SEPT 22, 14:00 (50 MINUTES)
ROOM D

2038 The Best of Both Worlds: Flexible Data Structures for Heterogeneous Computing

Learn how to switch between array of structs (AoS) and struct of arrays (SoA) storage without having to change the data access syntax. A few changes to the struct and container definitions will enable you to evaluate the performance of AoS vs. SoA on your existing AoS code. We present a simple abstraction that retains the more intuitive AoS syntax array[index]component, yet allows you to switch between AoS and SoA storage with a single template parameter at class definition.

Speaker(s):	Robert Strzodka (Senior Researcher, Max Planck Institut Informatik)
Topic(s):	Algorithms & Numerical Techniques, Tools & Libraries

WEDNESDAY, SEPT 22, 14:00 (50 MINUTES)
ROOM A3

2041 PyCUDA: Even Simpler GPU Programming with Python

Explore PyCUDA, a robust, open-source toolkit that lets you control your GPU from the comfort of Python, a Matlab-like scripting language. Learn about Fermi tuning with PyCUDA, the new interfaces for CUBLAS and CUFFT, the ecosystem of third-party libraries built on PyCUDA, and examples illustrating PyCUDA's benefits to large-scale applications.

Speaker(s):	Andreas Kloeckner (Courant Instructor, Courant Institute, NYU)
Topic(s):	Tools & Libraries, Computational Fluid Dynamics, Physics Simulation


Your HPC project success is determined by the strength of your infrastructure.
Introducing The New AMAX "ClusterMax" SuperG Platform
With integrated GPGPU technology and parallel architecture, the new AMAX ClusterMax SuperG platform allows your infrastructure to achieve accuracy, attain efficiency and accelerate applications. Scalable and reliable, the ClusterMax SuperG delivers the strength your organization relies on to solve the most complex computing challenges.
To find out how AMAX can build an HPC solution for your specific application, visit us at GPU Technology Conference 2010, Table #89/91 or at www.amax.com.
AMAX
AMAX Information Technologies

WEDNESDAY, SEPT 22, 14:00 (20 MINUTES)
ROOM K

2045 Roe-Pike Scheme for 2D Euler Equations

Hear how we are improving our elsA and CEDRE computational fluid dynamics software by working on solving the Euler equations set on the GPU. We discuss how our implementation considers the associated Riemann problem and the Roe-Pike differencing scheme at several orders in space while also introducing immerse boundary conditions. Covers the significant speedup obtained through algorithmic and computational optimizations.

Speaker(s):	Matthieu Lefebvre (PhD student, ONERA)
Topic(s):	Computational Fluid Dynamics, Algorithms & Numerical Techniques

WEDNESDAY, SEPT 22, 14:00 (50 MINUTES)
ROOM B

2068 Parallelizing FPGA Technology Mapping using GPUs

FPGA technology mapping is an algorithm that is heavily data parallel, but contains many features that make it unattractive for GPU implementation. The algorithm uses data in irregular ways since it is a graph-based algorithm. It also makes heavy use of constructs like recursion which is not supported by GPU hardware. In this paper, we take a state-of-the-art FPGA technology mapping algorithm within Berkeley's ABC package and attempt to parallelize it on a GPU. We show that runtime gains of 3.1x are achievable while maintaining identical quality as demonstrated by running these netlists through Altera's Quartus II place-and-route tool.

Speaker(s):	Doris Chen (Student, University of Toronto)
Topic(s):	Algorithms & Numerical Techniques,

WEDNESDAY, SEPT 22, 14:00 (50 MINUTES)
ROOM L

2120 High Performance Complex Event Processing on GPGPU

Complex Event processing (CEP),a crucial component in enterprise-scale applications, is the key element in that it allows applications to process the incoming event streams and apply relevant techniques in real-time for quicker decisions, making it easy to identify complex patterns in the events. Much of the time, this system is consumed by the event matching algorithms. Our work utilizes the highly parallel GPU for event matching algorithm wherein every incoming event is worked upon by this algorithm and results in high throughput.

Speaker(s):	Murali Krishna (Junior Research Associate, Infosys Technologies Limited), Sudeep Mallick (Principle Research Scientist, Infosys)
Topic(s):	Databases & Data Mining, Finance

WEDNESDAY, SEPT 22, 14:00 (50 MINUTES)
ROOM A1

2125 Developing GPU Enabled Visual Effects For Film And Video

The arrival of fully programable GPUs is now changing the visual effects industry, which traditionally relied on CPU computation to create their spectacular imagery. Implementing the complex image processing algorithms used by VFX is a challenge, but the payoffs in terms of interactivity and throughput can be enormous. Hear how The Foundry's novel image processing architecture simplifies the implementation of GPU-enabled VFX software and eases the transition from a CPU based infrastructure to a GPU based one.

Speaker(s):	Bruno Nicoletti (CTO, The Foundry)
Topic(s):	Film, Tools & Libraries, Video Processing

WEDNESDAY, SEPT 22, 14:00 (50 MINUTES)
ROOM E

2137 CUDA for Real-Time Multigrid Finite Element Simulation of Soft Tissue Deformations

The take-away of this presentation is an efficient CUDA implementation of a finite hexahedra multigrid solver for simulating elastic deformable models in real time. Due to the regular shape of the numerical stencil induced by the hexahedral regime, computations and data layout can be restructured to avoid execution divergence and to support memory access patterns enabling the hardware to coalesce multiple memory accesses into single memory transactions. This enables to effectively exploit the GPU's parallel processing units and high memory bandwidth. Performance gains of up to a factor of 12 compared to a highly optimized CPU implementation are demonstrated.

Speaker(s):	Christian Dick (PostGraduate Fellow ,Technische Universität München), Joachim Georgii (PostDoc, Technische Universität München)
Topic(s):	Physics Simulation, Algorithms & Numerical Techniques, High Performance Computing

WEDNESDAY, SEPT 22, 14:00 (50 MINUTES)
ROOM A7

2139 Interactive Histology of Large-Scale Biomedical Image Stacks

Get the latest information on leveraging GPU computing to process and visualize large-scale biomedical image stacks. We will discuss both display-aware processing and GPU-accelerated texture compression for histology applications on the GPU.

Speaker(s):	Won-Ki Jeong (Research Scientist, Harvard University), Jens Schneider (Postdoctoral Fellow, King Abdullah University of Science and Technology)
Topic(s):	Medical Imaging & Visualization, Imaging, Life Sciences

WEDNESDAY, SEPT 22, 14:00 (50 MINUTES)
ROOM A5

2140 Superfast Nearest Neighbor Searches Using a Minimal kd-tree

Learn how to adapt a kd-tree spatial data structure for efficient nearest neighbor (NN) searches on a GPU. Although the kd-tree is not a natural fit for GPU implementation, it can still be effective with the right engineering decisions. By bounding the maximum height of the kd-tree, minimizing the memory footprint of data structures, and optimizing the GPU kernel code, multi-core GPU NN searches with tens of thousands to tens of millions of points run 10-40 times faster than the equivalent single-core CPU NN searches.

Speaker(s):	Shawn Brown (Graduate Student, UNC, Chapel Hill)
Topic(s):	Algorithms & Numerical Techniques, Databases & Data Mining, Machine Learning & Artificial Intelligence

WEDNESDAY, SEPT 22, 14:00 (50 MINUTES)
ROOM A2

2142 Complex Geophysical Imaging Algorithms Enabled by GPU technology

Learn how computational expensive geophysical methods with 100s of TB of data become a commercial reality through the

adoption of GPUs. The first part of the talk will give an overview of the computational challenges for imaging facing the oil and gas industry. The second part will show how the current most advanced methods are taking advantage of the GPU technology.

Speaker(s):	David Nichols (Research Director, Schlumberger)
Topic(s):	Energy Exploration, Algorithms & Numerical Techniques, High Performance Computing

WEDNESDAY, SEPT 22, 14:00 (50 MINUTES)
ROOM C

2164 Analytical Performance Models to Improve the Efficiency of GPU Computing

Dive deep into a simple analytical model that provides insight into performance bottlenecks of parallel applications on GPU architectures. We will discuss how the model estimates the execution time of massively parallel programs. We will also cover how to optimize applications based on our developed performance analysis models.

Speaker(s):	Hyesoon Kim (Assistant Professor, Georgia Tech)
Topic(s):	Tools & Libraries

WEDNESDAY, SEPT 22, 14:00 (50 MINUTES)
MARRIOTT SAN JOSE BALLROOM

2204 Bridging GPU Computing and Neuroscience to Build Large-Scale Face Recognition on Facebook.

Biologically-inspired computer vision algorithms – those that aim to mirror the computations performed by the brain's visual system – have emerged as exceptionally promising candidates in object and face recognition research, achieving performance on a range of object and face recognition tasks. Recently, we have begun harnessing the newly-available power of NVIDIA GPUs to tackle the problem of biologically-inspired model selection within a largescale model search framework, drawing inspiration from high-throughput screening approaches in molecular biology and genetics where a large number of organisms are screened in parallel for a given property of interest.

As the available computational power provided by massively paralleltechnology from NVIDIA continues to expand, w e hope that this research will hold great potential for new social networking applications in addition to rapidly accelerating progress in artificial vision, and for generating new, experimentally testable hypotheses for the study of biological vision.

Speaker(s):	Nicolas Pinto (PhD Student, MIT)
Topic(s):	Computer Vision, High Performance Computing, Machine Learning & Artificial Intelligence, Neuroscience

WEDNESDAY, SEPT 22, 14:00 (50 MINUTES)
ROOM A8

2246 The Challenges of Integrating CUDA Engines Into an Existing Package, Yet Not Sinking the Boat

Based on a true story, come listen to a daring tale about the process of integrating a large CUDA component (physical engine) into an existing product (3D engine) replacing some of its functionality. The architectural difficulties and finer points that needed to be addressed. The tuning and testing of such a large system, while not effecting the stability of the original system.

Speaker(s):	Eri Rubin (Senior CUDA R&D developer, OptiTex)
Topic(s):	Physics Simulation, Tools & Libraries

WEDNESDAY, SEPT 22, 14:00 (50 MINUTES)
ROOM N

2248 Parallel Processing on GPUs at the University of Utah

The University of Utah is a CUDA Center of Excellence. We have been doing both basic and applied research using CUDA. In this session, we plan to give 3-4 talks on ongoing research. Most of the work that we will be presenting has been peered reviewed at top conferences.

Speaker(s):	Claudio Silva (Professor, University of Utah), Huy Vo (Research Assistant, University of Utah)
Topic(s):	High Performance Computing, Life Sciences, Medical Imaging & Visualization, Tools & Libraries

WEDNESDAY, SEPT 22, 14:00 (50 MINUTES)
ROOM M

2287 Internal GPUs on Dedicated x16 Slots - Are They Needed For HPC? (Sponsored by Dell)

We have benchmarked the real performance impact on a series of GPU accelerated applications to understand the benefits and drawbacks of different system level configurations. Come hear about the effects on performance of GPUs in shared slots and of GPUs that are externally connected.

Speaker(s):	Mark Fernandez (Computer Scientist, Dell)
Topic(s):	High Performance Computing

WEDNESDAY, SEPT 22, 14:00 (50 MINUTES)
KEYNOTE HALL

4002 Emerging Companies: CEO on Stage featuring Allegorithmic SAS, Bunkspeed, and miGenius

See the hottest new technologies from startups that are transforming computing.

In a lively and fast-paced exchange, the "Emerging Companies Summit - CEO on Stage" sessions will feature CEOs from three startups who will each have 15 minutes to introduce their companies and interact with a panel of leading venture capitalists, technology executives, and industry analysts.

Panelist(s):	Drew Lanza (Partner, Morgenthaler), Dan'l Lewin (Corporate VP and Strategic and Emerging Business Development, Microsoft), Jon Peddie (President, JPR), Jeff Herbst (Vice President of Business Development, NVIDIA)
Speaker(s):	Philip Lunn (CEO, Bunkspeed), Dr Sébastien Deguy (Founder and CEO, Allegorithmic), Chris Blewitt (Director, miGenius Limited)
Topic(s):	General Interest, Cloud Computing, Computer Graphics, Mobile & Tablet & Phone

WEDNESDAY, SEPT 22, 14:30 (20 MINUTES)
ROOM K

2049 Deflated Preconditioned Conjugate Gradient on the GPU

Explore how to use deflation as a second level preconditioning technique to speed up Block Incomplete Cholesky Preconditioned Conjugate Gradient Method. We use it to solve the Pressure correction equation involved in the solution of the Two-Phase Fluid Flow problem. Our implementation reaches speedup factors between 25-30, for more than 260,000 unknowns, when compared to the CPU.

Speaker(s):	Rohit Gupta (Researcher/Teacher, Delft University of Technology), Kees Vuik (Professor, Delft University of Technology)
Topic(s):	Computational Fluid Dynamics, Algorithms & Numerical Techniques

WEDNESDAY, SEPT 22, 15:00 (50 MINUTES)
ROOM A1

2029 Computer Vision Algorithms for Automating HD Post-Production

Discover how post-production tasks can be accelerated by taking advantage of GPU-based algorithms. In this talk we present computer vision algorithms for corner detection, feature point tracking, image warping and image inpainting, and their efficient implementation on GPUs using CUDA. We also show how to use these algorithms to do real-time stabilization and temporal re-sampling (re-timing) of high definition video sequences, both common tasks in post-production. Benchmarking of the GPU implementations against optimized CPU algorithms demonstrates a speedup of approximately an order of magnitude.

Speaker(s):	Hannes Fassold (Scientist, JOANNEUM RESEARCH)
Topic(s):	Computer Vision, Video Processing

WEDNESDAY, SEPT 22, 15:00 (50 MINUTES)
MARRIOTT GUADALUPE ROOM

2044 GRASSY: Leveraging GPU Texture Units for Asteroseismic Data Analysis

Learn how to use the hidden computation capability of GPU texture units for general purpose computation. We describe GRASSY, a system for stellar spectral synthesis where the core problem is interpolation between pre-computed intensity value. We map these pre-computed tables to the GPU's texture memory. Interpolation then becomes a texture lookup where the hardware automatically performs the interpolation, albeit at very low precision. Our mathematical framework reasons about the impact of this precision and our performance results show 500X speedups. This work generalizes the GPU texture units as computation engines and opens up new problems for GPU acceleration.

Speaker(s):	Matt Sinclair (Research Assistant, UW-Madison)
Topic(s):	Astronomy & Astrophysics, High Performance Computing

WEDNESDAY, SEPT 22, 15:00 (50 MINUTES)
ROOM N

2050 Copperhead: Data-Parallel Python for the GPU

Learn how to write Python programs that execute highly efficiently on GPUs using Copperhead, a data-parallel Python runtime. Using standard Python constructs like map and reduce, we will see how to construct data-parallel computations and embed them in Python programs that interoperate with numerical and visualization libraries such as NumPy, SciPy and Matplotlib. We will examine how to express computations using Copperhead, explore the performance of Copperhead programs running on GPUs, and discuss Copperhead's runtime model, which enables data-parallel execution from within Python.

Speaker(s):	Bryan Catanzaro (PhD Candidate, University of California, Berkeley)
Topic(s):	Tools & Libraries

WEDNESDAY, SEPT 22, 15:00 (50 MINUTES)
ROOM M

2080 Tackling Multi-Gigabit Design Challenges with a Practical Virtual EMI/ESD Lab

Learn about efficient methodologies for performant and cost-effective EMI and ESD suppression techniques by means of massive GPU parallel processing for simulations. We will discuss solving ever more complicated EMI and ESD challenges very early in the design process using in a so called 'Virtual EMI/ESD lab'.

Speaker(s	Davy Pissoort (Professor, KHBO-FMEC), Amolak Badesha (Senior Application Expert & Strategist, Agilent Technologies), Hany Fahmy (Director, SI/EMC Engineering, NVIDIA)
Topic(s):	Physics Simulation, Tools & Libraries

WEDNESDAY, SEPT 22, 15:00 (50 MINUTES)
ROOM C

2122 Using GPUs for Real-Time Brain-Computer Interfaces

Learn how GPU processing can provide researchers with an inexpensive and versatile alternative to dedicated signal processing hardware for real-time neural prosthetics. Topics will include an overview of algorithms, current state-of-the-art hardware, GPU processing in a real-time environment, multi-platform processing, and future directions in BCIs using GPU processing.

Speaker(s):	Adam Wilson (Postdoctoral Fellow, University of Cincinnati)
Topic(s):	Neuroscience, Algorithms & Numerical Techniques, Signal processing

WEDNESDAY, SEPT 22, 15:00 (50 MINUTES)
ROOM E

2170 Lattice Boltzmann Multi-Phase Simulations in Porous Media using GPUs

Learn how a very efficient implementation of multiphase lattice Boltzmann methods (LBM) based on CUDA delivers significant benefits for predictions of properties in rocks. This simulator on NVIDIA hardware enables us to perform pore scale multi-phase (oil-water-matrix) simulations in natural porous media and to predict important rock properties like absolute permeability, relative permeabilites, and capillary pressure. We will show videos of these simulations in complex real world porous media and rocks.

Speaker(s):	Jonas Toelke (Chief Computational Software Development, Ingrain)
Topic(s):	Computational Fluid Dynamics, Energy Exploration

WEDNESDAY, SEPT 22, 15:00 (50 MINUTES)
ROOM A2

2174 Reverse Time Migration on GPUs

Learn how GPUs can be used to accelerate subsurface imaging for Oil & Gas exploration. We will discuss results and lessons learned while implementing a Reverse Time Migration algorithm on GPUs achieving significant performance improvements over a comparable CPU implementation.

Speaker(s):	Alex Loddoch (Research Scientist, Chevron)
Topic(s):	Energy Exploration, High Performance Computing

WEDNESDAY, SEPT 22, 15:00 (50 MINUTES)
ROOM A7

2211 Modern Architecture for Massively Parallel Medical Tomographic Image Reconstruction on a GPU Cluster

Learn how to combine GPU and Cluster Programming with a real-world example. Many aspects of medical tomographic image reconstruction are embarrassingly parallel, but require massive compute power. We distribute the load onto a cluster of multi-GPU equipped nodes using Message Passing Interface (MPI) and CUDA. The Thrust library allows for a modern object-oriented approach.

Speaker(s):	Sven Prevrhal (Staff Research Scientist, Philips), Jingyu Cui (Graduate Student, Stanford University)
Topic(s):	Medical Imaging & Visualization, Algorithms & Numerical Techniques, High Performance Computing, Tools & Libraries

WEDNESDAY, SEPT 22, 15:00 (50 MINUTES)
ROOM B

2218 Redesigning Molecular Dynamics for GPUs and GPU Clusters

Generalized Born and Particle Mesh Ewald (PME) molecular dynamics are two computationally intensive algorithms for simulating biological molecules. While several adaptations of Generalized Born have attained excellent speedup on GPUs, high performance Particle Mesh Ewald has been more elusive. Here we describe in detail a recent port of PME implemented within AMBER 11 that has achieved performance on par with up to 128 nodes of a top ten supercomputer.

Speaker(s):	Scott Le Grand (Principal Engineer, NVIDIA)
Topic(s):	Molecular Dynamics, Algorithms & Numerical Techniques, High Performance Computing, Life Sciences

WEDNESDAY, SEPT 22, 15:00 (50 MINUTES)
ROOM A3

2234 Unstructured Finite Volume Code on a Cluster with Multiple GPUs per Node

Explore how a code written to run in parallel using OpenMP and on a single GPU was modified to run across multiple GPUs and nodes on a multi-CPU, multi-GPU cluster installed at the Naval Research Laboratory. We will discuss the performance of this code running in parallel using MPI/OpenMP and MPI/CUDA.

Speaker(s):	Keith Obenschain (Computer Scientist, Naval Research Lab), Andrew Corrigan (Naval Research Laboratory & George Mason University)
Topic(s):	Computational Fluid Dynamics, High Performance Computing

WEDNESDAY, SEPT 22, 15:00 (50 MINUTES)
ROOM L

2237 Accelerating Business Intelligence Applications with Fast Multidimensional Aggregation

In this research session, we present an approach using NVIDIA GPUs as massively parallel coprocessors for in-memory OLAP computations. Early tests have shown speedup factors of more than 40x compared to optimized sequential algorithms on a CPU. In addition to the data structures and algorithms involved, we describe a method to extend the approach to systems with more than one GPU in order to scale it to larger data sets.

Speaker(s):	Tobias Lauer (Researcher, University of Freiburg), Christoffer Anselm (Software Developer, Jedox Business Intelligence)
Topic(s):	Databases & Data Mining

WEDNESDAY, SEPT 22, 15:00 (50 MINUTES)
ROOM A5

2238 Better Performance at Lower Occupancy

It is usually advised to optimize CUDA kernels for higher occupancy to hide memory and arithmetic latencies better. In this presentation, I show that increasing occupancy is not the only way and not always the best way to hide latency on GPU. Instead, it may be advantageous to rely on the parallelism within threads-instruction-level parallelism. This insight yields a simple optimization technique that is used in later versions of CUBLAS and CUFFT. I discuss the rationale behind the technique and illustrate it by speeding up matrix multiplication, starting with the basic implementation found in the NVIDIA GPU Computing SDK.

Speaker(s):	Vasily Volkov (Student, UC Berkeley)
Topic(s):	High Performance Computing

WEDNESDAY, SEPT 22, 15:00 (20 MINUTES)
ROOM K

2251 TotalView Debugger for CUDA

Hear how the TotalView debugger is being extended to support GPU computation with CUDA. In addition to the basic challenges associated with debugging parallel programming, CUDA programming introduces a number of new concepts for which developers need visibility in debugging: a hierarchical memory, near-SIMD warps, streams, and kernels, among others. How do we create a tool that handles it all? We'll be discussing the status of our work and the challenges encountered in bringing this all together into a single package, TotalView for CUDA.

Speaker(s):	Chris Gottbrath (Principal Product Manager, TotalView Technologies, Inc., a Rogue Wave Software company)
Topic(s):	Tools & Libraries

WEDNESDAY, SEPT 22, 15:00 (50 MINUTES)
ROOM A8

2273 GPUs In the Front Line of our Defenses (Sponsored by GE)

Find out how GPUs are accelerating defense and aerospace applications and providing superior information processing to drive the next generation of capabilities to protect both homelands and soldiers. Learn how rugged VPX hardware and software architectures are able to scale from small power- & weight-constrained vehicles through to large complex processing arrays, on platforms as diverse as unmanned aerial vehicles (UAV), through tracked ground vehicles, and to ship borne radar.

Speaker(s):	Simon Collins (Product Manager, GE Intelligent Platforms)
Topic(s):	High Performance Computing

WEDNESDAY, SEPT 22, 15:00 (50 MINUTES)
MARRIOTT SAN JOSE BALLROOM

2281 Domain-Specific Languages

Computer graphics has introduced several domain-specific languages (DSLs) that enable high performance and parallelism for narrow problem domains - RenderMan, Cg, GLSL, and recently OpenRL and OptiX. We think that similar approaches can benefit other areas of GPU computing - visualization, animation, physics simulation, or scientific data analysis. In this talk, we present Shadie, a domain-specific shading language for rapid development of complex custom volume visualizations in radiation oncology. The shaders are written in a high-level Python-like language and translated to CUDA for efficiency. We will explain how you can develop your own DSLs using source-to-source translation and a suitable backend library.

Speaker(s):	Hanspeter Pfister (Professor, Harvard University)
Topic(s):	Programming Languages & Techniques, General Interest

WEDNESDAY, SEPT 22, 15:00 (50 MINUTES)
ROOM D

2296 CUDA Optimization for Ninjas: A Case Study of High-Performance Sorting

In this presentation, we use our implementation for high performance radix sorting as a case study for illustrating advanced design patterns and idioms. These techniques have allowed us to demonstrate Fermi sorting rates that exceed 1.0 billion 32-bit keys per second (and over 770 million key-value pairs per second), making it the fastest fully-programmable micro-architecture for this genre of sorting problems.

Although the CUDA programming model is elegantly decoupled from any particular hardware configuration, we present techniques for exploiting knowledge of the NVIDIA GPU machine model in order to produce more efficient implementations. Our design patterns enable the compiler to specialize a single program text for a variety of architectures, resulting in target code that "fits" the underlying hardware significantly better than more general approaches. In particular, we discuss strategies for kernel fusion, warp-synchronous programming, flexible granularity via meta-programming, algorithm serialization, and data-movement tuning.

Speaker(s):	Duane Merrill (PhD Candidate, University of Virginia)
Topic(s):	Programming Languages & Techniques

WEDNESDAY, SEPT 22, 15:00 (50 MINUTES)
KEYNOTE HALL

4003 Emerging Companies Summit Panel: GPUs for Computer Vision

Moderated by Jon Peddie, President at Jon Peddie Research

The GPU (graphics processing unit) runs advanced applications which are transforming existing industries and creating new ones. Join our panel of leading industry experts as they discuss the latest technology advances in the usage of GPUs for Computer Vision, they will cover facial, gesture, human motion, and biometrics recognition, augmented reality, robotic computing and more.

Panelists:	Sam Cox (CEO, Milabra), Tom Dean (Research Scientist, Google) Janko Mrsic-Flogel (CTO, MirriAd), Joe Stam (Sr. Applications Engineer, NVIDIA), Yoram Yaacovi (CTO & General Manager, Technologies at Microsoft)
Topic(s):	Computer Vision

WEDNESDAY, SEPT 22, 15:30 (20 MINUTES)
ROOM K

2143 CUDA Fortran Programming for NVIDIA GPUs

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

Speaker(s):	Brent Leback (Engineering Manager, The Portland Group)
Topic(s):	Tools & Libraries, High Performance Computing, Programming Languages & Techniques

WEDNESDAY, SEPT 22, 16:00 (50 MINUTES)
ROOM D

2020 GPU-Accelerated Data Expansion for the Marching Cubes Algorithm

Learn how to accelerate marching cubes on the GPU by taking advantage of the GPU's high memory bandwidth and fast on-chip shared memory in a data expansion algorithm that can extract the complete iso-surface mesh from (dynamic) volume data without requiring any data transfers back to the CPU.

Speaker(s):	Chris Dyken (Research Scientist, SINTEF), Gernot Ziegler (Developer Technology (Compute), NVIDIA)
Topic(s):	Algorithms & Numerical Techniques, Imaging, Medical Imaging & Visualization

WEDNESDAY, SEPT 22, 16:00 (50 MINUTES)
ROOM K

2069 GPU-Accelerated Business Intelligence Analytics

Join us and learn why GPU computing is a game changer for business intelligence (BI). We will discuss how GPUs can be used to accelerate BI analytics at much lower cost, higher performance, and better power efficiency than other alternatives.

Speaker(s):	Ren Wu (Senior Scientist, HP Labs)
Topic(s):	Databases & Data Mining, Finance, High Performance Computing

WEDNESDAY, SEPT 22, 16:00 (50 MINUTES)
ROOM A1

2072 GPUs at the Computer Animation Studio

Learn five simple ways in which GPUs have been adopted in the production pipeline at Blue Sky Studios. Covers how we use GPUs to improve animation tools, add real-time anaglyph support, and accelerate noise functions including code samples from production tools.

Speaker(s):	Hugo Ayala (Sr. Research Associate, Blue Sky Studios)
Topic(s):	Film, Stereoscopic 3D, Tools & Libraries

WEDNESDAY, SEPT 22, 16:00 (50 MINUTES)
ROOM B

2073 High Performance Molecular Simulation, Visualization, and Analysis on GPUs

This talk will present recent successes in the use of GPUs to accelerate interactive visualization and analysis tasks on desktop computers, and batch-mode simulation and analysis jobs on GPU-accelerated HPC clusters. We'll present Fermi-specific algorithms and optimizations and compare with those for other devices. We'll also present performance and performance/watt results for NAMD molecular dynamics simulations and VMD analysis calculations on GPU clusters, and conclude with a discussion of ongoing work and future opportunities for GPU acceleration, particularly as applied to the analysis of petascale simulations of large biomolecular complexes and long simulation timescales.

Speaker(s):	John Stone (Senior Research Programmer, University of Illinois at Urbana-Champaign)
Topic(s):	Molecular Dynamics, Algorithms & Numerical Techniques, High Performance Computing, Life Sciences

WEDNESDAY, SEPT 22, 16:00 (50 MINUTES)
ROOM E

2083 GPU Accelerated Solver for the 3D Two-phase Incompressible Navier-Stokes Equations

This demonstrates the potential of GPUs for solving complex free surface flow problems using level set methods. These methods are capable of producing complex surface deformations, and therefore are used widely in computer graphics, as well as engineering applications. This work demonstrates that GPUs can be used to accelerate the most computationally expensive part of free surface flow calculations, and therefore allows much larger problems to be solved on workstation machines than was previously possible. These techniques will be exemplified by our current project to port our in-house fluid solver NaSt3DGPF to the GPU.

Speaker(s):	Peter Zaspel (Research Assistant University of Bonn)
Topic(s):	Computational Fluid Dynamics, Algorithms & Numerical Techniques, High Performance Computing, Physics Simulation

WEDNESDAY, SEPT 22, 16:00 (50 MINUTES)
ROOM C

2093 Computational Photography: Real-Time Plenoptic Rendering

Get the latest information on GPU-based plenoptic rendering including a demonstration of refocusing, novel view generation, polarization, high dynamic range, and stereo 3D. Learn how GPU hardware enables plenoptic rendering tasks with high-resolution imagery to be performed interactively, opening up entirely new possibilities for modern photography.

Speaker(s):	Andrew Lumsdaine (Professor, Indiana University), Georgi Chunev (Research Assistant, Indiana University), Todor Georgiev (Senior Research Scientist II, Adobe Systems)
Topic(s):	Imaging, Computer Vision, Stereoscopic 3D

WEDNESDAY, SEPT 22, 16:00 (50 MINUTES)
MARRIOTT GUADALUPE ROOM

2108 Binary Black Holes Simulations using CUDA

Get the latest information on how to evolve binary black holes simulations on GPUs.

Speaker(s):	Abdul Mroue (Post-Doc Fellow, CITA, University of Toronto)
Topic(s):	Astronomy & Astrophysics, Algorithms & Numerical Techniques, Physics Simulation

WEDNESDAY, SEPT 22, 16:00 (50 MINUTES)
ROOM A8

2118 Large-scale Gas Turbine Simulations on GPU Clusters

This talk describes a strategy for implementing structured grid PDE solvers on GPUs. Techniques covered include the use of source-to-source compilation and the use of sparse matrix vector multiplications for complicated boundary conditions. A new production-quality solver for flows in turbomachines called Turbostream that uses these techniques is presented. The impact of the use of GPUs on the turbomachinery design process is demonstrated by two 64-GPU simulations that have recently been performed on the University of Cambridge's GPU cluster.

Speaker(s):	Tobias Brandvik (PhD Student, University of Cambridge)
Topic(s):	Computational Fluid Dynamics,

WEDNESDAY, SEPT 22, 16:00 (50 MINUTES)
MARRIOTT SAN JOSE BALLROOM

2135 Processing Petabytes per Second with the ATLAS Experiment at the Large Hadron Collider at CERN

Learn how GPUs could be adopted by the ATLAS detector at the Large Hadron Collider (LHC) at CERN. The detector, located at one of the collision points, must trigger on unprecedented data acquisition rates (PB/s), to decide whether to record the event, or lose it forever. In the beginning, we introduce the ATLAS experiment and the computational challenges it faces. The second part will focus on how GPUs can be used for algorithm acceleration - using two critical algorithms as exemplars. Finally, we will outline how GPGPU acceleration could be exploited and incorporated into the future ATLA computing framework.

Speaker(s):	Philip Clark (Reader (Associate Professor) in Particle Physics, University of Edinburgh), Andy Washbrook (Postdoctoral Research Assistant, University of Edinburgh)
Topic(s):	High Performance Computing, Algorithms & Numerical Techniques, Physics Simulation

WEDNESDAY, SEPT 22, 16:00 (50 MINUTES)
ROOM A7

2144 Large-Scale Visualization Using A GPU Cluster

Learn how to visualize extremely large-scale scientific data using GPGPU techniques on a GPU-accelerated visualization cluster. Recent advances in general-purpose GPU (GPGPU) computing provide a promising solution to compute-intensive scientific visualization. However, the largest scientific simulations produce datasets that are orders of magnitude larger than the memory available on current GPUs. Many distributed GPUs must be used in parallel. We present Longhorn, currently the world's largest GPU-enhanced cluster dedicated for visualization and data analysis, and describe the distributed memory architecture and GPGPU techniques to interactively visualize massive datasets using distributed GPUs on Longhorn.

Speaker(s):	Byungil Jeong (Visualization Scientist, TACC / UT-Austin), Paul Navratil (Visualization Scientist, Texas Advanced Computing Center)
Topic(s):	Medical Imaging & Visualization, High Performance Computing

WEDNESDAY, SEPT 22, 16:00 (50 MINUTES)
ROOM A5

2154 The Impact of Data Movement on GPU Performance

GPU computing has taken the scientific computing landscape by storm, fueled by the massively parallel arithmetic hardware. When coding, researchers rely on best practices that have been developed in the short timespan of GPGPU. This session challenges a widely held belief that transfers to/from the GPU device must be minimized to achieve the best performance by presenting a case study on CULA, our library for dense linear algebra. The topics to be discussed include the relationship between computation and transfer time for synchronous/ asynchronous transfers, and impact that data allocations have on memory performance and overall solution time.

Speaker(s):	John Humphrey (Senior Engineer, EM Photonics, Inc), Daniel Price (Engineer, EM Photonics, Inc.)
Topic(s):	High Performance Computing, Algorithms & Numerical Techniques, Tools & Libraries

WEDNESDAY, SEPT 22, 16:00 (50 MINUTES)
ROOM A3

2201 A Case Study of Accelerating Matlab Based Applications using GPUs

Learn how to accelerate Matlab based applications using GPUs. We cover a popular neuro-imaging software called SPM and show how to use CUDA and Jacket to speedup computationally intensive Matlab applications.

Speaker(s):	Aniruddha Dasgupta (Graduate Student, Georgia Institute of Technology)
Topic(s):	Medical Imaging & Visualization

WEDNESDAY, SEPT 22, 16:00 (50 MINUTES)
ROOM N

2217 GPU-Based Conjugate Gradient Solvers for Lattice QCD

Learn how to perform state-of-the-art quantum chromodynamics (QCD) computation using NVIDIA GPUs at 1% of the cost of a conventional supercomputer and 10% of its power consumption. We will discuss how physicists around the world are using GPU clusters to solve QCD. We will focus upon how TWQCD have been using a large GPU cluster (200 GPUs) to simulate QCD, attaining 36 Teraflops (sustained).

Speaker(s):	Ting-Wai Chiu (Professor, National Taiwan University)
Topic(s):	High Performance Computing, Physics Simulation

WEDNESDAY, SEPT 22, 16:00 (50 MINUTES)
ROOM A2

2226 Reverse Time Migration with GMAC

Get a close look at implementing Reverse Time Migration (RTM) applications across multiple GPUs. We will focus on how RTM applications can be scaled using the GMAC asymmetric distributed shared memory (ADSM) library to break the problem into manageable chunks. We will provide an introduction to GMAC and discuss handling boundary conditions and using separate kernels to improve efficiency.

Speaker(s):	Javier Cabezas (Researcher, Barcelona Supercomputing Center), Mauricio Araya (Senior Researcher, Barcelona Supercomputing Center)
Topic(s):	Energy Exploration, Algorithms & Numerical Techniques, High Performance Computing

WEDNESDAY, SEPT 22, 16:00 (50 MINUTES)
ROOM L

2252 Simulating Housefly Vision Elements Using OpenCL

An OpenCL GPU based computer simulation of a biologically motivated model, based on the anatomy of housefly's first optic ganglion, the lamina ganglionaris (the lamina layer) is presented. Specific to GPU technology, the computer model demonstrates: the implementation of a 2nd Order Runga-Kutta method to approximate coupled differential equations using GPU hardware; the mapping of a non-Cartesian coordinate system onto the Cartesian layout of the threads. Testing examined usage and access across device memory spaces to determine the optimal usage/access method for the ANN. This result was generalized for OpenCL GPU devices, using the capabilities of OpenCL.

Speaker(s):	Karen Haines (Professor, WASP/The University of Western Australia)
Topic(s):	Neuroscience, Algorithms & Numerical Techniques, Signal processing

WEDNESDAY, SEPT 22, 16:00 (50 MINUTES)
ROOM M

2302 Microsoft Technologies for HPC

NVIDIA Parallel Nsight provides access to the power of the GPU from within the familiar environment of Microsoft Visual Studio. In this session, we will expand on the computational power of Visual Studio 2010, Windows HPC Server and the Technical Computing Libraries and show how to increase your performance.

Speaker(s):	Calvin Clark (Senior Consultant, Microsoft)
Topic(s):	High Performance Computing

WEDNESDAY, SEPT 22, 16:00 (50 MINUTES)
KEYNOTE HALL

4004 Emerging Companies: CEO on Stage featuring Cooliris, empulse GmbH, and Playcast Media Systems

See the hottest new technologies from startups that could transform computing.

In a lively and fast-paced exchange, the "Emerging Companies Summit - CEO on Stage" sessions will feature CEOs from three startups who will each have 15 minutes to introduce their companies and interact with a panel of leading venture capitalists, technology executives, and industry analysts.

Panelist(s):	Nathan Brookwood (Research Fellow, Insight64), Charles Carmel (VP of Corporate Business Development, Cisco), Flip Gianos (General Partner, InterwestInterWest Partners), Jeff Herbst (Vice President of Business Development, NVIDIA)
Speaker(s):	Austin Shoemaker (Cooliris), Michael Hummel (Managing Director, empulse GmbH), Natan Peterfreund (CTO, Playcast Media Systems)
Topic(s):	General Interest, Databases & Data Mining, Video Processing, Computer Graphics

WEDNESDAY, SEPT 22, 17:00 (50 MINUTES)
ROOM A3

2005 Porting Large-Scale Legacy Fortran Codes

Explore a new automatic Fortran translator which has been developed and used to port the numerical subroutines of FEFLO, a general-purpose legacy Computational Fluid Dynamics code operating on unstructured grids, to run on the GPU. Data transfer to the CPU is minimized throughout the course of a CFD run. Benchmarks of large-scale production runs will be presented.

Speaker(s):	Andrew Corrigan (Research Mathematician Naval Research Laboratory & George Mason University), Rainald Löhner (Professor, George Mason University)
Topic(s):	Algorithms & Numerical Techniques, Computational Fluid Dynamics, Machine Learning & AI

WEDNESDAY, SEPT 22, 17:00 (50 MINUTES)
ROOM B

2006 Short-Range Molecular Dynamics on GPU

Learn how to accelerate short-range molecular dynamics using CUDA C. We will cover building the neighbor list and calculating the forces on the GPU. To handle the case where a few particles have significantly more neighbors than most other particles, we propose a hybrid data structure for the neighbor list that can achieve a good balance between performance and storage efficiency. A CUDA C implementation of the technique for Leonard-Jones forces can be found in the LAMMPS molecular dynamics open source code.

Speaker(s):	Peng Wang (Developer Technology Engineer, NVIDIA)
Topic(s):	Molecular Dynamics

WEDNESDAY, SEPT 22, 17:00 (60 MINUTES)
MARRIOTT SAN JOSE BALLROOM

2011 Fundamental Performance Optimizations for GPUs

This presentation covers the major CUDA optimizations. Topics will include: maximizing memory throughput, kernel launch configuration, using shared memory, and improving GPU/CPU interaction. While C for CUDA is used for illustration, the concepts covered will apply equally to programs written with OpenCL and DirectCompute APIs.

Speaker(s):	Paulius Micikevicius (Developer Technology Engineer, NVIDIA)
Topic(s):	Programming Languages & Techniques, Tools & Libraries

WEDNESDAY, SEPT 22, 17:00 (50 MINUTES)
ROOM A2

2014 Scalable Subsurface Data Visualization Framework

Mental Images' DiCE-based geospatial library is a CUDA and cluster-based visualization framework that enables scalable processing and rendering of huge amounts of subsurface data for interactive seismic interpretation.

Geospatial exploration in the oil and gas industries is concerned with scanning the earth's subsurface structure for detecting oil and for cost-effective drilling of detected oil reservoirs.

Efficient seismic interpretation requires the interpreters to be able to interactively explore huge amounts of volumetric seismic information with embedded stacked horizons to gain visual insight into the subsurface structure and to determine where oil recovery facilities and drilling infrastructure shall be built.

Speaker(s):	Tom-Michael Thamm (VP Products, mental images GmbH), Marc Nienhaus (Sen. Graphics Software Engineer, mental images GmbH)
Topic(s):	Energy Exploration, Databases & Data Mining, Imaging, Tools & Libraries

WEDNESDAY, SEPT 22, 17:00 (50 MINUTES)
ROOM C

2021 Efficient Volume Segmentation on the GPU

Explore a new technique in the detection of common regions in a 2D/3D data array. Connected components along the axes are linked before actual label propagation starts. The algorithm is completely gather-based, which allows for several optimizations in the CUDA C implementation. It enables real-time frame rates for the analysis of typical 2D images and interactive frame rates for the analysis of typical volume data.

Speaker(s):	Allan Rasmusson (PhD candidate/NVIDIA Intern, University of Aarhus), Gernot Ziegler (Developer Technology, Compute, NVIDIA)
Topic(s):	Algorithms & Numerical Techniques, Computer Vision, Imaging, Medical Imaging & Visualization

WEDNESDAY, SEPT 22, 17:00 (50 MINUTES)
ROOM A8

2077 Catastrophic Risk Management: Fast and Flexible with GPU Analytics

RMS will describe our experience leveraging GPUs and simple software architectural principles to deliver both spectacular performance gains and enhanced flexibility in next generation portfolio risk management applications.

Speaker(s):	Philippe Stephan (CTO, RMS)
Topic(s):	Finance

WEDNESDAY, SEPT 22, 17:00 (50 MINUTES)
ROOM A5

2089 Analyzing CUDA Accelerated Application Performance at 20 PFLOP/s

Learn how applications can be executed over multiple GPUs located in multiple hosts, what the challenges are to scale one application to a 20 PFLOP/s machine and why tool support is a necessity. Receive an overview on the available performance analysis tools that support CUDA developers in generating applications with outstanding speedups.

Speaker(s):	Guido Juckeland (Senior System Engineer (HPC), Leader Hardware Accelerator Group, TU Dresden - ZIH), Jeremy Meredith, (Computer Scientist, Oak Ridge National Laboratory)
Topic(s):	High Performance Computing, Tools & Libraries

WEDNESDAY, SEPT 22, 17:00 (50 MINUTES)
ROOM E

2128 Hybrid Quantum Mechanics/Electrodynamics (QM/ED) Modeling of Solar Cells on a CUDA Cluster

One of the greatest challenges of the twenty-first century is the utilization of renewable energy. In providing a theoretical explanation and guidelines for computer-aided design of dye-sensitized solar cell (DSSC), we recently developed a hybrid multi-scale quantum mechanics/classical electrodynamics (QM/ED) methodology.

Our numerical simulations were tested on a CUDA enabled Linux cluster using CP2K. We extended its CUDA implementation to MPI parallel environment. Our preliminary results demonstrated a superior performance advantage of hybrid MPI/GPGPU programming that could potentially shorten the total simulation wall time by an order of magnitude.

Speaker(s):	Hanning Chen (Research Associate, Northwestern University)
Topic(s):	Quantum Chemistry, Energy Exploration, Molecular Dynamics, Physics Simulation

WEDNESDAY, SEPT 22, 17:00 (50 MINUTES)
ROOM A1

2162 Real-time Reyes: Programmable Rendering on Graphics Processors

We present a discussion of ideas and techniques behind programmable graphics pipelines on modern GPUs, specifically the example design of a real-time Reyes renderer. Walking through this example, we address the philosophy beneath programmable GPU graphics, the broad strategy for the specific pipeline, and algorithmic and implementation-level details for key rendering stages. We cover several issues concerning GPU efficiency, including those involving work scheduling, parallelization of traditional stages, and balancing of rendering workloads. We expect the audience to gain an in-depth exposure of the state of research in programmable graphics, and an insight into efficient pipeline design for irregular workloads.

Speaker(s):	Anjul Patney (Graduate Student, University of California, Davis), Stanley Tzeng (Graduate Student, University of California, Davis)
Topic(s):	Computer Graphics, Film

WEDNESDAY, SEPT 22, 17:00 (50 MINUTES)
ROOM D

2167 Designing a Geoscience Accelerator Library Accessible from High Level Languages

Explore a library for geoscience applications on CUDA and OpenCL platforms. Target applications span atmosphere, ocean, geomorphology and porous media flows. These areas are linked by common numerical techniques encapsulated in our library. We will review the scope of the library, its meta-programming approaches, and its key design attributes. We will also demonstrate its support for multi-GPU parallelism within and across address spaces and provide examples of is use from high level languages including C, Fortran, and Python.

Speaker(s):	Chris Hill (Principle Research Scientist, M.I.T), Alan Richardson (Graduate Student, M.I.T)
Topic(s):	Programming Languages & Techniques, Algorithms & Numerical Techniques, Computational Fluid Dynamics, Tools & Libraries

WEDNESDAY, SEPT 22, 17:00 (50 MINUTES)
MARRIOTT GUADALUPE ROOM

2178 Using GPUs to Track Changes in the Sun

Learn how GPU computing is enabling astrophysicists to study our closest star. NASA's recently launched Solar Dynamics Observatory is continuously streaming full-disk images of the Sun at visible, UV and EUV wavelengths. This presentation will discuss ways that GPU computing is helping scientists cope with the analysis of the immense data volumes as well as in numerical modeling of the Sun.

Speaker(s):	Mark Cheung (Physicist, Lockheed Martin Solar & Astrophysics Laboratory)
Topic(s):	Astronomy & Astrophysics, Computer Vision, Computational Fluid Dynamics, Physics Simulation

WEDNESDAY, SEPT 22, 17:00 (50 MINUTES)
ROOM N

2242 Swarming Bacteria and Diffusing Particles: High-Throughput Analysis of Microscopic 3D Motion

Ever since the 1827 discovery of Brownian motion by observing pollen grains, quantifying motion under the microscope has led to breakthroughs in physics, biology and engineering. Here, I present methods we have developed using confocal microscopy to deduce 3D structure and dynamics from 2D image sequences. We analyze the motion of diffusing colloidal particles and swarms of bacteria free to swim in 3D, which we observe at the single-organism level. We rely heavily on GPU computing to process our large data sets, making extensive use of NPP, CuFFT and optical-flow CUDA algorithms originally developed for machine vision in automobiles.

Speaker(s):	Peter Lu (Post-Doctoral Research Fellow, Harvard University)
Topic(s):	Computer Vision, Imaging, Life Sciences

WEDNESDAY, SEPT 22, 17:00 (50 MINUTES)
ROOM A7

2243 Microsoft RemoteFX - GPU Virtualization for Desktop Centralization

Learn about Microsoft's upcoming GPU Virtualization feature, RemoteFX, which will ship in Windows Server 2008 R2 SP1. Microsoft RemoteFX enables GPUs to be hosted in the datacenter as a service that can be shared by multiple users for streaming the real-time and complete Windows 7 desktop experience to ultra-lightweight client devices anywhere on the corporate network. With Microsoft RemoteFX, users will be able to work remotely in a Windows Aero desktop environment, watch full-motion video, enjoy Silverlight animations, and run 3D applications – all with the fidelity of local-like performance.

Speaker(s):	Tad Brockway (Product Unit Manager, Microsoft)
Topic(s):	Cloud Computing, Computer Graphics

WEDNESDAY, SEPT 22, 17:00 (50 MINUTES)
ROOM K

2282 GPU-Enabled Biomedical Imaging

The purpose of this presentation is to describe several novel biomedical imaging applications which make extensive use of GPUs. In CT iterative reconstructions, for example, high performance computing is allowing us to see details and structures we previously were not able to discern.

Speaker(s):	Homer Pien (Director of the Laboratory for Medical Imaging and Computations, Massachusetts General Hospital / Harvard Medical School)
Topic(s):	Medical Imaging & Visualization, High Performance Computing, Imaging, Life Sciences

WEDNESDAY, SEPT 22, 17:00 (20 MINUTES)
ROOM L

2285 Walt Disney Animation Studios' GPU-Accelerated Animatic Lighting Process with Soft Shadows and Depth of Field

See how Walt Disney Animation's software uses OpenGL and GLSL shaders to interactively display depth of field, accurate lighting, and soft shadows in the Maya viewport. Learn how this improved our animatic process and helps us make better animated movies.

We'll show the tools in action and show the progression of a shot from standard Maya to final animatic look, and will compare the result with a production Renderman render. We'll also walk you through the GLSL shader render passes it uses to do deferred lighting and shadowing.

Speaker(s):	David Adler (Principal Software Engineer, Walt Disney Animation Studios)
Topic(s):	Film, Digital Content Creation (DCC)

WEDNESDAY, SEPT 22, 17:00 (50 MINUTES)
ROOM M

2294 GPU.NET with TidePowerd

Join TidePowerd for a demonstration of GPU.NET, our innovative new product which dramatically cuts the time needed to develop and maintain a GPU-based application by extending Microsoft's .NET Framework onto GPUs. With GPU.NET, your device-accelerated code can be written in any .NET-supported language (e.g., C#, F#, IronPython) and called like any other method - so it's easy to create new GPU-based applications without having to retrain your developers. You'll learn how to use GPU.NET to quickly develop a financial calculator in C#, use the built-in Visual Studio unit-testing tools to ensure the correctness of the code, and seamlessly deploy the application into a mixed Windows / Linux environment. We'll also discuss how GPU.NET expands the frontiers of GPU computing into lucrative new markets such business intelligence, database processing, and data visualization.

Speaker(s):	Jack Pappas (Co-founder, CEO, TidePowerd)
Topic(s):	Programming Languages & Techniques

WEDNESDAY, SEPT 22, 17:00 (50 MINUTES)
KEYNOTE HALL

4005 Emerging Companies: CEO on Stage featuring Jedox Business Intelligence, Rocketick, and Softkinetic

See the hottest new technologies from startups that are transforming computing.

In a lively and fast-paced exchange, the "Emerging Companies Summit - CEO on Stage" sessions will feature CEOs from three startups who will each have 15 minutes to introduce their companies and interact with a panel of leading venture capitalists, technology executives, and industry analysts.

Panelist(s):	Flip Gianos (General Partner, Interwest), Charles Carmel (VP of Corporate Business Development, Cisco), Nathan Brookwood (Research Fellow, Insight64), Jeff Herbst (Vice President of Business Development, NVIDIA)
Speaker(s):	Kristian Raue (CEO & Founder, Jedox Business Intelligence), Uri Tal (CEO, Rocketick), Michel Tombroff (CEO, Softkinetic)
Topic(s):	General Interest, Computer Vision, Databases & Data Mining, High Performance Computing

WEDNESDAY, SEPT 22, 17:30 (20 MINUTES)
ROOM L

2284 GPU Implementation of Collision-Based Deformation

Addressing the production needs for the upcoming Disney animated movie, we are in the process of developing a new Maya deformer that incorporates state-of-the-art collision-based deformations. Our deformer includes both dynamic and quasi-static solutions. Our solvers conserves volume and constrains surface area by solving linear systems in a graded volume mesh. To achieve realistic deformation in production-ready data at interactive rates, we leverage the computational power of the NVIDIA GPU architecture using CUDA. Our underlying data structure is specifically designed and optimized for CUDA (i.e. coalescing data access, minimizing CPU-GPU interaction, utilizing shared memory).

Speaker(s):	Dmitriy Pinskiy (Sr Software Engineer, Walt Disney Animation Studios)
Topic(s):	Film

WEDNESDAY, SEPT 22, 18:00 (120 MINUTES)
KEYNOTE HALL

1006 Exhibits Open / Networking Reception

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.

Topic(s):	General Interest

Designed By NVIDIA, Delivered By NextIO

NextIO vCORE™ Express

NextIO vCORE™ Express

NextIO is delivering the next generation of GPU rack solutions based on the new NVIDIA CUDA™ architecture codenamed "Fermi"

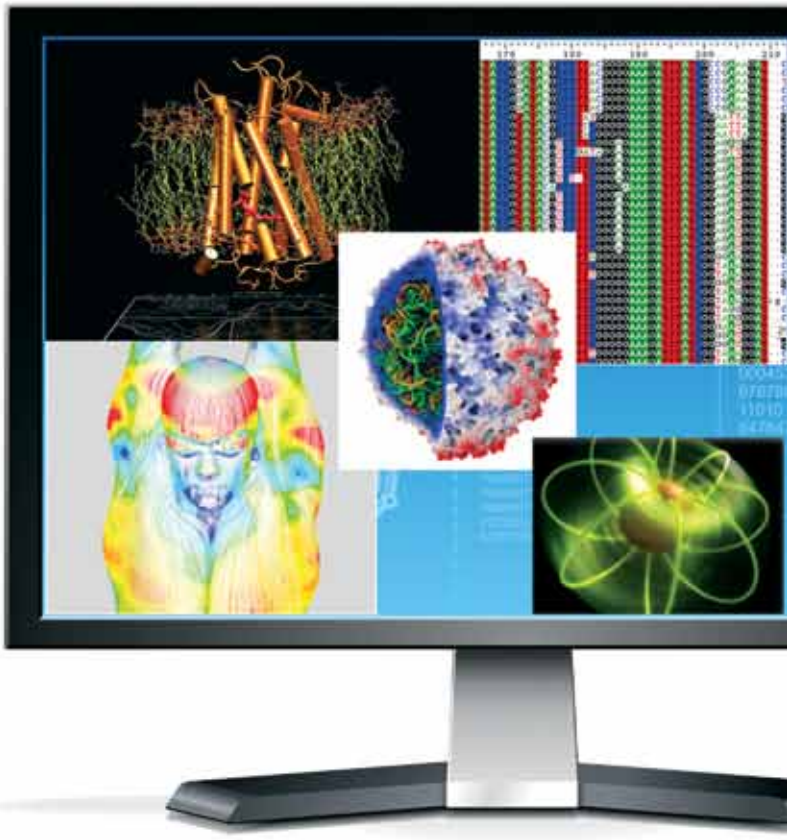

Designed by NVIDIA and delivered by NextIO, vCORE Express is a 1U computing system for high performance computing. vCORE Express delivers "must have" features for the technical and enterprise computing space including ECC memory for uncompromised accuracy and scalability, and 7X the double precision performance compared to NVIDIA Tesla 10-series GPU computing products. Compared to typical quad-core CPUs, vCORE Express computing systems deliver equivalent performance at 1/10th the cost and1/20th the power consumption.

Designed with four Fermi-based Tesla 2050 or 2070 computing processors in a standard 1U chassis, vCORE Express computing system scales to solve the world's most important computing challenges – more quickly and accurately.

For more information on vCORE Express, please contact sales@nextio.com, 1-877-7NEXTIO.

For a demonstration of vCORE Express and other NextIO vCORE™ products, please visit us at GTC Booth #63.

NextIO

WWW.NEXTIO.COM

The next generation of 1U workstation and workgroup GPU solutions developed by NVIDIA and delivered by NextIO.

THURSDAY, SEPT 23, 09:00 (50 MINUTES)
ROOM A5

2027 GPU-Based Image Processing in Military Applications

There are more than 6000 Unmanned Aerial Vehicles (UAVs) in use in the US Military. The US Army alone has flown more than 1 million UAV flight hours. Every UAV captures at least one stream of video; some as many as 9. All this video needs to be processed and analyzed both during the mission, and post-mission. Traditionally, custom ASICs, and FPGAs were required for even the most rudimentary image processing tasks. Now, GPUs provide orders of magnitude more compute at a fraction of the cost. Hear how MotionDSP uses GPUs to provide previously impossible capabilities to military imaging.

Speaker(s):	Sean Varah (CEO, MotionDSP Inc.)
Topic(s):	Video Processing, High Performance Computing

THURSDAY, SEPT 23, 09:00 (50 MINUTES)
ROOM L

2030 High-Throughput Cell Signaling Network Learning with GPUs

Explore how GPUs are being used to enable high-throughput cell signaling network discovery and data-intensive computational systems biology more generally. Systems biology is transitioning from a largely reductive discipline to one focused on building predictive models of large-scale biological systems. New instrumentation will provide the necessary raw data for such an approach, the key challenge now is building the hardware and software tools to efficiently and interactively build these models. This session will describe how GPUs can and will play a key role in these efforts.

Speaker(s):	Michael Linderman (Engineering Research Associate, Stanford University)
Topic(s):	Life Sciences, Algorithms & Numerical Techniques, Machine Learning & Artificial Intelligence

THURSDAY, SEPT 23, 09:00 (50 MINUTES)
ROOM A7

2033 Accelerating Pricing Models with virtual GPUs

Join Citadel to explore our three year undertaking on the feasibility of GPGPU computing for option pricing. We will discuss our 140X performance boost and the hurdles we had to overcome to integrate GPUs into our existing infrastructure. Please note that our talk will not get into the details of the model (that's proprietary information), but we will share our innovative solution to drive a grid of virtual GPUs.

Speaker(s):	Scott Donovan (System Architect, Citadel Investment Group)
Topic(s):	Finance, High Performance Computing

THURSDAY, SEPT 23, 09:00 (50 MINUTES)
ROOM C

2048 H.264/AVC Video Encoding with CUDA and OpenCL

Join experts from MainConcept, a leading provider of video codecs to the professional market, as they demonstrate the latest version of their CUDA-based H.264/AVC Encoder.

Speaker(s):	Thomas Kramer (VP Product Management, MainConcept)
Topic(s):	Video Processing, Tools & Libraries

THURSDAY, SEPT 23, 09:00 (50 MINUTES)
ROOM N

2076 Implementing CUDA Audio Networks

Learn how to implement a commercial software library that exploits CUDA for audio applications. We focus on the overall threading architecture and the underlying math for implementing general purpose audio processing in CUDA devices. Covers the use of inter-process communication to make a plug-in implementation loadable in 32 bit hosts installed in 64 bit systems, distributing the GPU load on remote servers, and creating a CUDA network for high-end purposes such as a big recording facility.

Speaker(s):	Giancarlo Del Sordo (Chief Developer and Product Manager, Acustica Audio)
Topic(s):	Audio Processing, Signal processing

THURSDAY, SEPT 23, 09:00 (50 MINUTES)
ROOM A3

2138 Faster, Cheaper, Better – Hybridization of Linear Algebra for GPUs

Learn how to develop faster, cheaper and better linear algebra software for GPUs through a hybridization methodology that is built on (1) Representing linear algebra algorithms as directed acyclic graphs where nodes correspond to tasks and edges to dependencies among them, and (2) Scheduling the execution of the tasks over hybrid architectures of GPUs and multicore. Examples will be given using MAGMA, a new generation of linear algebra libraries that extends the sequential LAPACK-style algorithms to the highly parallel GPU and multicore heterogeneous architectures.

Speaker(s):	Hatem Ltaief (Sr. Research Associate, University of Tennessee), Stan Tomov (Research Scientist, University of Tennessee)
Topic(s):	High Performance Computing, Algorithms & Numerical Techniques, Tools & Libraries

THURSDAY, SEPT 23, 09:00 (50 MINUTES)
ROOM K

2145 Photo Editing on the GPU with MuseMage

See how MuseMage greatly accelerates image processing and editing while providing real-time feedback by harnessing the power of GPUs. We will discuss the majority of MuseMage tools which are fully implemented on GPUs.

Speaker(s):	Kaiyong Zhao (Graduate Student, HKBU), Yubo Zhang (PhD student, UC Davis)
Topic(s):	Imaging

THURSDAY, SEPT 23, 09:00 (50 MINUTES)
MARRIOTT SAN JOSE BALLROOM

2156 GMAC: Global Memory For Accelerators

Learn how to use GMAC, a novel run-time for CUDA GPUs. GMAC unifies the host and device memories into a unified virtual address space, enabling the host code to directly access the device memory, and removing the need for data transfers between host and device memories. Moreover, GMAC also allows pointers to be used by both, the host and device code indistinctly.

This session will present the GMAC run-time and show how to use it in current applications. This session will cover from the basics of GMAC to multi-threaded applications using POSIX threads, OpenMP and MPI.

Speaker(s):	Isaac Gelado (Lecturer and Researcher, Universitat Politecnica de Catalunya)
Topic(s):	Tools & Libraries

THURSDAY, SEPT 23, 09:00 (50 MINUTES)
ROOM A1

2202 A Programming Model and Tool for Automatic High-Performance C to CUDA Mapping

Discover our automatic C-to-CUDA mapper prototype, and how it optimizes execution and data movement for a broad class of loop codes. Coupled with our powerful mapper, C as an input language does not only offer portability but also performance and performance portability. Learn about our optimizations and some of the performance obtained through different uses of the mapper.

Speaker(s):	Benoit Meister (Senior Engineer, Reservoir Labs)
Topic(s):	Tools & Libraries

THURSDAY, SEPT 23, 09:00 (20 MINUTES)
ROOM A8

2206 Accelerated Computational Fluid Dynamics Employing GPUs

Speaker(s):	Daniel Gaudlitz (Project Manager, FluiDyna)
Topic(s):	Computational Fluid Dynamics, High Performance Computing

THURSDAY, SEPT 23, 09:00 (50 MINUTES)
ROOM B

2236 A Work-Efficient GPU Algorithm for Level Set Segmentation

Explore a novel GPU level set segmentation algorithm that is both work-efficient and step-efficient. Our algorithm has O(logn) step-complexity, in contrast to previous GPU algorithms which have O(n) step-complexity. We apply our algorithm to 3D medical images and we show that in typical clinical scenarios, our algorithm reduces the total number of processed level set field elements by 16x and is 14x faster than previous GPU algorithms with no reduction in segmentation accuracy.

Speaker(s):	Mike Roberts (Research Assistant, Hotchkiss Brain Institute, University of Calgary, Canada)
Topic(s):	Medical Imaging & Visualization, Algorithms & Numerical Techniques, Computer Vision, Computer Graphics

THURSDAY, SEPT 23, 09:00 (50 MINUTES)
ROOM D

2272 GStream: A General-Purpose Data Streaming Framework on GPUs

We present GStream, a general-purpose, scalable and C++ template run-time framework amenable to both the streaming problem and GPU architectures. GStream offers transparent streaming data transmissions and automatic memory synchronization over a rich collection of computing resources that are transparently allocated and reused. Various problems other than streaming application, such as scientific computing, numerical codes and text processing, can be easily expressed using GStream and subsequently integrated with our GStream library. GStream's ease of use combined with efficient exploitation of GPU resources have the potential to lead to higher coding productivity and application performance through our data-centric specification paradigm.

Speaker(s):	Frank Mueller (Associate Professor, North Carolina State University), Xing Wu (Research Assistant, North Carolina State University)
Topic(s):	Tools & Libraries, High Performance Computing

THURSDAY, SEPT 23, 09:00 (20 MINUTES)
ROOM M

2278 Strategies for Code Encapsulation in GPU Implementations

Code encapsulation is a common technique used to reduce code complexity that a given programmer has to understand. It allows the use of increasingly complex systems of hardware, software, and algorithms to tackle increasingly difficult scientific problems. Unfortunately, code encapsulation is not easily attainable in current GPU environments. We will share our OpenCL development experiences for achieving partial encapsulation in GPU implementations, and discuss best practices in this area.

Speaker(s):	Brian Cole (Developer, OpenEye Scientific Software)
Topic(s):	Programming Languages & Techniques, High Performance Computing, Life Sciences

THURSDAY, SEPT 23, 09:00 (50 MINUTES)
ROOM A2

2301 GPU Cluster Computing: Accelerating Scientific Discovery

We propose holding a research roundtable focussed on using GPU clusters to support scientific research. The roundtable will bring together researchers that have recently deployed or are interested in deploying GPU clusters to enable scientific research. At the research roundtable they will be able to share their experiences in deploying this new technology and discuss the future of this technology in supporting research to tackle the world's most challenging scientific problems.

To open discussion we will provide a brief presentation about deployment of the CSIRO's latest supercomputer cluster, which is among the world's first to combine traditional CPUs with more powerful NVIDIA GPUs, that is providing a world class computational and simulation science facility to advance priority CSIRO science.

Speaker(s):	John Taylor (Science and Business Leader, CSIRO), Dragan Dimitrovici (XENON Systems Pty Ltd)
Topic(s):	High Performance Computing

THURSDAY, SEPT 23, 09:00 (50 MINUTES)
KEYNOTE HALL

4006 Fireside Chat with Jen-Hsun Huang - Co-founder & CEO, NVIDIA

Jen-Hsun Huang will take part in a fireside chat by Quentin Hardy, National Editor at Forbes Magazine. They will discuss the rise of GPUs, current trends in visual and parallel computing, and the transformational changes ahead for the industry.

Speaker(s):	Quentin Hardy (National Editor, Forbes Magazine), Jen-Hsun Huang (CEO & President, NVIDIA)
Topic(s):	General Interest

THURSDAY, SEPT 23, 09:30 (20 MINUTES)
ROOM A8

2037 Numtech & GPGPU, a SME Point of View

Hear why and how Numtech, a french SME working in the field of atmospheric dispersion and expertise of meteorological events, is

benchmarking GPGPU for its futures applications. A compressible and an incompressible interactive flow solvers are described.

Speaker(s):	Emmanuel Buisson (CEO, Numtech)
Topic(s):	Computational Fluid Dynamics, Physics Simulation

THURSDAY, SEPT 23, 10:00 (50 MINUTES)
ROOM L

2001 Acceleration of the Freesurfer Suite for Neuroimaging Analysis

See how GPU technology has dramatically accelerated the Freesurfer suite of tools used by thousands of researchers for the analysis of neuroimaging data.

Speaker(s):	Richard Edgar (Assistant in Neuroscience, Massachusetts General Hospital, Harvard University)
Topic(s):	Medical Imaging & Visualization, Imaging, Tools & Libraries

THURSDAY, SEPT 23, 10:00 (50 MINUTES)
ROOM A3

2002 CUDA Debugging on Linux and MacOS with cuda-gdb

Boost your development speed by mastering the CUDA debugging tools NVIDIA provides. In this session you will learn the basics of cuda-gdb and cuda-memcheck, as well as their more advanced features with live demonstrations on Linux and MacOS.

Speaker(s):	Satish Salian (Manager CUDA Debugger Tools, NVIDIA)
Topic(s):	Tools & Libraries

THURSDAY, SEPT 23, 10:00 (50 MINUTES)
ROOM A7

2032 Practical Methods Beyond Monte Carlo in Finance

Murex will share its practical experience using GPUs to accelerate high-performance analytics based on GPU-enabled Monte Carlo and PDE methods. We will also briefly describe Murex's experience developing a high-level payoff scripting language that allows user-definable payoffs for single and cross-asset instruments.

Speaker(s):	Pierre Spatz (Head of Quantitative Research, Murex SAS)
Topic(s):	Finance, Algorithms & Numerical Techniques

THURSDAY, SEPT 23, 10:00 (50 MINUTES)
ROOM K

2053 Pixel Bender: Building a Domain Specific Language on the GPU

Examine the challenges and advantages of building the Pixel Bender domain specific language for image processing for the GPU. We will examine how Pixel Bender was made to work within several Adobe applications across a wide range of hardware systems and platforms.

Speaker(s):	Bob Archer (Senior Computer Scientist, Adobe Systems Inc)
Topic(s):	Tools & Libraries

THURSDAY, SEPT 23, 10:00 (50 MINUTES)
ROOM A2

2055 Application of Fermi GPU to Flow Cytometry and Cancer Detection

Learn how a Tesla C2050 enabled scientists to explore cancer data sets 400 times faster than a PC-only implementation. Discusses how the results of this work may lead to better diagnostics for detecting leukemia in blood cells.

Speaker(s):	Robert Zigon (Sr Staff Development Engineer, Beckman Coulter)
Topic(s):	Life Sciences

THURSDAY, SEPT 23, 10:00 (20 MINUTES)
ROOM A8

2110 Acceleration of a Novel Rotorcraft Wake Simulation

Dive deep as we present the details of a new CUDA-based algorithm for accurate rotorcraft wake simulations. We use a vortex particle method, accelerated with a multipole tree algorithm, combined with a traditional grid-based CFD code. This CUDA algorithm can evaluate the velocity and velocity-gradient with an effective throughput approaching 300 billion interactions per second on a C1060. This gives 10x speed-up and 2.5x better accuracy compared to the parallel CPU version.

Speaker(s):	Christopher Stone (Research Scientist, Intelligent Light)
Topic(s):	Computational Fluid Dynamics, Algorithms & Numerical Techniques

THURSDAY, SEPT 23, 10:00 (50 MINUTES)
ROOM N

2116 Real-time Multichannel Audio Convolution

Learn how a synthesis of 3D sound scenes can be achieved using a peer-to-peer music streaming environment and GPU. We will discuss the technical and cost benefits to this approach, while noting that it frees the CPU for other tasks.

Speaker(s):	Jose Antonio Belloch (MsC, Institute of Telecommunications and Multimedia Applications, Universidad Politecnica de Valencia), Alberto Gonzalez (Professor, Universidad Politecnica de Valencia), Antonio M. Vidal (Professor, Universidad Politecnica de Valenc)
Topic(s):	Audio Processing, Signal processing

THURSDAY, SEPT 23, 10:00 (50 MINUTES)
ROOM A5

2124 Operating System Abstractions for GPU Programming

GPGPU frameworks such as CUDA improve programmability, but GPU parallelism remains inaccessible in many application domains. This session argues that poor OS support causes this problem. OSes do not provide the kind of high-level abstractions for GPUs that applications expect for other resources like CPUs and file systems. We advocate reorganizing kernel abstractions to support GPUs as first-class computing resources, with traditional guarantees such as fairness and isolation. We demonstrate shortcomings in Windows 7 GPU support, and show that better OS abstractions can accelerate interactive workloads like gesture recognition by a factor of 10X over a CUDA implementation.

Speaker(s):	Christopher Rossbach (Researcher, Microsoft Research), Emmett Witchel (Professor, University of Texas at Austin)
Topic(s):	Programming Languages & Techniques, Tools & Libraries

THURSDAY, SEPT 23, 10:00 (20 MINUTES)
ROOM B

2149 Overview of Parallel Nsight for Visual Studio

NVIDIA Parallel Nsight provides access to the power of the GPU from within the familiar environment of Microsoft Visual Studio. This session is an entry level overview of the GPU computing and graphics development features of Parallel Nsight as well as a glimpse into the future of this powerful tool.

Speaker(s):	Kumar Iyer (Product Manager, NVIDIA)
Topic(s):	Tools & Libraries

THURSDAY, SEPT 23, 10:00 (50 MINUTES)
ROOM A1

2176 Easy GPU Meta-programming: A Case Study in Biologically-Inspired Computer Vision

Learn how to let the computer optimize your CUDA and OpenCL code for you with easy GPU Meta-programming and Scripting (e.g. PyCUDA). We will present a case study in which we consider the step-wise optimization of a 3D filter bank convolution, using a suite of open-source tools.

Speaker(s):	Nicolas Pinto (PhD Student, MIT)
Topic(s):	Tools & Libraries, Computer Vision, High Performance Computing, Neuroscience

THURSDAY, SEPT 23, 10:00 (50 MINUTES)
ROOM C

2215 Extending OpenCV with GPU Acceleration

OpenCV is a widely popular computer vision library, with millions of downloads and hundreds of thousands of users. Applications span many industries including robotics, industrial machine vision, automotive, film & broadcast, medical, and consumer applications. NVIDIA and the OpenCV development team are collaborating to provide CUDA implementations of the most demanding algorithms, thus enabling a new level of real-time capability and higher quality results.

This talk with introduce OpenCV, and summarize the new CUDA enabled capabilities, and provide an overview of future plans.

Speaker(s):	Joe Stam (Sr. Applications Engineer, NVIDIA)
Topic(s):	Computer Vision, Imaging, Stereoscopic 3D, Video Processing

THURSDAY, SEPT 23, 10:00 (50 MINUTES)
MARRIOTT SAN JOSE BALLROOM

2269 Bringing GPUs to Mainstream Molecular Dynamics Packages

Recent work in close collaboration with NVIDIA has produced a GPU accelerated version of the AMBER Molecular Dynamics Code PMEMD that runs between 20 and 130 times the speed of a single 2.8GHz Intel Nehalem Processor, with even higher performance on multiple GPUs, but which does not make sacrifices in the accuracy or validity of such calculations to achieve this. The GPU accelerated version supports both explicit solvent particle mesh ewald (PME) and implicit solvent simulations and is available as part of the new AMBER 11 package. This talk will provide an overview of the AMBER software, background behind this GPU work, benchmarks, the impact that GPU accelerated MD can have on the field, the techniques used to achieve the performance seen without sacrificing accuracy and finally the validation methods used to ensure simulations are directly equivalent to CPU based calculations. Ensuring that a GPU implementation of a MD package provides results that are indistinguishable from the CPU code is extremely tricky and often the desire to take shortcuts to boost performance can affect accuracy with unpredictable results. We have developed a comprehensive validation suite that can be used to perform the detailed testing that is required to ensure the approximations necessary for GPU performance do not impact the scientific results. Additionally we will discuss how we have made careful use of mixed single and double precision arithmetic in the AMBER implementation to achieve equivalence in the results without excessively compromising performance. Finally we provide examples of recent breakthrough simulations conducted using GPU enabled AMBER 11.

Speaker(s):	Ross Walker (Research Professor, San Diego Supercomputer Center)
Topic(s):	Molecular Dynamics

THURSDAY, SEPT 23, 10:00 (50 MINUTES)
KEYNOTE HALL

4007 Emerging Companies: CEO on Stage featuring Aqumin, RTT, and Scalable Display Technologies

See the hottest new technologies from startups that are transforming computing.

In a lively and fast-paced exchange, the "Emerging Companies Summit - CEO on Stage" sessions will feature CEOs from three startups who will each have 15 minutes to introduce their companies and interact with a panel of leading venture capitalists, technology executives, and industry analysts.

Panelist(s):	Norman Winarsky (VP of Ventures, Licensing and Strategic Programs, SRI), Savitha Srinivasan (Corporate Venture Partner, IBM), and Rob Enderle (Analyst, Enderle Group), Jeff Herbst (Vice President of Business Development, NVIDIA
Speaker(s):	Andrew Jamison (CEO, Scalable Display Technologies), Jeroen Snepvangers (CEO, RTT), Michael Zeitlin (CEO, Aqumin)
Topic(s):	General Interest, Finance, Imaging, Computer Graphics

THURSDAY, SEPT 23, 10:30 (20 MINUTES)
ROOM A8

2061 Accelerating Explicit FEM Shock & Blast Simulations

Explicit finite element codes are widely used to simulate the response of structures and mechanical equipment subjected to shock, blast and wave propagation phenomena. High resolution models require run times ranging from a few seconds to a few months are common and hence the payoff from GPU acceleration is tremendous. We describe the acceleration of our commercial finite element code NLFLEX using CUDA. We developed GPU kernels in CUDA based on our production code NLFLEX, for linear elasticity, explosives, elasto-plasticity and large deformation elasticity. We attained order of magnitude (10X) acceleration in single precision and approximately (5X) in double precision mode.

Speaker(s):	Nachiket Gokhale (Senior Research Engineer, Weidlinger Associates Inc)
Topic(s):	Algorithms & Numerical Techniques, Computational Fluid Dynamics, Physics Simulation

THURSDAY, SEPT 23, 10:30 (20 MINUTES)
ROOM B

2292 Implementation of High-Order Adaptive CFD Methods on GPUs

A discontinuous high-order formulation named the Correction Procedure via Reconstruction (CPR) is recently implemented on Nvidia GPUs. The CPR formulation is related to the discontinuous Galerkin (DG) method, and unifies several methods such as the DG, spectral volume and spectral difference into a single framework efficient for hybrid meshes. In preliminary 2D inviscid flow computations, a single GPU has been able to deliver a speedup of 44 over a CPU of the same generation. Extension is being made for viscous flow computation, and results will be presented at the final presentation.

Speaker(s):	Arun Somani (Anson Marston Professor, Iowa State University), Z.J. Wang (Professor, Iowa State University), Lizandro Solano (Iowa State University)
Topic(s):	Computational Fluid Dynamics

THURSDAY, SEPT 23, 11:00 (50 MINUTES)
ROOM B

2007 Folding@home: Petaflops on the Cheap Today; Exaflops Soon?

Learn how Folding@home has used petascale computing with GPUs to make fundamental breakthroughs in computational biology and how this technology can make an impact in your work.

Speaker(s):	Vijay Pande (Director, Folding@home Distributed Computing Project, Stanford University)
Topic(s):	Life Sciences, Cloud Computing, High Performance Computing, Molecular Dynamics

THURSDAY, SEPT 23, 11:00 (50 MINUTES)
ROOM A5

2023 Processing Device Arrays with C++ Metaprogramming

I will describe tricks for building APIs using C++ metaprogramming that generate custom kernels for complex manipulation of device-side arrays in CUDA. Using a variation of Expression Templates, multiple operations can be fused into a single kernel that executes with reasonable efficiency.

Speaker(s):	Jonathan Cohen (Senior Research Scientist, NVIDIA Research)
Topic(s):	Programming Languages & Techniques, Tools & Libraries

THURSDAY, SEPT 23, 11:00 (50 MINUTES)
ROOM N

2042 Interactive 3D Audio Rendering Systems

Learn how to leverage GPUs for interactive audio rendering. This session will give a short overview of the architecture of current GPUs, emphasizing some key differences between GPU and CPUs programming models for audio processing. We will illustrate the benefits of GPU-accelerated audio rendering with results from 3D audio processing and sound scattering simulations. Finally, we will discuss best practices for GPU implementations as well as future opportunities for audio rendering on massively parallel architectures.

Speaker(s):	Nicolas Tsingos (Senior Staff Engineer, Dolby Laboratories)
Topic(s):	Audio Processing, Ray Tracing, Signal processing

THURSDAY, SEPT 23, 11:00 (50 MINUTES)
ROOM L

2043 Disparity Map Generation

Explore the algorithms and implementation of disparity maps on the GPU. We will discuss how a disparity map facilitates stereoscopic content creation, applications and approaches tried, and final results of real time calculations on GPUs.

Speaker(s):	Henry Gu (CTO, GIC)
Topic(s):	Stereoscopic 3D, Computer Vision, Imaging

THURSDAY, SEPT 23, 11:00 (50 MINUTES)
ROOM K

2051 GPGPU in Commercial Software: Lessons From Three Cycles of the Adobe Creative Suite

Learn about leveraging GPUs for commercial software. We will discuss lessons learned creating and using the Adobe Image Foundation libraries to accelerate image and video processing using GPUs and multi-core. These libraries are used by most of Adobe's applications as well as integrated by hobbyist and professional applications with different levels of experience with GPUs and diverse user bases.

Speaker(s):	Kevin Goldsmith (Senior Engineering Manager, Adobe Systems, Incorporated)
Topic(s):	Imaging, Video Processing

THURSDAY, SEPT 23, 11:00 (50 MINUTES)
ROOM A3

2070 CUSPARSE Library: A Set of Basic Linear Algebra Subroutines for Sparse Matrices

The CUSPARSE library can impact and enable software solutions for computational science and engineering problems in the fields of energy exploration, physical simulations and life sciences among many others. It provides sparse linear algebra primitives that can be used to implement iterative linear system and eigenvalue solvers and can also serve as a building block for the state-of-the-art sparse direct solvers. CUSPARSE library is implemented using CUDA parallel programming model and provides sparse analogs to BLAS level-1,2,3 operations, such as matrix-vector multiplication, triangular solve and format conversion routines.

Speaker(s):	Maxim Naumov (Software Engineer, NVIDIA)
Topic(s):	Tools & Libraries, Algorithms & Numerical Techniques, High Performance Computing

THURSDAY, SEPT 23, 11:00 (50 MINUTES)
ROOM A1

2075 GPU-Accelerated Video Encoding

Learn how to accelerate video encoding using the GPU. We will give an overview of the typical video encoding pipeline and discuss how different parts of the pipeline can be ported to GPU using various approaches. We will focus on block-based Motion Estimation, in particular, as it is the corner stone of video encoding algorithms. The efficiency of its implementation on the GPU is crucial to the speed and quality of the encoder.

Speaker(s):	Anton Obukhov (Developer Technology Engineer, NVIDIA)
Topic(s):	Video Processing

THURSDAY, SEPT 23, 11:00 (50 MINUTES)
ROOM A7

2098 Enabling On Demand Value-At-Risk for Financial Markets

Learn how financial market risk managers can increase their ability to preempt exposure limit breaching and tighten risk control to increase investor confidence. Gain insight into the techniques for obtaining high performance Monte-Carlo based market value-at-risk (VaR) estimates over a hierarchy of risk aggregation levels. This session will focus on how the new Fermi platform can be used by financial institutions to enable on-demand estimates of the market VaR, and discuss important software architecture decisions, the benefits of the new GigaThread Engine and Parallel DataCache, as well as the guiding principles for constructing efficient algorithms on GPUs.

Speaker(s):	Matthew Dixon (Professor, UC Davis), Jike Chong (Principal Software Architect, Parasians, LLC)
Topic(s):	Finance, Algorithms & Numerical Techniques

THURSDAY, SEPT 23, 11:00 (20 MINUTES)
ROOM A8

2171 Parallel Algorithms for Interactive Mechanical CAD

The broad objective of our research is to develop mechanical Computer-Aided Design tools that provide interactive feedback to the designer. We have developed GPU algorithms for fundamental CAD operations (NURBS evaluation, surface-surface intersection, separation distance computation, moment computation, etc.) that are one to two orders of magnitude faster, and often more accurate, than current commercial CPU implementations. We will touch on strategies we have employed to meet GPU programming challenges, such as the separation of CPU/GPU operations, imposing artificial structure on computations, and transforming problem definitions to suit GPU-computation models.

Speaker(s):	Sara McMains (Associate Professor, University of California Berkeley), Adarsh Krishnamurthy (Student, University of California Berkeley)
Topic(s):	Algorithms & Numerical Techniques, Tools & Libraries, Computer Graphics

THURSDAY, SEPT 23, 11:00 (50 MINUTES)
ROOM C

2173 Enabling Large-Scale CCTV Face Recognition

Learn how to use CUDA and GPGPU to perform large scale face search for both forensics as well as CCTV face recognition.

Speaker(s):	Ben Lever (Senior Research Engineer, NICTA), Abbas Bigdeli (Senior Researcher and Technology Manager, NICTA)
Topic(s):	Computer Vision, Video Processing

THURSDAY, SEPT 23, 11:00 (50 MINUTES)
ROOM A2

2203 Modeling Evolution Computing the Tree of Life

Learn how GPUs are being used to accelerate our understanding of the tree of life. This session will cover BEAGLE, which is an open API and library for evaluating phylogenetic likelihoods of biomolecular sequence evolution. BEAGLE uses novel algorithms and methods for evaluating phylogenies under arbitrary molecular evolutionary models on GPUs, making use of the large number of processing cores to efficiently parallelize calculations.

Speaker(s):	Daniel Ayres (PhD Candidate, University of Maryland)
Topic(s):	Life Sciences

THURSDAY, SEPT 23, 11:00 (50 MINUTES)
MARRIOTT SAN JOSE BALLROOM

2219 High-Productivity CUDA Development with the Thrust Template Library

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 session we'll discuss Thrust's features and explain the basic design philosophy of the library.

Speaker(s):	Nathan Bell (Research Scientist, NVIDIA Research)
Topic(s):	Tools & Libraries

THURSDAY, SEPT 23, 11:00 (50 MINUTES)
KEYNOTE HALL

4008 Emerging Companies: CEO on Stage featuring ICD and Universal Robotics

See the hottest new technologies from startups that are transforming computing.

In a lively and fast-paced exchange, the "Emerging Companies Summit - CEO on Stage" sessions will feature CEOs from three startups who will each have 15 minutes to introduce their companies and interact with a panel of leading venture capitalists, technology executives, and industry analysts.

Panelist(s):	Rob Enderle (Analyst, Enderle Group), Jeff Herbst (Vice President of Business Development, NVIDIA), Savitha Srinivasan (Corporate Venture Partner, IBM), Norman Winarsky (VP of Ventures, Licensing and Strategic Programs, SRI)
Speaker(s):	David Peters (Founder and CEO, Universal Robotics), David Hayes (CEO, ICD)
Topic(s):	General Interest, Machine Learning & Artificial Intelligence, Mobile Devices

THURSDAY, SEPT 23, 11:30 (20 MINUTES)
ROOM A8

2106 Particleworks: Particle-based CAE Software on Multi-GPU

Prometech Software, Inc. is an university launched technology venture in Japan and has been working in the field of particle-based computational fluid dynamics for several years. Through collaboratinos with major automotive and material companies in Japan, Prometech has implemented our Particle technology on Multi-GPU and delivered as a CAE software, "Particleworks". In this session, we will discuss the theoretical background of our simulation (MPS; Moving Particle Simulation method), Multi GPU programming techniques of sparse matrix solver, performance results of Particleworks and the analysis examples of the Auto and Material.

Speaker(s):	Issei Masaie (Chief GPU Engineer, Prometech Software, Inc.)
Topic(s):	Computational Fluid Dynamics, High Performance Computing

THURSDAY, SEPT 23, 11:30 (20 MINUTES)
ROOM D

2298 Accelerated Image Quality Assessment using Structural Similarity

Explores the GPU porting and performance analysis of the image quality assessment algorithm based on structural similarity index(SSI). This index is a powerful tool for image quality assessment and the algorithm is highly suitable for GPU architecture, offering a rapid image quality assessment in many image restoration applications.

Speaker(s):	Mahesh Khadtare (Member, Technical Staff, Computational Research)
Topic(s):	Computer Vision, Imaging

THURSDAY, SEPT 23, 12:00 (120 MINUTES)
EXHIBIT HALL

1004 Exhibits Open / Networking Lunch

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.

Topic(s):	General Interest

THURSDAY, SEPT 23, 14:00 (50 MINUTES)
ROOM A7

2040 Derivatives & Bond Portfolio Valuation in a Hybrid CPU/GPU Environment

Learn how to compute traditional end of day computations in real time through the use of a hybrid GPU/CPU computing environment. We will detail how computing intensive tasks are delegated to the GPU while interface issues are dealt with by the CPU. We will discuss our methodology consisting of the following three components: (1) valuations; (2) by tenor risk measures; and (3) full distributions allowing for more complex analytics such as exotic options products valuation and counterparty value adjustments calculation.

Speaker(s):	Peter Decrem (Director, Rates Products, Quantifi)
Topic(s):	Finance, Algorithms & Numerical Techniques, High Performance Computing

THURSDAY, SEPT 23, 14:00 (50 MINUTES)
MARRIOTT SAN JOSE BALLROOM

2054 NAMD, CUDA, and Clusters: Taking GPU Molecular Dynamics Beyond the Desktop

A supercomputer is only as fast as its weakest link. The highly parallel molecular dynamics code NAMD was one of the first codes to run on a GPU cluster when G80 and CUDA were introduced in 2007. Now, after three short years, the Fermi architecture opens the possibility of new algorithms, simpler code, and easier optimization. Come learn the opportunities and pitfalls of taking GPU computing to the petascale.

Speaker(s):	James Phillips (Senior Research Programmer, University of Illinois)
Topic(s):	Molecular Dynamics, High Performance Computing, Life Sciences, Physics Simulation

THURSDAY, SEPT 23, 14:00 (50 MINUTES)
ROOM K

2087 Fast High-Quality Panorama Stitching

We present a panorama stitching application implemented with CUDA C on the GPU. The image processing pipeline consist of SIFT feature detection and matching and Graphcut image stitching to achieve high-quality results. We demonstrate live panorama creation with a Webcam.

Speaker(s):	Timo Stich (Developer Technology Engineer, NVIDIA)
Topic(s):	Video Processing, Algorithms & Numerical Techniques, Computer Vision, Imaging

THURSDAY, SEPT 23, 14:00 (50 MINUTES)
ROOM A2

2115 Modified Smith-Waterman-Gotoh Algorithm for CUDA Implementation

It is axiomatic that computational throughput can be increased by exploiting the parallelism of GPU hardware — but what if the computational algorithm is not easy to implement in parallel? We have modified one such algorithm — the Smith-Waterman-Gotoh dynamic programming algorithm for local sequence alignment — so as to make it more amenable to data-parallel computation. The result is a successful CUDA implementation that fully exploits GPU parallelism.

Speaker(s):	Richard Wilton (Research Scientist, The Johns Hopkins University)
Topic(s):	Life Sciences, Algorithms & Numerical Techniques

THURSDAY, SEPT 23, 14:00 (50 MINUTES)
ROOM L

2121 Maximizing Throughput of Barco's GPU-Enabled Video Processing Server

Find out how Imec middleware realizes the full potential of GPU-enabled video processing servers to manage multiple video processing pipelines. We will discuss how the middleware monitors GPU and CPU execution to best balance the load. Covers how we achieved a 30% increase in throughput with only a minimal 0.05% overhead on Barco's GPU-enabled video processing server.

Speaker(s):	Maja D'Hondt (Program Manager, imec)
Topic(s):	Video Processing, Tools & Libraries

THURSDAY, SEPT 23, 14:00 (20 MINUTES)
ROOM A8

2155 GPGPU in the real world. The ABAQUS experience

We describe the ABAQUS experience in integrating GPGPU acceleration into a complex, high performance commercial engineering software. In particular we discuss the trade-off we had to make and the benefits we obtained from this technology.

Speaker(s):	Luis Crivelli (Director Solver Development, Dassualt Systems Simulia Corporation)
Topic(s):	Physics Simulation, Algorithms & Numerical Techniques, Computational Fluid Dynamics, High Performance Computing

THURSDAY, SEPT 23, 14:00 (50 MINUTES)
ROOM N

2175 Hello GPU: High-Quality, Real-Time Speech Recognition on Embedded GPUs

In this presentation, we will talk about our experiences of implementing an end-to-end automatic speech recognition system that runs in faster than real-time on embedded GPUs, targeted towards small form-factor consumer devices. Focusing specifically on some of the challenges encountered during the design process, a major portion of our talk will focus on giving insights into modifications we made to well-established speech algorithms to fit well within the GPU programming model. We will show how these changes helped us in realizing a highly optimized system on platforms with limited memory bandwidth and compute resources.

Speaker(s):	Kshitij Gupta (Graduate Student Researcher, UC Davis)
Topic(s):	Embedded & Automotive, Audio Processing, Signal processing, Mobile & Tablet & Phone

THURSDAY, SEPT 23, 14:00 (50 MINUTES)
ROOM C

2209 Accelerating Computer Vision on the Fermi Architecture

GPUS have evolved from fixed function to general purpose, and continue to evolve with new features being added in every generation. This talk will discuss how to exploit the new features introduced by the Fermi architecture (such as concurrent kernel execution, writes to texture) to accelerate computer vision algorithms.

Speaker(s):	James Fung (Developer Technology, NVIDIA)
Topic(s):	Computer Vision, Tools & Libraries

THURSDAY, SEPT 23, 14:00 (50 MINUTES)
ROOM A3

2210 GPU-Ocelot: An Open Source Debugging and Compilation Framework for CUDA

Learn how to debug and profile CUDA applications using GPU-Ocelot. Ocelot is a compilation and emulation framework for CUDA that includes debugging and profiling tools as well as backend compilers for NVIDIA GPUs and x86 CPUs. We will present examples of applications developed on x86 CPUs and deployed on NVIDIA GPUs. We will also discuss memory checking, race detection, and deadlock detection tools available within Ocelot.

Speaker(s):	Gregory Diamos (PhD Student, Georgia Institute of Technology), Andrew Kerr (PhD Student, Georgia Institute of Technology), Sudhakar Yalamanchili (Professor, Georgia Institute of Technology)
Topic(s):	Tools & Libraries

THURSDAY, SEPT 23, 14:00 (50 MINUTES)
ROOM B

2220 Thrust by Example: Advanced Features and Techniques

Thrust is a parallel template library for developing CUDA applications which is modeled after the C++ Standard Template Library (STL). In this session we'll show how to implement decompose problems into the algorithms provided by Thrust. We'll also discuss the performance implications of "kernel fusion" and "array of structs" vs. "structure of arrays" memory layouts and how they relate to Thrust. Lastly, we'll present evidence that Thrust implementations are fast, while remaining concise and readable.

Speaker(s):	Jared Hoberock (Research Scientist, NVIDIA)
Topic(s):	Tools & Libraries

THURSDAY, SEPT 23, 14:00 (50 MINUTES)
ROOM A1

2241 Standing Out: Implementing a Great Stereo UI

Learn how to make S3D compatible user interfaces, HUDs, and in-game menus. The first part of this session will outline the common problems users encounter when displaying traditional 2D UI in stereoscopic 3D. The second part will focus on the different techniques, tips/tricks, and best practices developers can use to create high-quality S3D interfaces. The presentation will highlight examples from several shipped titles, as well as showcase a complete 3D UI game demo running in S3D on multiple devices including PC and mobile.

Speaker(s):	Brendan Iribe (President, Scaleform)
Topic(s):	Stereoscopic 3D, Tools & Libraries, Computer Graphics, Mobile & Tablet & Phone

THURSDAY, SEPT 23, 14:00 (50 MINUTES)
KEYNOTE HALL

4009 Emerging Companies Summit Panel: The "New Normal" For Building Emerging Companies Based On Disruptive Technologies

Moderated by Jeff Herbst – Vice President of Business Development, NVIDIA

Start-ups are facing unique challenges as aresult of the current economic and business environment. Not only is the venture funding environment very difficult, but small companies are finding it increasingly difficult to "break out" of the pack through IPO's and attractive M&A exits. This panel of experts (which includes VC and corporate investors) will attempt to assess the current state of both the public and private markets, and will explore various strategies and options for building successful companies in this "new" environment. Topics will include traditional forms of equity and debt, angel financing, as well as other creative/strategic financing options (eg. NRE arrangements, strategic partnerships etc.). The discussing promises to be both lively and provocative.

Panelist(s):	Gerald Brady (Managing Director, Silicon Valley Bank), Bill Frauenhofer (Managing Director, Citigroup Global Markets), Garrett Herbert (Partner, M&A Transaction Services, Deloitte & Touche LLP), Eric Jensen (Partner, Business Department Chair, Cooley LLP), Andrew T. Sheehan (Managing Director, Sutter Hill Ventures)
Topic(s):	Finance, General Interest

THURSDAY, SEPT 23, 14:30 (20 MINUTES)
ROOM A8

2240 Accelerating LS-DYNA with MPI, OpenMP, and CUDA

When solving implicit problems, the computational bottleneck in LS-DYNA is the multifrontal linear solver. These operations are performed with double precision arithmetic, hence until the arrival of the Tesla 2050, experiments with GPU acceleration were only a curiosity. This is no longer the case, and in this talk we will describe how LS-DYNA's hybrid (MPI and OpenMP) solver is further accelerated using GPUs to factor large dense frontal matrices.

Speaker(s):	Bob Lucas (Computational Sciences Division Director, University of Southern California)
Topic(s):	High Performance Computing, Algorithms & Numerical Techniques

THURSDAY, SEPT 23, 15:00 (50 MINUTES)
ROOM K

2003 Using CUDA to Accelerate Radar Image Processing

Come see how current GPU technology provides the means for the first portable real-time radar image processing algorithm. This session will outline how the GPU has afforded nearly three orders of magnitude improvement in performance for Synthetic Aperture Radar's (SAR) hallmark image processing algorithm. We will present algorithm details and further improvements.

Speaker(s):	Aaron Rogan (Research Scientist and System Adminstrator, Neva Ridge Technologies)
Topic(s):	Signal processing, Algorithms & Numerical Techniques, Imaging, Video Processing

THURSDAY, SEPT 23, 15:00 (120 MINUTES)
MARRIOTT SAN JOSE BALLROOM

2012 Analysis-Driven Performance Optimization

The goal of this session is to demystify performance optimization by transforming it into an analysis-driven process. There are three fundamental limiters to kernel performance: instruction throughput, memory throughput, and latency. In this session we will describe:

•how to use profiling tools and source code instrumentation to assess the significance of each limiter;

•what optimizations to apply for each limiter;

•how to determine when hardware limits are reached.

Concepts will be illustrated with some examples and are equally applicable to both CUDA and OpenCL development. It is assumed that attendees are already familiar with the fundamental optimization techniques.

Speaker(s):	Paulius Micikevicius (NVIDIA)
Topic(s):	Tools & Libraries

THURSDAY, SEPT 23, 15:00 (50 MINUTES)
ROOM A1

2016 VDPAU: PureVideo on Unix

Learn about VDPAU (Video Decode and Presentation API for Unix). VDPAU provides GPU-accelerated video decoding, post-processing, UI compositing, and display on Unix. VDPAU also supports sharing surfaces with OpenGL and CUDA ("interop"). This allows developers to implement their own post-processing algorithms or scene analysis, or to use decoded video surfaces as part of a scene rendered using OpenGL.

Speaker(s):	Stephen Warren (Snr Linux Software Engineer, NVIDIA)
Topic(s):	Video Processing, Tools & Libraries

THURSDAY, SEPT 23, 15:00 (50 MINUTES)
ROOM N

2046 Efficient Automatic Speech Recognition on the GPU

Learn about how the GPU is able to meet the challenges of implementing automatic speech recognition (ASR), gain insights into the data-parallel implementation techniques that can provide 10x faster performance compared to sequentially processing ASR on a CPU. The state-of-art algorithm for ASR performs a graph traversal on a large, irregular graph with millions of states and arcs, guided by speech input only known at runtime. We present four generalizable techniques including: dynamic data-gather buffer, find-unique, lock-free data structures using atomics, and hybrid global/local task queues. When used together, these techniques can effectively resolve ASR implementation challenges on a GPU.

Speaker(s):	Jike Chong (Principal Software Architect, Parasians, LLC)
Topic(s):	Machine Learning & Artificial Intelligence, Algorithms & Numerical Techniques, Audio Processing

THURSDAY, SEPT 23, 15:00 (50 MINUTES)
ROOM A5

2062 HOOMD-blue: Fast and Flexible Many-Particle Dynamics

See the newest capabilities and performance enhancements in HOOMD-blue, a general-purpose many-particle dynamics application written for GPUs. Speedups of 80-100x are attained for a wide range of simulation types. Topics for this presentation include an overview of HOOMD-blue, design and implementation details of the underlying algorithms, and a discussion on how generality is maintained without sacrificing performance.

Speaker(s):	Joshua Anderson (Research Area Specialist, University of Michigan)
Topic(s):	Molecular Dynamics, High Performance Computing, Life Sciences, Physics Simulation

THURSDAY, SEPT 23, 15:00 (20 MINUTES)
ROOM A7

2064 Correlated Paths for Monte Carlo Simulations

Learn how the GPU can be deployed to generated correlated paths for Monte Carlo simulation. Using Asian Basket options as an example, the session shows the generation of correlated paths with a local volatility model for each of the underlying assets. Once the paths have been computed, the payoff in each scenario is computed and reduced to determine the expected value, all on the GPU.

Speaker(s):	Thomas Bradley (Developer Technology Engineer, NVIDIA)
Topic(s):	Finance

THURSDAY, SEPT 23, 15:00 (50 MINUTES)
ROOM B

2081 Morphing a GPU into a Network Processor

Modern Internet routers must meet two conflicting objectives, high performance and good programmability, to satisfy the ever-increasing bandwidth requirements under fast changing network protocols. A few recent works prove that GPUs have great potential to serve as the packet processing engine for software routers. However, current GPU's batched execution model cannot guarantee quality-of-service (QoS) requirement. In this work, we show how to convert a GPU into an effective packet processor through minimal changes in both hardware architecture and scheduling mechanism. Experimental results proved that the new GPU architecture could meet stringent QoS requirements, but maintain a high processing throughput.

Speaker(s):	Yangdong Deng (Associate Professor, Tsinghua University)
Topic(s):	General Interest, High Performance Computing

THURSDAY, SEPT 23, 15:00 (50 MINUTES)
ROOM A2

2105 CUDA-FRESCO: An Efficient Algorithm for Mapping Short Reads

Learn about CUDA-FRESCO and how it addresses issues with MUMmerGPU. We will detail how CUDA-FRESCO overcomes MUMmerGPU's problems processing reads with errors or mismatches and delivers additional performance beyond MUMmerGPU's 5-12x speedup with less than 100bp query length.

Speaker(s):	Chun-Yuan Lin (Assistant Professor, Department of CSIE, Chang Gung University)
Topic(s):	Life Sciences, Algorithms & Numerical Techniques, Tools & Libraries

THURSDAY, SEPT 23, 15:00 (50 MINUTES)
ROOM L

2107 Accelerating Stereographic and Multi-View Images Using Layered Rendering

Explore applications of geometry shaders in improving the performance of stereo pair or multi-viewer image generation. This session will cover the basic approach of single-pass stereo-pair creation and provides guidelines for when layered rendering can be used to increase performance. A particular emphasis will be placed on virtual reality and scientific visualization, but the techniques discussed apply to a wide range of rendering environments. Results will be shown for three GPU architectures, including the new GF100 GPU.

Speaker(s):	Jonathan Marbach (Director of Software Architecture and Engineering, TerraSpark Geosciences, LLC)
Topic(s):	Stereoscopic 3D

THURSDAY, SEPT 23, 15:00 (50 MINUTES)
ROOM C

2123 Enabling Augmented Reality with GPU Computing

This talk will take a detailed look at Sportvision's "First and 10" system, perhaps the most widely experienced example of AR ever, with 106 million viewers during the 2010 Superbowl alone. We'll examine the current implementation and the GPU features that enable low latency, video-rate performance.

Speaker(s):	Ryan Ismert (Director of Engineering, Sportvision, Inc.)
Topic(s):	Computer Vision

THURSDAY, SEPT 23, 15:00 (50 MINUTES)
ROOM A3

2153 CULA - A Hybrid GPU Linear Algebra Package

Get the latest information on CULA, an implementation of hybrid GPU/CPU linear algebra solvers for NVIDIA GPUs. CULA launched at GTC2009 and has since received large speedups and many new features. We will cover all the features, old and new, along with performance, inner workings, and how users can integrate CULA into their applications. Learn how your existing linear algebra applications can benefit from a high quality library. Much more information is available at www.culatools.com and at our presentation and booth.

Speaker(s):	John Humphrey (Senior Engineer, EM Photonics, Inc)
Topic(s):	High Performance Computing, Algorithms & Numerical Techniques, Tools & Libraries

THURSDAY, SEPT 23, 15:00 (20 MINUTES)
ROOM A8

2213 BCSLIB-GPU: Significant Performance Gains for CAE

Hear product architects and developers describe the algorithmic depths and high level breath of the use of GPUs that have been employed to create BCSLIB-GPU, the GPU enablement of the industry standard sparse matrix software suite, BCSLIB-EXT. We provide a range of comparison data with Tesla and Fermi compared with multi-core CPU only systems and for a wide range of realisitic demanding real world test problems.

Speaker(s):	Danl Pierce (Partner, Access Analytics Int'l, LLC)
Topic(s):	Tools & Libraries, Algorithms & Numerical Techniques, High Performance Computing, Embedded & Automotive

THURSDAY, SEPT 23, 15:00 (50 MINUTES)
KEYNOTE HALL

4010 Emerging Companies: CEO on Stage featuring NaturalMotion Ltd, OptiTex, and Useful Progress

See the hottest new technologies from startups that are transforming computing.

In a lively and fast-paced exchange, the "Emerging Companies Summit - CEO on Stage" sessions will feature CEOs from three startups who will each have 15 minutes to introduce their companies and interact with a panel of leading venture capitalists, technology executives, and industry analysts.

Panelist(s):	Tim Bajarin (President, Creative Strategies), Jeff Herbst (Vice President of Business Development, NVIDIA), Bill Tai (General Partner, CRV), Paul Weiskopf (Sr. VP of Corporate Development, Adobe)
Speaker(s):	Yoram Burg (President, OptiTex.), Sylvain Ordureau (CEO, Useful Progress), Torsten Reil (CEO, NaturalMotion Ltd)
Topic(s):	General Interest, Medical Imaging & Visualization, Physics Simulation, Computer Graphics

THURSDAY, SEPT 23, 15:30 (20 MINUTES)
ROOM A7

2063 Banking on Monte Carlo... and Beyond

Last year NAG presented spectacular results for Monte Carlo techniques on GPUs using NAG's GPU library. This year we will talk about new projects in the areas of Monte Carlo and PDE techniques, delivering additional benefits to the finance industry for real-world problems, including credit modeling.

Speaker(s):	Ian Reid (Chief Commercial Officer, NAG)
Topic(s):	Finance

THURSDAY, SEPT 23, 15:30 (20 MINUTES)
ROOM A8

2208 Acceleration of SIMULIA's Abaqus Solver on NVIDIA GPUs

Learn about Acceleware's and Dassault Systemes' integrated solution that performs an LDL^T factorization on GPUs within the Abaqus software package. We will discuss efficient GPU

parallelization of the factorization algorithm and enabling the CPU and GPU to overlap their computations and data transfers. Includes an end user simulation case study and GPU performance measurements including 300 GFlops in single precision and 145 GFlops in double precision on NVIDIA Tesla C2050.

Speaker(s):	Chris Mason (Product Manager, Acceleware)
Topic(s):	High Performance Computing

THURSDAY, SEPT 23, 14:00 (50 MINUTES)
ROOM A5

2008 OpenCL Optimization

Learn how to optimize your OpenCL application to achieve maximum performance on NVIDIA GPUs. We will first briefly discuss how the OpenCL programming model maps onto NVIDIA GPU's architecture. We will then talk about memory, instruction, and NDRange optimization techniques, illustrating each with small code samples.

Speaker(s):	Peng Wang (Developer Technology Engineer, NVIDIA)
Topic(s):	Tools & Libraries, High Performance Computing

THURSDAY, SEPT 23, 16:00 (50 MINUTES)
ROOM B

2086 GPGPU DL_POLY

Discover DL_POLY.

DL_POLY: an MD code ICHEC has ported to CUDA. The presentation especially focuses on the auto-tuning of the work distribution between CPU and GPU

Speaker(s):	Gilles Civario (Head of Capability Computing and Novel Architecture Group, ICHEC)
Topic(s):	Molecular Dynamics, High Performance Computing

THURSDAY, SEPT 23, 16:00 (50 MINUTES)
ROOM A2

2088 Nucleotide String Matching Using CUDA-Accelerated Agrep

Dive deep into the intelligent utilization of various CUDA memory spaces to remarkably speedup approximate DNA/RNA nucleotide sequence matching algorithm in bioinformatics by an amazing factor of 67 compared to multi-threaded quad core CPU counterpart. Our talk provides a very good example to demonstrate how to use indexable array to save frequently updated variables directly into GPU registers, how to organize shared memory into a 2D array to avoid bank conflict, and how to shuffle the data structure to satisfy the requirement for coalesced global memory access. Our CUDA implementation employs online approach and can be applied in real time.

Speaker(s):	Hongjian Li (Graduate Student, The Chinese University of Hong Kong)
Topic(s):	Life Sciences, Algorithms & Numerical Techniques

THURSDAY, SEPT 23, 16:00 (50 MINUTES)
ROOM N

2091 The GPU in the Reactive Control of Industrial Robots

Universal Robotics is using GPUs for real-time visual sensing in the reactive control of industrial robots. For a robot to work in a complex dynamic environment to achieve a more loosely specified goal, such as moving arbitrary boxes from a pallet to a conveyor, requires reactivity. Reactive control requires intensive, concurrent, low-latency computation for motion planning, exception handling, and sensing. We describe and demonstrate how GPU-based computation enables visual servoing and box moving. We also discuss the potential of the GPU to solve more difficult sensory problems such as multi-robot cooperation, multimodal sensor binding, attention, sensitization, and habituation.

Speaker(s):	Dr. Alan Peters (CTO, Universal Robotics, Inc.)
Topic(s):	Machine Learning & Artificial Intelligence

THURSDAY, SEPT 23, 16:00 (50 MINUTES)
ROOM A1

2095 Building High Density Real-Time Video Processing Systems

Learn how GPU Direct can be used to effectively build real time, high performance, cost effective video processing products. We will focus especially on how to optimize bus throughput while keeping CPU load and latency minimal.

Speaker(s):	Ronny Dewaele (Director Technology Center, Barco)
Topic(s):	Video Processing, Imaging

THURSDAY, SEPT 23, 16:00 (50 MINUTES)
ROOM A3

2100 Hybrid GPU/Multicore Solutions for Large Linear Algebra Problems

Large linear algebra problems may be solved using recursive block decomposition in which GPUs efficiently compute the sub-blocks and multicore CPUs put the sub-blocks back together within a large shared memory space. This talk will present benchmark results for such a hybrid approach, implemented in Matlab® and using Jacket® to access the GPU compute power.

Speaker(s):	Nolan Davis (Research Scientist, SAIC)
Topic(s):	High Performance Computing, Algorithms & Numerical Techniques, Signal processing

THURSDAY, SEPT 23, 16:00 (50 MINUTES)
ROOM C

2114 Cascaded HOG on GPU

We propose a real time HOG based object detector implemented on GPU. To accelerate the detection process, the proposed method uses two serially-cascaded HOG detectors. The first low dimensional HOG detector discards detection windows obviously not showing target objects. It reduces the computational cost of the second high dimensional HOG detector. This method tested on 640x480 color image and the same size movie. The computation time decreases to 70ms per image. That is 4 times faster than a case of single detector. This method provides real time performance even on middle end GPUs such as GeForce GTS 250.

Speaker(s):	Kento Tarui (Researcher, AquaCast Corporation)
Topic(s):	Computer Vision, Machine Learning & Artificial Intelligence

THURSDAY, SEPT 23, 16:00 (50 MINUTES)
ROOM K

2126 Accelerating Signal Processing: Introduction to GPU VSIPL

Learn how to use the Vector Signal Image Processing Library to accelerate signal processing applications without needing to understand platform-specific programming and optimization techniques. We will discuss how GPU VSIPL implements the VSIPL API and uses CUDA-capable GPUs to maximize performance of several example applications.

Speaker(s):	Dan Campbell (Research Engineer, Georgia Tech Research Institute)
Topic(s):	Signal processing, Tools & Libraries

THURSDAY, SEPT 23, 16:00 (20 MINUTES)
ROOM A8

2133 3D Full Wave EM Simulations Accelerated by GPU Computing

3D Full Wave Electromagnetic simulations of RF components, antennas, printed circuit boards, can be quite time consuming. Computer Simulation Technology (CST) toolsuite includes the capability to activate GPU Computing. Examples will be shown of using Tesla C1060 and S1070 configurations to provide significant performance improvement of complex simulations.

Speaker(s):	Fabrizio Zanella (Systems Manager, CST of America)
Topic(s):	High Performance Computing,

THURSDAY, SEPT 23, 16:00 (20 MINUTES)
ROOM A7

2136 Pseudo Random Number Generators for Massively Parallel Apps

Learn how to select the best and fastest pseudo random number generator for your massively parallel Monte Carlo simulation. Pseudo random numbers generators (PRNG) are a fundamental building block of these simulations and it is thus required to select suitable PRNGs with regard to the specific problem at hand while considering the parallel hardware architecture. Recent developments in random number generations provide a wide variety of choices, each with different properties and trade-offs. We provide a comprehensive survey of the current state of the art for massively parallel PRNG and show a broad range of applications.

Speaker(s):	Holger Dammertz (PhD Student, Ulm University)
Topic(s):	Algorithms & Numerical Techniques, Finance

THURSDAY, SEPT 23, 16:00 (50 MINUTES)
ROOM A5

2271 Compose CUDA Masterpieces! Write better, Leverage More

Not all CUDA code is created equally. Learn how to step up your CUDA game. Also, learn how to build large, multi-person CUDA projects for your organization. In very clear descriptions, learn the difference between naïve GPU code, intermediate GPU code, and advanced GPU mastery. We show how careful construction of CUDA kernels can affect application performance. We also discuss how Jacket tools greatly facilitate the development of CUDA-based projects. Finally, we will debut the Jacket runtime's new C/C++ library. With this library, the technical computing functions in Jacket's MATLAB engine are made available in C/C++.

Speaker(s):	James Malcolm (VP of Engineering, AccelerEyes)
Topic(s):	Tools & Libraries

THURSDAY, SEPT 23, 16:00 (50 MINUTES)
ROOM D

2279 Working Man's Guide to 3D Video Editing

Video editing is currently at two simultaneous inflections points: use of GPUs for video processing and the beginning of wide spread adoption of 3D. At this time however, identifying and navigating through the necessary tools and equipment to create compelling 3D video content is challenging. This session is intended to provide a pragmatic guide to creating prosumer 3D video content and how the GPU greatly assists and speeds up this process. The intended audience is anyone interested in how to create compelling 3D movies at a prosumer level.

Speaker(s):	Rudy Sarzo (Principal, SMI), Ian Williams (Director PSG Applied Engineering, NVIDIA), Kevan O'Brien (NVIDIA)
Topic(s):	Digital Content Creation (DCC)

THURSDAY, SEPT 23, 16:00 (50 MINUTES)
ROOM L

2283 500 Teraflops Heterogeneous Cluster

HPC Affiliated Resource Center (ARC) will be host of a very large interactive HPC. The large cluster (CONDOR) will integrate cell broadband engine processors, GPGPUs and powerful x86 server nodes, with a combined capability of 500 Teraflops. Applications will include neuromorphic computing, video synthetic aperture radar backprojection, matrix multiplications, and others. This presentation will discuss progress on performance optimization using the Heterogeneous Cluster and lessons learned from this research.

Speaker(s):	Mark Barnell (HPC Director, Air Force Research Lab (AFRL))
Topic(s):	High Performance Computing

THURSDAY, SEPT 23, 16:00 (50 MINUTES)
KEYNOTE HALL

4011 Emerging Companies: CEO on Stage featuring Cinnafilm Inc., Perceptive Pixel, and Total Immersion

See the hottest new technologies from startups that are transforming computing.

In a lively and fast-paced exchange, the "Emerging Companies Summit - CEO on Stage" sessions will feature CEOs from three startups who will each have 15 minutes to introduce their companies and interact with a panel of leading venture capitalists, technology executives, and industry analysts.

Panelist(s):	Tim Bajarin (President, Creative Strategies), Jeff Herbst (Vice President of Business Development, NVIDIA), Bill Tai (General Partner, CRV), Paul Weiskopf (Sr. VP of Corporate Development, Adobe)
Speaker(s):	Lance Maurer (CEO, Cinnafilm, Inc.), Bruno Uzzan (Founder and CEO, Total Immersion), Jeff Han (Founder and Chief Scientist, Perceptive Pixel)
Topic(s):	General Interest, Computer Vision, Film, Imaging

THURSDAY, SEPT 23, 16:30 (20 MINUTES)
ROOM A8

2066 Accelerating System Level Signal Integrity Simulation

Discuss how GPU acceleration for key parts of the ANSYS Nexxim Simulator resulted in significant speedup over multi-core processors. We will cover time consumption and data parallelism exposure considerations, and focus on key areas where GPU acceleration was applied including convolution and Eye rendering.

Speaker(s):	Danil Kirsanov (Scientist, ANSYS), Ekanathan Palamadai (Research & Development Engineer, ANSYS)
Topic(s):	Physics Simulation, Algorithms & Numerical Techniques, Signal processing

THURSDAY, SEPT 23, 16:30 (20 MINUTES)
ROOM A7

2101 Pricing American Options Using GPUs

This presentation focuses on the challenging problem of Pricing High-Dimensional American Options (PHAO) and how GPUs can be involved in this task. On the one hand, we present a method based on Malliavin calculus which is effective for parallel architecture. On the other hand, we compare this method with Longstaff & Schwartz method which is more dedicated to sequential architecture. We will conclude with some ideas about the parallelization of the former method on a cluster of machines and finally we will discuss this method considering it as a reformulation of a non-linear parabolic problem using BSDEs.

Speaker(s):	Lokman A. Abbas-Turki (PhD Student in Applied Mathematics, Paris-Est University)
Topic(s):	Finance, Physics Simulation

THURSDAY, SEPT 23, 17:00 (60 MINUTES)
KEYNOTE HALL

1003 Closing Ceremonies and Keynote with Dr. Sebastian Thrun, Stanford University

What really causes accidents and congestion on our roadways? How close are we to fully autonomous cars? In his keynote address, Stanford Professor and Google Distinguished Engineer, Dr. Sebastian Thrun, will show how his two autonomous vehicles, Stanley (DARPA Grand Challenge winner), and Junior (2nd Place in the DARPA Urban Challenge) demonstrate how close yet how far away we are to fully autonomous cars. Using computer vision combined with lasers, radars, GPS sensors, gyros, accelerometers, and wheel velocity, the vehicle control systems are able to perceive and plan the routes to safely navigate Stanley and Junior through the courses. However, these closed courses are a far cry from everyday driving. Find out what the team will do next to get one step closer to the "holy grail" of computer vision, and a huge leap forward toward the concept of fully autonomous vehicles.

Sebastian Thrun is a professor of computer science and electrical engineering at Stanford, where he directs the Stanford AI Lab. He is also a distinguished engineer at Google. Thrun's team won the DARPA Grand Challenge, a US-Government sponsored autonomous robot race that took place in 2005. Thrun also pioneered the scientific field of probabilistic robotics, and he co-invented Google Street View. In recognition of his contributions, Thrun has been elected into the US National Academy of Engineering and the German Academy of Sciences. He is an elected fellow of the AAAI, ECCAI, and WTN. Popular Science included Thrun in their "Brilliant Ten," Forbes Magazine in their "E-Gang" members, Scientific American in their list of 50 world technology and policy leaders, and Fortune selected him as one of the 50 smartest people in tech. Wired Magazine awarded Thrun's robot Stanley the top spot in the most influential robots of all times. The robot is now part of a permanent exhibition in the Smithsonian Museum of American History. Thrun has authored 11 books and over 300 scientific articles.

Topic(s):	General Interest, Computer Vision, Machine Learning & Artificial Intelligence
Speaker(s):	Sebastian Thrun (Professor/Distinguished Engineer, Stanford/Google)

Appro GPU Computing Solutions

Supercomputer Solutions

Appro 1U Tetra GPU Solution

The most powerful 1U server that integrates **2 CPUs/4 GPUs** delivering 1,792 GPU cores per server.
Buy online today!

Appro HyperPower™ GPU Cluster

A modular hybrid blade cluster delivering **480 CPU/35,840 GPU** cores per rack cabinet offering flexibility, price/performance and easy deployment.

Don't miss out!

Attend Appro's presentation at the GPU Technology Conference this week and enter for a chance to win a prize.

Session 2270 - Appro's GPU Computing Solutions

Date: Tuesday, September 21, 3:00-3:50 pm

For more information, please visit www.appro.com
or call Appro Sales at 800-927-5464

ALGORITHMS & NUMERICAL TECHNIQUES

A01 - Communication-Avoiding QR Decomposition for GPUs

Communication-Avoiding QR is a recent algorithm for solving a QR decomposition, which is optimal with regard to the amount of communication performed. We've implemented the CAQR algorithm on the GPU and found that it performs exceptionally well, especially for the challenging case of tall-skinny matrices.

Author: Michael Anderson (University of California, Berkeley)

A02 - Accelerating Symbolic Computations on NVIDIA Fermi

We present the first implementation of a complete modular resultant algorithm on the graphics hardware. Our recent developments taking advantage of new NVidia Fermi GPU architecture and instruction set allowed us to achieve about 150x speed-up over a modular resultant algorithm from Maple 13.

Author: Pavel Emeliyanenko (Max-Planck Institute for Informatics)

A03 - Particle-In-Cell Simulations on the GPU

Particle-In-Cell simulations represent an important technique in the field of kinetic plasma simulations. 2D particle pushing and conserved current aggregation has been implemented in CUDA. On a TESLA C1060 the CUDA code is 4 times faster than SSE2 optimized code on a quad core INTEL XEON processor.

Author: Hartmut Ruhl (Ludwig-Maximilians-University)

A04 - Parallel Ant Colony Optimization with CUDA

The Ant Colony Optimization (ACO) Algorithm is a metaheuristic that is used to find shortest paths in graphs. By using CUDA to implement an ACO algorithm, we achieved significant improvement in performance over a highly-tuned sequential CPU implementation. The construction step of the ACO algorithm consists of each ant creating an independent solution, and this step is where most of the computation is spent. Since the construction step is the same for most ACO variations, parallelizing this step will also allow for easy adaptation to different pheromone updating functions. Currently, our research tests this hypothesis on the travelling salesmen problem.

Author: Octavian Nitica (University of Delaware)

A05 - High Performance and Scalable Radix Sorting for GPU Stream Architectures

The need to rank and order data is pervasive, and sorting operations are fundamental to many algorithms. This poster presents a very efficient method for sorting large sequences of fixed-length keys (and values) using GPU stream processors. Compared to the state-of-the-art, our implementation demonstrates multiple factors of speedup (up to 3.8x) for all NVIDIA GPGPUs. For this domain of sorting problems, we believe our sorting primitive to be the fastest available for any fully-programmable microarchitecture: our stock NVIDIA GTX480 sorting results exceed the 1G keys/sec average sorting rate (i.e., one billion 32-bit keys sorted per second).

Author: Duane Merrill (University of Virginia)

A06 - Task Management for Irregular Workloads on the GPU

We explore software mechanisms for managing irregular tasks on graphics processing units. Traditional GPU programming guidelines teaches us how to efficiently program the GPU for data parallel pipelines with regular input and output. We present a strategy for solving task parallel pipelines which can handle irregular workloads on the GPU. We demonstrate that dynamic scheduling and efficient memory management are critical problems in achieving high efficiency on irregular workloads. We showcase our results on a real time Reyes rendering pipeline.

Author: Stanley Tzeng (University of California, Davis)

A07 - A Hybrid Method for Solving Tridiagonal Systems on GPU

Tridiagonal linear systems are of importance to many problems in numerical analysis and computational fluid dynamics, as well as to computer graphics applications in video games and computer-animated films. This poster presents our study on the performance of multiple tridiagonal algorithms on a GPU. We design a novel hybrid algorithm that combines a work-efficient algorithm with a step-efficient algorithm in a way well-suited for a GPU architecture. Our hybrid solver achieves 8x and 2x speedup respectively in single precision and double precision over a multi-threaded highly-optimized CPU solver and a 2x speedup over a basic GPU solver.

Author: Yao Zhang (University of California, Davis)

A08 - Development of Desktop Computing Applications and Engineering Tools on GPUs

A GPU competence center and laboratory for research and collaboration within academia and partners in industry has been established in 2008 at section for Scientific Computing, DTU informatics, Technical University of Denmark. In GPULab we focus on the utilization of GPUs for high-performance computing applications and software tools in science and engineering, inverse problems, visualization, imaging, dynamic optimization. This poster illustrates the latest and most interesting projects that have been developed at our center.

Author: Hans Henrik B. Soerensen (Technical University of Denmark)

A09 - Ballot Counting for Optimal Binary Prefix Sum

This poster describes a new technique for performing binary prefix sums using Fermi's new __ballot() and __popc() functions. These instructions greatly increase intra-warp communication, allowing for an 80% speedup over standard GPU methods in applications like Radix Sort. It also points to future research that will enable suffix array construction, Burrows-Wheeler Transform, and the BZIP algorithm to take advantage of these instructions for efficient GPU compression.

Author: David Whittaker (University of Alabama at Birmingham)

A10 - Deriving Parallelism and GPU Acceleration of Algorithms with Inter-Dependent Data Fields

This poster presents an approach to derive parallelism in algorithms that involve building sparse matrix that represents relationships between inter-dependent data fields and enhancing its performance on the GPU. This work compares the algorithm performance on the GPU to its CPU variant that employs the traditional sparse matrix-vector multiplication (SpMV) approach. We have also compared our algorithm performance with CUSP SpMV on GPU. The softwares used in this work are MATLAB and Jacket – GPU engine for MATLAB

Author: Jaideep Singh (Accelereyes)

A11 - Parallelizing the Particle Level Set Method

The particle level set is widely used as an accurate interface tracking tool in simulation, computer vision and other related fields. However, high computation cost prevents applying this method to real-time and interactive scenarios. This work intensively used parallel design patterns that are implemented in the thrust library, like compaction, reduction and scattering, to parallelize the particle level set method in order to attain real-time performance.

Author: Wen Zheng (Stanford University)

A12 - Accelerating Cuda Graph Algorithms at Maximum Warp

Graphs are powerful data representations favored Graphs are powerful data representations favored in many computational domains. GPUs have showed promising results in this domain, but their performance when the graph is highly irregular. In this study, we propose three general schemes to accelerate graph algorithms on a modern GPU architecture: (i) deferred processing of outliers, (ii) efficient dynamic workload balancing and (iii) warp-based execution exploiting threads in a SIMD-like manner. Our evaluation reveals that our schemes exhibit up to 9x speedup over previous GPU algorithms and 23x over single CPU execution on irregular graphs.They also yield up to 30% improvement,even for regular graphs

Author: Sungpack Hong (Stanford University)

A13 - Implementation of Adaptive Cross Approximation on NVIDIA GPUs

The Method of Moments is a popular computational method for solving integral equations in electromagnetics. However, it suffers from high computational and memory costs since it requires the solution of a dense linear system. The Adaptive Cross Approximation (ACA) is an effective technique for compressing the system matrix thereby reducing the necessary storage as well as the number of operations required to solve the system. Acceleration of the ACA MoM with NVIDIA GPUs can finally enable the solution of "real world" scattering problems on a personal workstation in a practical timeframe.

Author: Daniel Faircloth (Georgia Tech Research Institute)

A14 - A GPU Accelerated Continuous-based Discrete Element Method for Elastodynamics Analysis

The Continuum-based Distinct Element Method (CDEM) is the combination of Finite Element Method (FEM) and Discrete Element Method (DEM), which is mainly used in general structural analyses, as well as landslide stability evaluations, coal and gas outburst analyses. By means of CUDA and a GTX-285 VGA card, the GPU version achieves hundreds times speedup ratio.

Author: Zhaosong Ma (Institute of Mechanics, Chinese Academy of Sciences)

A15 - GPU Algorithms for NURBS Minimum Distance and Clearance Computations

We present GPU algorithms and strategies for accelerating distance queries and clearance computations on models made of trimmed NURBS surfaces. We provide a generalized framework for using GPUs as co-processors in accelerating CAD operations. The accuracy of our algorithm is based on the model space precision, unlike earlier graphics algorithms that were based only on image space precision. Our algorithms are at least an order of magnitude faster and about two orders of magnitude more accurate than the commercial solid modeling kernel ACIS.

Author: Adarsh Krishnamurthy (University of California, Berkeley)

A16 - Gate-Level Simulation with GP-GPUs

This poster describes my research work on how to leverage the GP-GPU execution parallelism to achieve high performance in the time consuming problem of gate-level simulation of digital hardware designs.

Author: Debapriya Chatterjee (University of Michigan)

SGI® Altix® UV

accelerating performance

SGI Altix UV is the largest shared memory system offering the Tesla-20 series. SGI integrates NVIDIA Tesla solutions to accelerate the pace of our customers' work to solve the most computationally-intensive challenges including structural design, drug research, oil and gas exploration, and computational finance.

- Massive parallel compute power
- Optimized system performance
- Faster time-to-results

Learn more about Altix UV at www.sgi.com/altixuv

NVIDIA

sgi®

www.sgi.com

A17 - CUDA Implemenation of Barrier Option Valuation using Jump-Diffusion Model and Browning Bridge

Impressive speedups up to 100x using GPUs compared to CPUs are achieved by taking advantage data parallelism, increased bandwidth and the ability to hide latency. We have implemented a Monte Carlo valuation of a barrier option modeled by a standard diffusion process with a jump diffusion term obeying an underlying Poisson process to account for rare events. In addition, a Brownian Bridge is incorporated to account for barrier crossings in between diffusion trajectories and to reduce bias. This option is representative of exotic options which lack a closed-form solution and are amenable to Monte Carlo type methods for valuation.

Author: Vincent Natoli (Stone Ridge Technology)

ASTRONOMY & ASTROPHYSICS

B01 - Black Holes in Galactic Nuclei Simulated with Large GPU Clusters in CAS

Many, if not all galaxies harbour supermassive black holes. If galaxies merge, which is quite common in the process of hierarchical structure formation in the universe, their black holes sink to the centre of the merger remnant and form a tight binary. Depending on initial conditions and time supermassive black hole binaries are prominent gravitational wave sources, if they ultimately come close together and coalesce. We model such systems as gravitating N-body systems (stars) with two or more massive bodies (black holes), including if necessary relativistic corrections to the classical Newtonian gravitational forces (Kupi et al. 2006, Berentzen et al.2009).

Author: Rainer Spurzem (National Astronomical Obersvatories, Chinese Academy of Sciences)

AUDIO PROCESSING

C01 - Exploring Recognition Network Representations for Efficient Speech Inference on the GPU

We explore two contending recognition network representations for speech inference engines: the linear lexical model (LLM) and the weighted finite state transducer (WFST) on NVIDIA GTX285 and GTX480 GPUs. We demonstrate that while an inference engine using the simpler LLM representation evaluates 22x more transitions per second than the advanced WFST representation, the simple structure of the LLM representation allows 4.7-6.4x faster evaluation and 53-65x faster operands gathering for each state transition. We illustrate that the performance of a speech inference engine based on the LLM representation is competitive with the WFST representation on highly parallel GPUs.

Author: Jike Chong (Parasians, LLC)

C02 - Efficient Automatic Speech Recognition on the GPU

Automatic speech recognition (ASR) technology is emerging as a critical component in data analytics for a wealth of media data being generated everyday. ASR-based applications contain fine-grained concurrency that has great potential to be exploited on the GPU. However, the state-of-art ASR algorithm involves a highly parallel graph traversal on an irregular graph with millions of states and arcs, making efficient parallel implementations highly challenging. We present four generalizable techniques including: dynamic data-gather buffer, find-unique, lock-free data structures using atomics, and hybrid global/local task queues. When used together, these techniques can effectively resolve ASR implementation challenges on an NVIDIA GPU.

Author: Jike Chong (Parasians, LLC)

COMPUTATIONAL FLUID DYNAMICS

D01 - High-Order Unstructured Compressible Flow Solver on the GPU

The objective of this project is to develop a scalable and efficient high-order unstructured compressible flow solver for GPUs. The solver allows the achievement of arbitrary order of accuracy for flows over complex geometries. High-order solvers require more operations per degree of freedom, thus making them highly suitable for massively parallel processors. Preliminary results indicate speed-ups up to 70x with the Tesla C1060 compared to the Intel i7 CPU. Memory access was optimized using shared and texture memory.

Author: Patrice Castonguay (Stanford University)

D02 - Parallel 3D Geometric Multigrid Solver on GPU Clusters

An investigation of the performance and scalability of a multigrid pressure Poisson equation solver running on a GPU cluster.

Author: Dana Jacobsen (Boise State University)

D03 - Acceleration of mesh-free CFD using CUDA

In this work, the acceleration of a mesh-free Computational Fluid Dynamics (CFD) code is performed using CUDA. The poster gives an overview of the CUDA implementation strategy and the resulting performance increase.

Author: Ruairi Nestor (Irish Centre for High-End Computing)

D04 - Airblast Modelling on Multiple Tesla units

We used NVIDIA Tesla GPUs to accelerate the solution of hyperbolic partial differential equations, with application to modelling airblast generated by industrial bench mining operations. Parallelisation over multiple GPUs was achieved using MPI.

Author: Sean Lovett (University of Cambridge)

D05 - Implementation of High-Order Adaptive CFD Methods on GPGPUs

This poster describes our implementation of adaptive high-order CFD methods on GPUs. A speedup factor of up to 44 has been achieved for 2D flow problems.

Author: Z.J. Wang (Iowa State University)

D06 - Computational Fluid Dynamics on GPU

Computational Fluid Dynamics, an important branch in HPC field, has a history of seeking and requiring higher computational performance. The traditional way to satisfy this quest is to use faster machines or supercomputers. Yet these approaches seem inconvenient and costly to many individual researchers. We investigated the use of GPU to accelerate CFD codes and tested the performances on CUDA and OpenCL platform. We have ported 2D cave flow, 2D Riemann, and 2D flow over a RAE2882 airfoil to the GPU and explored some GPU-specific optimization strategies. In most cases, approximately 16 to 63 x speed up can be achieved.

Author: Long Wang (Supercomputing Center, Chinese Academy of Sciences)

COMPUTER GRAPHICS

E01 - Dynamic and Implicit Trees for Graphics and Visualization on the GPU

We propose a new way to represent trees that allows for faster algorithms, that are simple to implement (especially on the GPU), and with a lower memory overhead than previous approaches. Using our data structure, we have seen significant improvements in both volume ray casting and ray tracing applications over previous state-of-the-art methods.

Author: Nathan Andrysco (Purdue University)

E02 - Fragment-Parallel Composite and Filter

In this poster, we describe our recent work in the area of programmable graphics pipelines by presenting a fragment-parallel formulation of an A-buffer-style composite and filter equation, and describe its implementation on a modern GPU.

Author: Anjul Patney (University of California, Davis)

COMPUTER VISION

F01 - Architecture Aware Design for a Parallel Object Recognition System

We have developed a parallel object recognition system using CUDA, achieving 70x-80x speedup against the original serial implementation. In order to optimize our implementation, we evaluated the performance of different parallelization strategies on some key computations in the object recognition system. Finally we concluded that the parallel implementation performance is sensitive to input data properties. Therefore, we should dynamically adjust the parallelization strategy at runtime to optimize key computations.

Author: Bor-Yiing Su (University of California, Berkeley)

F02 - Dense Point Trajectories by GPU-Accelerated Large Displacement Optical Flow

In this poster we discuss a method for computing point trajectories based on a fast parallel implementation of a recent optical flow algorithm that tolerates fast motion. The parallel implementation of large displacement optical flow runs about 78x faster than the serial C++ version. We use this implementation is a point tracking application. Our resulting technique tracks up to three orders of magnitude more points and is 46% more accurate than the Kanade-Lucas-Tomasi tracker. Compared to the Particle Video tracker, we achieve 66% better accuracy while retaining the ability to handle large displacements while running an order of magnitude faster.

Author: Narayanan Sundaram (University of California, Berkeley)

F03 - Visual Cortex on a Chip: Large-scale, Real-Time Functional Models of Visual Cortex on a GPGPU

Los Alamos National Laboratory's Petascale Synthetic Visual Cognition project is exploring full-scale, real-time functional models of human visual cortex to understand how human vision achieves its accuracy, robustness and speed. Commercial-off-the-shelf hardware to support this modeling is rapidly improving, e.g., a teraflop GPGPU card costs ~$500 and is ~size of mouse cortex. We present results demonstrating image classification on UAV aerial video with a visual cortex model running on a 240-core NVIDIA GeForce GTX285, and see >x10 speed-up. As this technology continues to improve, cortical modeling on GPGPU devices has the potential to revolutionize computer vision.

Author: Steven Brumby (Los Alamos National Laboratory)

F04 - Fermi in Action: Robust Background Subtraction for Real-time Video Analysis

Background subtraction is one of the important image processing steps for video surveillance and many computer vision problems such as tracking & recognition. However, robust background subtraction that adapts well to variable environment changes is highly computational and consumed large amount of memory. Thus, its practical application is often limited. Here, we aimed to expand its usage and tackle vision problems that requires high frame rate camera such as real-time sports analysis, real-time object detection and recognition. Using recent advances in accelerator hardware – NVIDIA Fermi Architecture and taking advantage of heterogeneous computing , we are able to gain good performance that allows to use in these practical applications.

Author: Melvin Wong (Institute for Infocomm Research)

F05 - Bridging Neuroscience and GPU Computing to Build General Purpose Computer Vision

The construction of artificial vision systems and the study of biological vision are naturally intertwined as they represent simultaneous efforts to forward- and reverse-engineer systems with similar goals. Here, we present a high-throughput approach to more expansively explore biologically-inspired models by leveraging GPUs. We show that this approach can yield significant gains in performance on object and face recognition (including "Labeled Faces in the Wild" challenge and faces from Facebook), consistently outperforming the state-of-the-art. We highlight how the application of flexible programming tools, such as high-level scripting, template metaprogramming/auto-tuning, can enable large performance gains, while managing complexity for the developer.

Author: Nicolas Pinto (Massachusetts Institute of Technology)

F06 - CUDA for Vision and Imaging Library

CUVI Lib (CUDA for Vision and Imaging Library) is a software library that provides a set of GPU accelerated computer vision and image processing functions. CUVI can both be utilized as an add-on library for the NVIDIA's NPP (NVIDIA Performance Primitives) as it compliments the functionality present in NPP as well as it can be used as a standalone library ready to be plugged into end-user C/C++ applications.

Author: Salman Ul Haq (TunaCode)

F07 - GPU-Friendly Multi-View Stereo Reconstruction Using Surfel Representation and Graph Cuts

We present a new surfel (surface element) based multi-view stereo algorithm which runs entirely on GPU. We utilize flexibility of surfel-based 3D shape representation and global optimization by graph cuts in a same framework.The orientation of the constructed surfel candidates imposes an effective constraint that reduces the effect of the minimal surface bias. The entire processing pipeline is implemented on the latest GPU to speed up the processing significantly. Experimental results show that the proposed approach reconstructs the 3D shape of an object accurately and efficiently, which runs more than 100 times faster than on CPU.

Author: In Kyu Park (Inha University)

F08 - CUDA Accelerated Face Recognition

A GPU based implementation of a face recognition solution using PCA with Eigenfaces algorithm.

Author: Jayadeep Vijayan (NeST Software)

F09 - GPU Driven Dense Reconstruction for Community Photo Collections

We present a system to reconstruct dense 3D models from community photo collections. First images are described using GIST and are clustered using hamming distances. Each of these clusters is geometrically verified and connected using Geotags. Connected clusters are bundle adjusted and the obtained registration is used to estimate depthmaps that are finally fused to obtain dense 3D models. Each of the above steps, except Bundle Adjustment, is implemented in CUDA and runs on multiple GPUs. The performance of our pipeline is two order of magnitude faster on one order more images compared to state of the art method.

Author: Jan-Michael Frahm (University of North Carolina, Chapel Hill)

F10 - Portable Central Vision Enhancement System for Macular Degeneration Patients

Vision enhancement systems is an alternative visual aid device to enhance the remaining vision for visual impairment subjects. Our aim is to develop a mobile central vision enhancement system for macular degeneration patients. Three different types of enhancement algorithms have been developed and their efficiency was tested on low vision patients. These three algorithms have been implemented on a portable low power devic. The Nvidia system-on-a-chip Tegra has been chosen for this implementation.

Author: Chloe Vaniet (Imperial College London)

F11 - Dense Stereo Vision on GPU

A dense stereo vision for a material handling dual-arm industrial robot have been implemented with the Rectification, Stereo Correspondence and 3D Pose from depth are ported out to GPU using CUDA.

Author: Esubalew Bekele (Universal Robotics Inc.)

F12 - Upsampling Range Data in Dynamic Environments

We present a flexible, parallelized method for fusing information from optical and range sensors based on an accelerated high-dimensional filtering approach. Our system takes as input a sequence of monocular camera images as well as a stream of sparse range measurements as obtained from a laser or other sensor system. Our method produces a dense, high-resolution depth map of the scene, automatically generating confidence values for every interpolated depth point. We describe how to integrate priors on object shape, motion and appearance and how to achieve an efficient implementation using parallel processing hardware such as GPUs.

Author: Jennifer Dolson (Stanford University)

F13 - GPU Accelerated Marker-less Motion Capture

In this work, we derive an efficient filtering algorithm for tracking human pose at 4-10 frames per second using a stream of monocular depth

images. The key idea is to combine an accurate generative model-which is achievable in this setting using state of the art GPU hardware-with a discriminative model that feeds data-driven evidence about body part locations. We describe a novel algorithm for propagating the noisy evidence about body part locations up the kinematic chain using the unscented transform.We provide extensive experimental results on 28 real-world sequences using automatic ground-truth annotations from a commercial motion capture system.
Author: Varun Ganapathi (Stanford University)

F14 - 3D Facial Feature Modeling with Active Appearance Models

Active Appearance Models (AAM) is a powerful tool for modeling and matching objects under shape deformations and texture variations. It learns characteristics of objects by building a compact statistical model from applying Principal Component Analysis (PCA) to a set of labeled data. Although AAM has been widely applied in the fields of computer vision, due to its flexible framework, it still cannot satisfy the requirement of real-time situations. To alleviate this problem, we address the computational complexity of the fitting procedure by running the AAM optimization algorithm on a GPU using a hybrid CPU / GPU block processing architecture.
Author: Tim Llewellynn (nViso / EPFL)

F15 - OpenCV on GPU

OpenCV is a free open source library of computer vision algorithms. Recently a new module consisting of functions implemented on GPU was introduced in OpenCV. It consists of several methods for calculating stereo correspondence between two images that is used to reconstruct a 3D scene. A simple block-matching algorithm works up to 10x faster compared to a CPU implementation in OpenCV providing real-time processing of HD stereo pairs on Tesla cards. Belief propagation-based algorithms show 20-50x speedup compared to a CPU implementation.
Author: Anatoly Baksheev (ITEEZ)

DATABASES & DATA MINING

G02 - Speculative Query Processing

With an increasing amount of data and user demands for fast query processing, the optimization of database operations continues to be a challenging task. A common optimization method is to leverage parallel hardware architectures. With the introduction of general-purpose GPU computing, massively parallel hardware has become available within commodity hardware. To efficiently exploit this technology, we introduce the method of speculative query processing. This speculative query processing works on index structures to efficiently support heavily used database operations. To show the benefits and opportunities of our approach, we present a fine and coarse grain implementation for multidimensional queries.
Author: Peter Volk (Technische Universität Dresden)

G03 - Virtual Local Stores

We propose a mechanism to provide the benefits of a software-managed memory hierarchy on top of a hierarchy of hardware-managed caches. A virtual local store (VLS) is mapped into the virtual address space of a process and backed by physical main memory, but is stored in a partition of the hardware-managed cache when active. This reduces context switch cost, and allows VLSs to migrate with their process thread. The partition allocated to the VLS can be rapidly reconfigured without flushing the cache, allowing programmers to selectively use VLS in a library routine with low overhead.
Author: Henry Cook (University of California, Berkeley)

EMBEDDED & AUTOMOTIVE

H01 - Driver Assistance: Speed-Limit Sign Recognition on the GPU

We investigate the use of differentGPU-based implementations for performing real-time speed limit sign recognition on a resource-constrainedembedded system. The system recognized US and European Union speed-limits at over 88% accuracy while running in real-time. The system is hardware-accelerated using CUDA and OpenGL. It introduces a novel technique for detecting speed-limit signs which is only possible with the aid of GPU processing.
Author: Vladimir Glavtchev (BMW)

H02 - Complex Automotive Applications

NVIDIA GPU architecture becomes a very interesting hardware target for complex automotive application. We implemented the same automotive application on several different hardware targets and analyzed the maximum frame rate and the effective CPU charge. This paper shows how real-time applications like pedestrian detection and driving assistance take benefits from a massively parallel "central" architecture like GPU/CUDA. Real-time performance and zero-delay transfers can be achieved using a full asynchronous implementation. The same approach can really multiply the application performance by the number of GPU devices present on the embedded system, at a reasonable power consumption.
Author: Marius Vasiliu (University of Paris Sud)

HIGH PERFORMANCE COMPUTING

I01 - A GPU-based Architecture for Real-Time Data Assessment at Synchrotron Experiments

Modern X-ray imaging cameras provide millions of pixels and several thousand frames per second.

To process such an amount of information we have optimized the reconstruction software employed at the tomography beamlines of ANKA and ESRF synchrotrons to use the computational power of modern graphic cards. Using GPUs as compute coprocessors we were able to reduce the reconstruction time by a factor 30 and process a typical data set of 20GB in 40 seconds. The time needed for the first evaluation of the reconstructed sample is reduced significantly and quasi real-time visualization is now possible.
Author: Suren Chilingaryan (Karlsruhe Institute of Technology)

I02 - Automatic High-Performance GPU code Generation using CUDA-CHiLL

This poster presents a system to automatically generate high-performance GPU code starting from an input sequential loop nest computation. The compiler analyzes input computation in C and automatically generates a set of equivalent code variants represented by transformation recipe. These recipes guide the underlying code transformation and generation framework to apply code transformations and ultimately produces CUDA code.
We use the system to generate high performing CUDA code for four BLAS functions, matrix transpose and convolution stencils. The results mostly outperform CUBLAS2.2/CUDA_SDK2.2 and naive GPU kernel and can achieve perform up to 435GF(mm) with avg speedup up to 1.78x.
Author: Malik M Khan (USC/ UoU)

I03 - CSIRO Advances in GPU Computing. What could you do with 256 GPUs?

The Commonwealth Scientific and Industrial Research Organisation (CSIRO) is Australia's national science agency. CSIRO is currently applying GPU Computing on a scale ranging from single GPU workstations through to their 256 GPU cluster. This poster showcases some of CSIRO's work in the areas GPU accelerated biological imaging, image deconvolution, synchrotron science and CT reconstruction, and statistical inference in complex environmental models. Speedups of between 8 to 230x have been seen across these applications areas using a broard range of GPU computing platforms.
Author: Luke Domanski (CSIRO)

I04 - High Performance Agent-Based Simulation with FLAME for the GPU

The Flexibile Large-scale Agent Modelling Environment for the GPU (FLAME GPU) addresses the performance and architecture limitations of previous work by presenting a flexible framework approach to ABM on the GPU. Most importantly it addresses the issue of agent heterogeneity through the use of state machine based agent representation. This representation allows agents to be separated into associated state lists which are processed in batches to allow very diverse population of agents whilst avoiding large divergence in parallel code kernels. The use of the GPU allows AB models to be visualised in real time, which further widens the application of ABM to real-time simulations.
Author: Paul Richmond (University of Sheffield)

I05 - The Scalable HeterOgeneous Computing (SHOC) Benchmark Suite

SHOC is a benchmark suite for heterogeneous systems. This poster describes the suite and presents recent performance measurements.
Author: Kyle Spafford (Oak Ridge National Laboratory)

I06 - HyperFlow: An Efficient Dataflow Architecture for Multi CPU-GPU Systems

We propose a new pipeline architecture that can take advantage of the many processing elements available in modern CPU-GPU systems to maximize performance in visualization and computational tasks. Our architecture is very flexible and allows the construction of classical parallel algorithms such as data streamers and map/reduce templates. We also discuss examples and performance benchmarks that demonstrate the potential of our system.
Author: Huy Vo (University of Utah)

I07 - MPI-CUDA Applications Checkpointing

We propose a checkpoint/restart tool for multi-GPU applications such as MPI-CUDA applications
Author: Nguyen Toan (Tokyo Institute of Technology)

I08 - Particle Simulations using DEM on GPUs

Particle based numerical methods are an emerging field since the GPU/CUDA technique became widely accepted in the last years. 80% of the whole material,used in pharmaceutical technology are powders. Numerical simulations of such material is possible by using the Discrete Element Method (DEM). The main restrictions here is compute power together with the problem size. Only a few ten-thousand particles lead to weeks to months of compute time in order to reflect processes of a few minutes in real time. DEM scales excelent with the massively-parallel CUDA environment, enabling us to access the million particle range in acceptable job runtimes.
Author: Charles Radeke (University Graz)

I09 - Mastering Multi-GPU Computing on a Torus Network

We describe APEnet+, the new generation of our 3D torus network which scales up to tens of thousands of cluster nodes with linear cost. The basic component is a custom PCIe adapter with six high-speed links, designed around a programmable HW component (FPGA), a nice environment for studying integration techniques between GPUs and network interfaces. The highlevel programming model is MPI, while a low-level RDMA API is also available.
Author: Davide Rossetti (National Institute of Nuclear Physics)

I10 - Atmospheric Modelling, Simulation and Visualization using CUDA

The Laboratory Meteorological Dynamics (LMD) by CNRS weather model is used extensively for research and weather forecasting purposes. Simulation of atmospheric climate is one of the most challenging computational tasks because of its numerical complexity and simulation time. The numerical simulations must be obviously achieved faster than in real time to use them in decision support.

Author: Priyanka Sah (Indian Institute of Technology, Delhi)

I11 - Automatic Program Generation for the Fermi - DFT Transform

The goal of SPIRAL is to push the limits of automation in software and hardware development and optimization numerical kernels beyond what is possible with current tools. In this research, we address the problem of an efficient high performance computing platform of libraries automatically generated by a computer forNVIDIA GPU architectures. Spiral generates code that automatically bypasses all the architectural restrictions on GPUs, shared memory bank conflicts, global memory coalescing and pushes code to the limits (maximum number of threads, register pressure, etc.). The procedure of code generation is fast, platform dependent, easy to rewrite and problem adaptable.

Author: Christos Angelopoulos (Carnegie Mellon University)

I12 - Fast N-body Algorithms for Dynamic Problems on the GPU

we present an extension of the earlier algorithm by Gumerov & Duraiswami (J. Comput. Phys., 2008) which adapts the FMM to the GPU, where the data structures are efficiently generated on the GPU as well. Details and performance on current architectures will be presented.

Author: Qi Hu (University of Maryland)

I13 - GPU Acceleration of Cube Calculus Operations

In our current work, we present the first massively parallel, GPU accelerated implementation of the Cube Calculus operations for multivalued and binary logic, also called Cube Calculus Machine (CCM). Substantial speedups upto the order of 85x are achieved using the CUDA enabled nVIDIA Tesla GPU compared to the CPU implementation on a sequential processor.CC is a very efficient and convenient mathematical formalism for representation, processing and synthesis of binary and multivalued logic which has significant applications in logic synthesis, image processing and machine learning. Thus, massive speedups achieved using GPUs are very encouraging to build future parallel VLSI EDA systems

Author: Vamsi Parasa (Portland State University)

I14 - An Atomic Tesla

We examined the possibility of using an Atom-based host system to control a Tesla S1070. Our simple benchmarks found that Atom-based systems should be viable for codes with serial portions small enough to make Amdahl's Law irrelevant. Such systems would have a much lower power draw than 'traditional' GPU clusters.

Author: Richard Edgar (Massachusetts General Hospital)

I15 - ICHEC's GPU Research: Porting of Scientific Application on NVIDIA GPU

ICHEC is the Irish National HPC centre, with a mission to provide both high performance computing resources and expertise for the Irish research community. In addition to its core mission of research enablement, ICHEC started in May 2009 an exploratory activity in GPGPU and CUDA programming. Quantum Espresso is an increasingly popular molecular dynamic package, mainly developed by the DEMOCRITOS group in Trieste (IT). PWscf is part of the Qauntum Espresso suite which performs electronic and ionic structure calculations. Interesting part on the porting of PWscf is an high performance [ZD] gemm which execute in parallel between CPU and GPU.

Author: Ivan Girotto (Irish Centre for High-End Computing)

I16 - Implementation of Smith-Waterman algorithm in OpenCL for GPUs

In the poster is presented the implementation of Smith-Waterman algorithm done in OpenCL. This implementation is capable of computing similarity indexes between query sequences and a reference sequence with or without sequence alignment paths. In accordance with the requirement for the target application in cancer research the implementation provides processing of very long reference sequences (in the order of millions of nucleotides). Performance compares favorably against CPU, being on the order of 14 - 610 times faster; 4.5 times faster than the Farrar's implementation. It is also on par with CUDASW++v2.0.1 performance, but with less constraints in sequence length.

Author: Dzmitry Razmyslovich (Institute of Computer Engineering, University of Heidelberg)

I17 - Computing Strongly Connected Components in Parallel on CUDA

The problem of decomposition of a directed graph into its strongly connected components is a fundamental graph problem inherently present in many scientific and commercial applications. We show how existing parallel algorithms can be reformulated in order to be accelerated by NVIDIA CUDA technology. We design a new CUDA-aware procedure for pivot selection and we redesign the parallel algorithms in order to allow for CUDA accelerated computation. We experimentally demonstrate that with a single

GTX 280 GPU card we can easily outperform optimal serial CPU algorithm.
Author: Milan Ceska (Masaryk University)

I18 - A CUDA Runtime Target for the Sequoia Compiler

We describe an implementation of the Sequoia Runtime interface in CUDA that enables the Sequoia compiler to target programs written in Sequoia for single and multiple GPU systems.
Author: Michael Bauer (Stanford University)

I19 - GPU Computing for Real-Time Optical Measurement Techniques

Measuring displacement and strains during deformation of advanced materials which are too small, big, compliant, soft or hot are typical scenarios where non-contact techniques are needed. Using Digital Image Correlation and Tracking, strain can be calculated from a series of consecutive images with sub pixel resolution. However, the image processing is a computation intensive task and can't be performed in real time using general purpose processors. We implemented 3 stage pipelined architecture: images are loaded, preprocessed using CPU, and correlated on GPUs. Using two GTX295 cards we were able to reach 35 times speedup compared to fastest Core i7 processor.
Author: Suren Chilingaryan (Karlsruhe Institute of Technology)

I20 - An MPI/CUDA Implementation of Discontinuous Galerkin Time Domain Method for Maxwell's Equations

We describe an MPI/CUDA approach to solve Maxwell's equations in time domain by means of an Interior Penalty Discontinuous Galerkin Time Domain Methods and a local time stepping algorithm. We show that MPI/CUDA provides 10x speed up versus MPI/CPU, in double precision. Moreover, we present scalability results and an 85% parallelization efficiency up to 40 GPUs on the Glenn cluster of Ohio Supercomputing Center. Finally, we study an electromagnetic cloaking example for a broad band signal(8-11GHz), to show the potential of our approach to solve real life examples in short simulation times.
Author: Stylianos Dosopoulos (Ohio State University)

IMAGING

J01 - Neurite Detection using CUDA, GPU Accelerated Biological Imaging for High-Content Analysis

The analysis of microscopic neurite structures in images is an important for studying the effects of lead compounds on brain diseases or the regeneration of brain cells after trauma. In High-Content Analysis (HCA) 100s to 1000s of microscopy images are processed during automated experiments. The speed of the image processing in these situations greatly affects the workflow throughput. We report some early results on GPU acceleration of the Neurite Detection module in our groups' HCA-Vision. The most time consuming algorithm steps are accelerated by up to 13.6x resulting in a 3.3x speedup for the entire algorithm (70% of theretical maximum).
Author: Luke Domanski (CSIRO)

J02 - Fast Radon Transform via Fast Non-uniform FFTs on GPUs

Fast Radon Transform is required in X-ray Phase Contrast Tomography performed at the Advanced Light Source, Lawrence Berkeley National Lab. We describe a fast implementation based on fast non-uniform FFTs on GPUs.
Author: Chao Yang (Lawrence Berkeley National Laboratory)

J03 - Projected Conjugate Gradient Solvers on GPU and its Applications

In this work, the focus is specifically on how to speedup the projected CG algorithm utilizing the GPU. It is shown that the projected CG method can be used within the single precision accuracy of the current GPU. One benefit gained through use of the projected CG is that it reduces the total number of matrix vector multiplications, which is usually a bottleneck for an efficient GPU-based Krylov-based algorithm. A modified projection based CG algorithm in the thesis is further proposed which shows a better performance. Numerical results using the GPU are provided to support the proposed algorithm.
Author: Youzuo Lin (Arizona State University)

J04 - Real-time Direct Georeferencing of Images from Airborne Line Scan Cameras

The Norwegian Defense Research Establishment (FFI) is developing a technology demonstrator for airborne real-time hyperspectral target detection. The system includes two nadir-pointing line scan cameras. The line scanned images are georeferenced in real-time by intersecting rays cast from the cameras with a 3D model of the terrain underneath. The georeferenced images may then easily be ortho-rectified (e.g by using texture mapping in OpenGL) and overlaid digital maps. This poster presents the performance of a cuda implementation of the georeferencing method.
Author: Trym Vegard Haavardsholm (Norwegian Defence Research Establishment (FFI))

J05 - CUDA Acceleration of Color Histogram Matching

Histogram matching techniques are methods for the adjustment of color in a pair of images. It can be used as a preliminary stage for several video applications as for example 3D content creation. In such application two cameras separated a known distance acquire video streams that can be

combined in order to compute a depth map. As both cameras take slightly different scenes they can be lit by different sources becoming a possible color shift between their streams and thus penalizing the quality and the user experience. Our approach considers the use of a NVIDIA 3D broadcast solution system with professional HD cameras.
Author: Antonio Sanz (Universidad Rey Juan Carlos)

LIFE SCIENCES

K01 - Generalized Linear Model (GLM) Based Quantitative Trait Locus (QTL) Analysis

Relating Genotype to Phenotype in Complex Environments has been identified as one of the grand challenges of plant sciences. Under the umbrella of the iPlant Collaborative funded by the Plant Science Cyberinfrastructure Collaborative program of the NSF, our goal is to develop GPU implementation of the General Linear Model (GLM) to statistically link genotype to phenotype and dramatically decrease the execution time for GLM analyses. GPU based highly parallelized Forward Regression stage of the GLM achieved 177x speedup over the Matlab based serial version. Results of this study will enable larger, more intensive genetic mapping analyses to be conducted.
Author: Ali Akoglu (University of Arizona)

K02 - GPU-REMuSiC: The Implementation of Constrain Multiple Sequence Alignment on Graphics Processing Unit

We implement RE-MuSiC tool on multi-GPUs (called GPU-REMuSiC) with NVIDIA CUDA. By a special model implementation, the DP computation time in GPU-REMuSiC running on single and two GeForce GTX 260 cards achieves more than 75 and 130 speedups comparing to that in sequential RE-MuSiC running on Intel i7 920 CPU, respectively.
Author: Chun-Yuan Lin (Chang Gung University)

MACHINE LEARNING & ARTIFICIAL INTELLIGENCE

L01 - CUDA Creatures

CUDA Creatures applies parallel algorithms to the iterated Prisoner's Dilemma, a classic study of the evolution of cooperation. We bring interactivity to parameter space exploration by achieving 600x to 800x speedups on GTX 260.
Author: Andrew Hershberger (Stanford University)

MEDICAL IMAGING & VISUALIZATION

M01 - Real-time Ultrasound Data Processing for Regional Anesthesia Guidance

Ultrasound imaging techniques such as Doppler flow imaging and acoustic radiation force impulse (ARFI) imaging require estimation of velocity or displacement from the received echoes. Real-time processing and display of images allows for real-time guidance of procedures, improving patient safety and efficacy. Using CUDA, the processing code has been implemented in pre-clinical regional anesthesia studies investigating new methods for localizing where fluid is being injected. The computation time has been reduced from 20 minutes to 18 seconds, resulting in the rapid display of dynamic images of the fluid being injected.
Author: Stephen Rosenzweig (Duke University)

M02 - GPU-Accelerated Texture Decompression of Biomedical Image Stacks

Histopathology is the microscopic examination of tissue in order to study the manifestations of disease. High resolutions images are vital for accurate diagnoses and a major obstacle to the use of digital imaging in histopathology has been the inability to display these large images at interactive rates. We have created a tool for interactive visualization of biomedical image stacks using GPU-accelerated on-the-fly texture decompression. The image stacks are compressed using a novel approach custom tailored for the data we are dealing with, i.e. data exhibiting exceptionally high coherence between the slices of each image stack.
Author: Chirantan Ekbote (Harvard University)

M03 - Accelerated Large Scale Spherical Model Forward Solutions for the EEG/MEG using CUDA

The study presented in the poster looks at the utility of a CUDA based approach to improve the computational speed of the spherical model EEG and MEG forward solution for large scale 3-D dipole grid (on order of 1000 and up) and sensor locations (on order of 100 and up). Fast computation of the forward solution is critical in improving the speed of the inverse solution in biosource imaging. The inverse solution gives the location of the epileptogenic foci from the EEG and MEG measurements.
Author: Nitin Bangera (MIND Research Network)

M04 - CUDA Accelerated Real Time Volumetric Cardiac Image Enhancement

CUDA enables high data rate real time volumetric cardiac ultrasound image enhancement. Substantial improvements in processing data rate and memory bandwidth demand over a CPU based approach were found with CUDA.
Author: Ismayil Guracar (Siemens Medical Solutons)

M05 - Efficient Visualization of Salient Manifolds in Scalar, Vector, and Tensor Fields

Our research focuses on harnessing the massively parallel compute power of the GPU to visually explore complex datasets. We propose adaptive GPU-based approaches that intertwines computation and rendering. Along side we present novel dynamic data structures for the GPU. Our research include the visualization of salient structures in vector fields using LCS, extraction of ridge and valley surfaces from volumetric scalar fields with scale analysis, and efficient volume / surface rendering.

Author: Samer Barakat (Purdue University)

M06 - Highly Parallel Image Reconstruction for Positron Emission Tomography (PET)

We present a novel method of computing line projection operations required for list-mode ordered-subsets expectation-maximization (OSEM) for fully 3-D PET image reconstruction on a GPU using the CUDA framework. Our method overcomes challenges such as compute thread divergence and exploits GPU capabilities such as shared memory and atomic operations. This new GPU-CUDA implementation is 120X faster than a reference CPU implementation. The image quality is preserved with root mean squared (RMS) deviation between the images generated using the CPU and the GPU being 0.08%, which has negligible effect in typical clinical applications.

Author: Jingyu Cui (Stanford University)

MOLECULAR DYNAMICS

N01 - Energy Evaluation of Rosetta Proteins Using CUDA

In this poster, we describe preliminary results using CUDA to accelerate the energy evaluation of proteins folded by the Rosetta software suite.

Author: Will Kohut (University of California, Davis)

N02 - GPU Accelerated Molecular Dynamics Algorithms for Soft Matter Systems using HOOMD-Blue

The rheological, thermodynamic, and self-assembly behavior of liquids, colloids, polymers, foams, gels, granular materials and biological systems are often studied in simulation by using coarse-grained models based on molecular dynamics algorithms. The open source general purpose particle dynamics code HOOMD-Blue has been expanded to include the simulation techniques and pair potentials used to study this class of problems.

Author: Carolyn Phillips (University of Michigan)

NEUROSCIENCE

O01 - Distributed Multi-Level Out-of-Core Volume Rendering

In neuroscience, scans of brain tissue are acquired using electron microscopy, resulting in extremely high-resolution volume data with sizes of many terabytes. To support the work of neurobiologists, interactive exploration of such volumes requires new approaches for distributed out-of-core volume rendering. A major goal of our distributed GPU volume rendering system is to sustain a pixel-to-voxel ratio of about 1:1. This display-aware approach effectively bounds the working set size required for ray-casting, which makes it largely independent of the volume resolution. Currently, our system achieves interactive volume rendering of 43GB and 92GB volumes on 1 to 8 Tesla nodes.

Author: Johanna Beyer (King Abdullah University of Science and Technology)

PROGRAMMING LANGUAGES & TECHNIQUES

P01 - GPU-to-CPU Callbacks

Our poster outlines GPU-to-CPU callbacks, a method for the GPU to request work from the CPU. We give some motivation, demonstrate the code architecture, and give samples of CPU and GPU code that show callbacks being executed.

Author: Jeff Stuart (University of California, Davis)

PHYSICS SIMULATION

Q01 - Acceleration of Computational Electromagnetics Physical Optics - Shooting and Bouncing Ray Method

Electromagnetic fields radiated by a 1964 Ford Thunderbird are calculated over 50 times faster than a standard CPU by using a Quadro FX 5800 GPU.

Author: Huan-Ting Meng (University of Illinois at Urbana-Champaign)

Q02 - Massively Parallel Micromagnetic FEM Calculations with Graphical Processing Units

We adapted our Micromagnetic Simulator "TetraMag" to NVIDIA's CUDA architecture, resulting in a significant increase in calculation speed and cost efficiency over the most recent PC-based machines. The poster gives an outline of the general challenges and the methods used to adapt the solutions to GPUs as well as benchmark results obtained using standard micromagnetic problems.

Author: Elmar Westphal (Forschungszentrum Juelich)

Q03 - Multiplying Speedups: GPU-Accelerated Fast Multipole BEM, for Applications in Protein Electrostatics

We have developed a fast multipole boundary element method (BEM) for biomolecular electrostatics. With GPU acceleration of the FMM, there is a multiplicative speed-up resulting from the fast O(N) algorithm and GPU hardware. With this method, we can obtain converged results for multi-million atom systems in less than an hour, using multi-GPU clusters.
Author: Lorena Barba (Boston University)

Q04 - GPU-Powered Control of a Compliant Humanoid Robot

The ECCEROBOT project deals with the construction and control of a robot with a humanoid skeleton and muscle-like compliant, elastic actuators. The nonlinear passive and active coupling between the skeletal elements, combined with the effect of environmental interaction, present an extremly complex control problem. Our solution; motor programs are found using physics-based simulation of both the robot and its environment to locate candidate movements. For real time control multiple copies of the simulation must be run in faster than real time, requiring the use of GPU acceleration. Further, in order to capture the environment we use GPU-accelerated dense reconstruction vision.
Author: Alan Diamond (University Of Sussex, UK)

PROGRAMMING LANGUAGES & TECHNIQUES

R01 - A Speech Recognition Application Framework for Highly Parallel Implementations on the GPU

Data layout, data placement, and synchronization processes are not usually part of a speech application expert's daily concerns. Yet failure to carefully take these concerns into account in a highly parallel implementation on the graphics processing units (GPU) could mean an order of magnitude of loss in application performance. We present an application framework for parallel programming of automatic speech recognition (ASR) applications that allows a speech application expert to effectively implement speech applications on the GPU, and demonstrate how the ASR application framework has enabled a Matlab/Java programmer to achieve a 20x speedup in application performance on a GPU.
Author: Jike Chong (Parasians, LLC)

R02 - Scalable Computer Vision Applications

We are developing a domain specific language for computer vision algorithms that facilitates rapid implementation of algorithms that are scalable and portable across CPU-GPU architectures. The presented approach significantly lowers the barrier of implementation of computer vision algorithms for heterogeneous CPU-GPU architectures, and enables a single implementation to automatically scale to use additional hardware as it becomes available.
Author: Rami Mukhtar (NICTA)

R03 - Language and Compiler Extensions for Heterogeneous Computing

GPGPU architectures offer large performance gains over their traditional CPU counterparts for many applications. However, current GPU programming models present numerous challenges to the programmer: lower-level languages, explicit data movement, loss of portability, and performance optimization challenges. In this paper, we present novel methods and compiler transformations that increase productivity by enabling users to easily program GPUs using the high productivity programming language Chapel.
Author: Albert Sidelnik (University of Illinois at Urbana-Champaign)

SIGNAL PROCESSING

S01 - Achieving 1 TFLOP for the Radio Astronomy Correlator

In this work we apply CUDA, using the Fermi architecture, to the problem of cross-correlation arising in radio astronomy. This accounts for the bulk of computation in radio astronomy, and essentially is described by vector outer-products. Traditionally this task is performed using FPGAs, and the goal of this work was to see how efficiently GPUs could be used for this task. We describe the tiling strategies and optimization techniques employed to maximize performance. We achieve in excess of 1 teraflop per second using a single GeForce GTX 480, which corresponds to 78% of peak performance,
Author: Michael Clark (Harvard University)

S02 - CUDA Implementation of Software for Identifying Post-Translational Modifications

InsPecT is a software for identifying post-translational modifications of protein. With the help of the MS-Alignment algorithm, InsPecT can search PTMs in unrestrictive mode, even reveal unknown types of modifications. However, the MS-Alignment has a tremendous time complexity and takes more than 99% computing time of InsPecT. We accelerated MS-Alignment on GPUs. After optimization and parallelization with MPI, cuda-InsPecT, a new open source software based on MPI+CUDA with high efficiency is born.
Author: Long Wang (Supercomputing Center, Chinese Academy of Sciences)

TOOLS & LIBRARIES

U01 - Mint: An OpenMP to CUDA Translator

We aim to facilitate GPU programming for finite difference applications. We have developed Mint, a source to source compiler to generate CUDA code from OpenMP code. Mint transforms omp parallel for loops into CUDA kernels and applies domain specific optimizations such as shared memory, register and kernel fuse optimizations. Since our translator targets structured grid problems, it optimizes the code better than the general purpose compilers. In this poster, we present translation and optimization steps along with our initial performance results.

Author: Didem Unat (University of California, San Diego)

U02 - Real-Time Particle Simulation in the Blender Game Engine with OpenCL

The goal of this project is to produce interactive scientific visualizations that can be used in educational games. We use the computational power of OpenCL to enable features in the Blender Game Engine that would otherwise not be possible in real-time. By adding an interactive particle system to the game engine, we set the stage to demonstrate many interesting scientific phenomena (molecular dynamics, fluid dynamics, statistics) with the added benefit of real-time special effects for games in general.

Author: Ian Johnson (Florida State University)

U03 - GStream: A General-Purpose Data Streaming Framework on GPU Clusters

In this poster, we propose GStream, a general-purpose, scalable data streaming framework on GPUs. The contributions of GStream are as follows: (1) We provide powerful, yet concise language abstractions suitable to describe conventional algorithms as streaming problems. (2) We project these abstraction onto GPUs to fully exploit their inherent massive data- parallelism. (3) We demonstrate the viability of streaming on accelerators. Experiments show that the proposed framework provides flexibility, programmability and performance gains for various benchmarks from a variety of domains, including but not limited to data streaming, data parallel problems, numerical codes and text search.

Author: Yongpeng Zhang (North Carolina State University)

U04 - NukadaFFT : An Auto-Tuning FFT Library for CUDA GPUs

We have released our FFT library for CUDA GPUs. Most of algorithms and auto-tuning technologies of FFT for CUDA are already published. The library now supports new Fermi architecture and works with CUDA 3.0 or later.

Author: Akira Nukada (Tokyo Institute of Technology)

VIDEO PROCESSING

V01 - Real-Time Color Space Conversion for High Resolution Video

Color space conversion or color correction is a widely used technique to adapt the color characteristics of video material to the display technology employed (e.g. CRT, LCD, projection) or to create a certain artistic look. As color correction often is an interactive task and colorists need a direct response, state-of-the-art real-time color correction systems for video are so far based on expensive dedicated hardware. This submission shows the feasibility to replace dedicated color correction systems by General Purpose GPUs. It is shown that a single Tesla C2050 GPU supports real-time color correction up to a resolution of 4096x2048 pixel.

Author: Klaus Gaedke (Technicolor)

V02 - 3D Object Detection in Digital Holographic Microscope Images

Digital Holographic Microscopy (DHM) is based on the classical holographic principle invented by Hungarian physicist Dennis Gabor. The holographic images are acquired by a CCD camera. Depth slices can be reconstructed using Fourier transform. The numerical reconstruction and further image processing for object detection is done using General Purpose Graphical Processor Units (GPGPU).

Author: Vilmos Szabo (Pazmany Peter Catholic University)

GDDR5

Experience Rich Entertainment With Hynix 1Gb/2Gb Graphics Memory

A new world of graphics is coming to you

hynix
Semiconductor

High Speed and High Density

Experience the joy of rich entertainment

ECO Eco Memory

Hynix Semiconductor America Inc.
3101, North First Street, San Jose, CA 95134, U.S.A
Tel . 1-408-232-8000
http://www.hynix.com

Abbas-Turki, Lokman A.
PhD Student in Applied Mathematics (Paris-Est University)

SUPELEC Engineer in Signal Processing and MS degree in Applied Mathematics and Mathematical Finance. My research interests are: High-Performance Computing, Mathematical Finances, Economics, Statistics.

- **Session(s): 2101 - Pricing American Options Using GPUs (Thursday, Sept 23, 16:30)**

Adler, David
Principal Software Engineer (Walt Disney Animation Studios)

Mr. Adler is a Principal Software Engineer at Walt Disney Animation Studios, where he develops software such as a GPU-enhanced viewport for Maya for the Layout and Animation production departments. Mr. Adler has 17 years experience developing software for 2D and 3D animated motion picture production.

- **Session(s): 2285 - Walt Disney Animation Studios' GPU-Acelerated Animatic Lighting Process with Soft Shadows and Depth of Field (Wednesday, Sept 22, 17:00)**

Albuz, Elif
(NVIDIA)

Elif Albuz joined NVIDIA 6 years ago to design software for video acceleration on GPUs. She is currently leading CUDA FFT Library development. At NVIDIA, she designed video encoder/decoder and post-processing algorithms and architected error resiliency and various parts of video encoder. Before NVIDIA, she worked at Sony Electronics leading DVD firmware team. She holds a masters from University of Delaware with focus on Multimedia Processing on Parallel Architectures. Her expertise is in video codecs, video processing, parallel algorithm design and low-level optimizations.

- **Session(s): 2216 - CUDA Libraries Open House (Wednesday, Sept 22, 11:00)**

Anderson, Joshua
Research Area Specialist (University of Michigan)

Joshua Anderson is a Research Area Specialist in the Laboratory for Computational Nanoscience & Soft Matter Simulation at the University of Michigan. Dr. Anderson holds a Ph.D. degree in Condensed Matter Physics from Iowa State University and is the the lead developer of HOOMD-blue, a high performance particle simulation tool. His current research interests include GPU computing, polymer physics, and nanoparticle self-assembly.

- **Session(s): 2062 - HOOMD-blue: Fast and Flexible Many-Particle Dynamics (Thursday, Sept 23, 15:00)**

Anselm, Christoffer
Software Developer (Jedox Business Intelligence)

Christoffer Anselm is a software developer trainee at Jedox Business Intelligence, Germany.

- **Session(s): 2237 - Accelerating Business Intelligence Applications with Fast Multidimensional Aggregation (Wednesday, Sept 22, 15:00)**

Aoki, Takayuki
Professor (Tokyo Institute of Technology)

Professor Aoki is a deputy director of Global Scientific Information and Computing center (GSIC), Tokyo Institute of Technology. His research area is Computational Fluid Dynamics (CFD) and he has developing high-accurate numerical schemes. He is a leader of GPU computing for large-scale CFD and organizes GPU Computing Study Group with 500 members and GPU sessions in many conferences. Last year, he published a CUDA programming textbook in Japanese.

- **Session(s): 2295 - Large-scale CFD Applications and a Full GPU Implementation of a Weather Prediction Code on the TSUBAME Supercomputer (Tuesday, Sept 21, 16:00)**

Araya, Mauricio
Senior Researcher (Barcelona Supercomputing Center)

Mauricio Araya Polo received the engineering degree in computer science in 2001
from the University of Chile. He received the master (2003) and PhD (2006) degrees from the University of Nice - Sophia-Antipolis at the Institut National de Recherche en Informatique et Automatique (INRIA), France. Since 2007, he is researcher, from 2009 senior researcher, on Computational Geophysics at the Barcelona Supercomputing Center (BSC). His research interests cover the areas of multi-core architectures, and programming models, numerical algorithms and code optimization techniques for HPC.

- **Session(s): 2226 - Reverse Time Migration with GMAC (Wednesday, Sept 22, 16:00)**

Archer, Bob
Senior Computer Scientist (Adobe Systems Inc)

Bob has been at Adobe for 6 years and is the lead for the Pixel Bender language and compiler. Prior to Adobe he worked on 3D simulation, virtual reality systems and computer games.

- **Session(s): 2053 - Pixel Bender: Building a Domain Specific Language on the GPU (Thursday, Sept 23, 10:00)**

Aubert, Dominique
Lecturer (Strasbourg University)

D. Aubert is a lecturer at the University of Strasbourg, France, and is a member of the Astronomical Observatory. His research covers topics such as cosmology, the early Universe and the formation of galaxies. To this aim, he uses and developps applications for the numerical simulation of astrophysical phenomena such as the gravitational N-Body problem, the transfer of light through the large scale structures of the Universe or the dynamics of astrophysical fluids. Lately he ported such applications on the large scale multi-GPUs cluster provided by the Titane supercomputer in collaboration with the CEA, France. In 2010, these numerical investigations on GPUs owed him to be chosen as one of the "Young Astrophysicist of the Year" by the French Astrophysical Society.

- **Session(s): 2099 - Cosmology Powered by GPUs Redux (Wednesday, Sept 22, 11:00)**

Augonnet, Cedric
PhD Candidate (INRIA)

Cédric Augonnet got his bachelor degree from the École Normale Supérieure in Lyon and his master degree from the Vrije Universiteit of Amsterdam and Bordeaux University. His research activities mainly focus on high performance computing on heterogeneous multicore architectures. He is currently a PhD candidate in the Runtime team in INRIA Bordeaux.

- **Session(s): 2160 - StarPU: a Runtime System for Scheduling Tasks (Wednesday, Sept 22, 10:00)**

Ayala, Hugo
Sr. Research Associate (Blue Sky Studios)

Hugo Ayala is a Senior Research Associate at Blue Sky Studios, where he has worked on the films "Robots", "Ice Age: The Meltdown", "Horton Hears a Who", "Ice Age: Dawn of the Dinosaurs", and the upcoming film,

"Rio". His contributions range from the occasional fluid simulation and effect, to the design and implementation of footprint and crowd workflows. He works closely with animators and especial effects artists to develop production related tools. Hugo studied Mechanical Engineering at the Massachusetts Institute of Technology, from where he holds a bachelor, master, and doctorate degrees. He studied Screen Writing and Acting at the Harvard Extension school, and studied Studio Art at the Museum of Fine Arts in Boston. In addition to Blue Sky Studios, Hugo has worked as a software developer at Apple, and Boris FX.

- **Session(s): 2072 - GPUs at the Computer Animation Studio (Wednesday, Sept 22, 16:00)**

Ayres, Daniel
PhD Candidate (University of Maryland)

Daniel Ayres is a Ph.D. student researching computational algorithms related to molecular evolution at the Center for Bioinformatics and Computational Biology at the University of Maryland. A major focus of his research is on GPU-computing of likelihood calculations in phylogenetic analysis. He has a Bachelor of Science in Computer Engineering from the University of Illinois at Urbana- Champaign where he studied topics such as scientific computing, algorithms, networks, and computer graphics.

- **Session(s): 2203 - Modeling Evolution Computing the Tree of Life (Thursday, Sept 23, 11:00)**

Badesha, Amolak
Senior Application Expert & Strategist (Agilent Technologies)

Amolak Singh Badesha is Senior Application Expert and Market Strategist with Agilent Technologies. He has BSEE and MSEE degrees with focus on High-frequency Wireless and Digital Design. Amolak plays a critical role in defining the roadmap and delivering new technologies for Agilent's High-Speed Digital Silicon and System customers. He is also GPU evangelist at Agilent, and stronlgy believes in the future of highly parallel programming and hardware. Amolak has published many papers on Wireless and High-Speed Digital design. He has received severa; awards from Agilent and external customers for outstanding performance and strategic vision.

- **Session(s): 2080 - Tackling Multi-Gigabit Design Challenges with a Practical Virtual EMI/ESD Lab (Wednesday, Sept 22, 15:00)**

Bailey, Dan
R&D (Double Negative)

Dan graduated from the University of Bristol in the UK with a First Class Master of Engineering degree in Computer Science in the Summer of 2007. He worked in Research and Development at The Moving Picture Company in London for two years, concentrating predominantly on improving and extending the existing pipeline. He has been working in Research and Development on the proprietary fluid solver at Double Negative in London for the last year, focusing heavily on GPU development.

- **Session(s): 2239 - Fast GPU Preconditioning for Fluid Simulations in Film Production (Tuesday, Sept 21, 17:00)**

Bajarin, Tim
President (Creative Strategies)

Tim Bajarin is recognized as one of the leading industry consultants, analysts and futurists, covering the field of personal computers and consumer technology. Mr. Bajarin has been with Creative Strategies since 1981 and has served as a consultant to most of the leading hardware and software vendors in the industry including IBM, Apple, Xerox, Hewlett Packard/Compaq, Dell, AT&T, Microsoft, Polaroid, Lotus, Epson, Toshiba and numerous others. His articles and/or analysis have appeared in USA Today, Wall Street Journal, The New York Times, Time and Newsweek magazines, BusinessWeek and most of the leading business and trade publications.

- **Session(s): 4010 - Emerging Companies: CEO on Stage featuring NaturalMotion Ltd, OptiTex, and Useful Progress (Thursday, Sept 23, 15:00)**
- **4011 - Emerging Companies: CEO on Stage featuring Cinnafilm Inc., Perceptive Pixel, and Total Immersion (Thursday, Sept 23, 16:00)**

Barba, Lorena
Assistant Professor (Boston University)

Dr. Barba is a computational scientist and a fluid dynamicist. Her research covers particle methods used for fluid simulation, the development of fast and efficient algorithms, the use of novel computer architectures, as well as fundamental and applied fluid dynamics. She obtained her PhD (2004) in Aeronautics from CALTECH, then joined the Dept. of Mathematics at the University of Bristol, UK. Since Fall of 2008, she has been an Asst. Professor of Mechanical Engineering at Boston University.

- **Session(s): 2166 - The Triad of Extreme Computing- Fast Algorithms, Open Software and Heterogeneous Systems (Wednesday, Sept 22, 10:00)**

Barnell, Mark
HPC Director (Air Force Research Lab)

Mr. Mark Barnell has more than 24 years of experience in HPC and is the Air Force HPC Director for Advanced Computing Architectures. He works at the Information Directorate of the Air Force Research Laboratory in Rome, NY. His area includes high performance computers, Urban ISR (SAR), distributed and next generation architectures. Current projects include design, build and integration of the 500 TFLOPS heterogeneous Condor Cluster.

- **Session(s): 2283 - 500 Teraflops Heterogeneous Cluster (Thursday, Sept 23, 16:00)**

Barrett, Simon
Senior Software Engineer (NVIDIA)

Simon Barrett is a Senior Software Engineer in the Developer Tools group at NVIDIA and is part of the Parallel Nsight development team. Previously, he worked on NVIDIA's PerfHUD, the NVIDIA Display Driver, Sony PlayStation 3, and at Transmeta on Code Morphing Software.

- **Session(s): 2212 - Parallel Nsight for Accelerated DirectX 11 Development [Advanced] (Tuesday, Sept 21, 17:00)**

Bednarz, Tomasz
3d Visualisation Software Engineer (CSIRO)

He received his B.Sc. in 2001 and M.Sc. in 2002 in Technical Physics from Faculty of Physics and Applied Computer Science, AGH University of Science and Technology, Poland. In 2005 he obtained his Ph.D. in Energy and Environmental Engineering in 2005 from Interdisciplinary Graduate School of Engineering Sciences, Kyushu University, Japan. From 2005, he worked as Research Associate and Senior Research Fellow at the School of Engineering & Physical Sciences, James Cook University in Townsville, Australia. Currently, he works as 3D Visualization Software Engineer for CSIRO Earth Science and Resource Engineering in Brisbane, Australia.

Session(s): 2058 - A Practical Introduction to Computational Fluid Dynamics on GPUs (Wednesday, Sept 22, 10:00)

Bell, Nathan
Research Scientist (NVIDIA Research)

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 Ph.D in Computer Science from the University of Illinois at Urbana-Champaign (UIUC).

Session(s): 2219 - High-Productivity CUDA Development with the Thrust Template Library (Thursday, Sept 23, 11:00)

Belloch, Jose Antonio
PhD Candidate (Institute of Telecommunications and Multimedia Applications, Universidad Politecnica de Valencia)

Jose A. Belloch was born in Requena, Spain (1983). He received the degree in Electrical Engineering from the Universidad Politécnica de Valencia, Spain, (2007). In 2008, he worked for the Company Getemed (Teltow, Germany) as a software developer. In 2009, he enrolled in a PhD program in 2009 at the Institute of Telecommunications and Multimedia Applications inside Universidad Politécnica de Valencia. His research interests are focused on Audio Signal Processing onto the CUDA environment.

Session(s): 2116 - Real-time Multichannel Audio Convolution (Thursday, Sept 23, 10:00)

Bernaschi, Massimo
Professor (Istituto Applicazioni del Calcolo - C.N.R.)

Massimo Bernaschi has spent ten years with IBM mostly working in the field of parallel and distributed computing. Currently he is:Chief Technology Officer at the Istituto per le Applicazioni del Calcolo "M. Picone" of C.N.R. (National Research Council) and Adjunct Professor of Computer Science, "La Sapienza" University, Rome. He is the author of more than 120 papers in international journals and proceedings of international conferences.

Session(s): 2112 - The Heisenberg Spin Glass Model on GPU: Myth versus Fact (Tuesday, Sept 21, 11:00)

Bigdeli, Abbas
Senior Researcher and Technology Manager, (NICTA)

Abbas is currently a Senior Researcher and Technology Manager for Advanced Surveillance Project at National ICT Australia lab. He has collaborated on industrial projects with various companies in New Zealand, Australia and USA. He has more than 10 years experience in consultancy, scientific research and technology leadership in the areas of digital signal and image processing, computer architecture, and information security. He has published over 60 papers in journals, book chapters and refereed international conferences. He has invented 4 patents in the areas of information security and computer vision.

Session(s): 2173 - Enabling Large-Scale CCTV Face Recognition (Thursday, Sept 23, 11:00)

Blackman, Sam
CEO and Co-Founder (Elemental Technologies, Inc.)

Sam brings extensive management experience and video processing expertise to the Elemental team as chief executive officer. In a few short years, Sam has grown Elemental into the leading provider of GPU-accelerated video processing solutions building innovative products used by leading video content providers. Elemental™ Live, the first ever massively parallel live video processing system, allows for simultaneous encoding of video streams, targeting the comprehensive specifications required for the four-screen experience of TV, PC, tablet and mobile. Prior to co-founding Elemental in 2006, Sam specified and architected next-generation products as an IC design manager for 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. Due to a growing reputation for deep industry knowledge, Sam has been tapped for presentations and contributions to leading trade and news organizations.

Session(s): 4001 - Emerging Companies: CEO on Stage featuring Elemental Technologies, Inc., Geomerics, and Milabra (Wednesday, Sept 22, 11:00)

Bleiweiss, Avi
Principal Architect (NVIDIA Corporation)

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 25 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.

Session(s): 2207 - Playing Zero-Sum Games on the GPU (Wednesday, Sept 22, 11:00)

Blewitt, Chris
CFO (miGenius)

Combining deep knowledge and experience in the development and implementation of 'cutting edge' 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 for the diverse range of users needing these capabilities. Chris has over 25 years of experience in developing effective and efficient solutions for the wider needs of the non-technical user.

Session(s): 4002 - Emerging Companies: CEO on Stage featuring Allegorithmic SAS, Bunkspeed, and miGenius (Wednesday, Sept 22, 14:00)

Bodin, Francois
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. Professor Francois Bodin holds a Master's in CS and a PhD in CS, both from University of Rennes I.

- **Session(s): 2117 - Migration of C and Fortran Apps to GPGPU using HMPP (Wednesday, Sept 22, 11:30)**

Bradley, Thomas
Developer Technology Engineer (NVIDIA)

Thomas Bradley MEng(Hons) MIET graduated with a first-class MEng degree in Computer Systems Engineering from the University of Bristol, UK, in 2000. He also completed the final year of the Diplôme d'Ingénieur at l'École Nationale Supérieure de Télécommunications in Brest, France. He has led architecture development for video encoding processors for general purpose parallel processors at STMicroelectronics and ClearSpeed. Since then, he has specialized in High Performance Computing software development at ClearSpeed and at NVIDIA.

- **Session(s): 2064 - Correlated Paths for Monte Carlo Simulations (Thursday, Sept 23, 15:00)**

Brandvik, Tobias
PhD Student (University of Cambridge)

Tobias Brandvik is a PhD student at the Whittle Laboratory under the supervision of Dr Graham Pullan. He obtained his MEng degree from the University of Cambridge in 2007. Tobias's current PhD research is focused on how to best use emerging multi-core architectures for scientific computing. This includes both creating tools to ease the porting of legacy applications, and to investigate the possibilities offered by the greater computational power of multi-core processors in real design settings. His main area of research has been the development of the Turbostream solver and the software framework that enables the solver to run on many different multi-core processors. Now that the majority of this work has been completed, Tobias's focus is on novel uses of the solver for turbomachinery simulations.

- **Session(s): 2118 - Large-scale Gas Turbine Simulations on GPU Clusters (Wednesday, Sept 22, 16:00)**

Brockway, Tad
Product Unit Manager (Microsoft)

Tad Brockway manages the RemoteFX engineering team for Microsoft. RemoteFX is a set of RDP technologies - most prominently graphics virtualization and the use of advanced codes - that are being added to Windows Server 2008 R2 Service Pack 1. These technologies are based on the IP that Microsoft acquired and continued to develop since acquiring Calista Technologies. Tad has been an engineer and manager in the desktop virtualization space for Microsoft since 1998.

- **Session(s): 2243 - Microsoft RemoteFX - GPU Virtualization for Desktop Centralization (Wednesday, Sept 22, 17:00)**

Brodtkorb, André Rigland
Research Scientist (SINTEF ICT)

André R. Brodtkorb recently submitted his Ph.D. thesis entitled Scientific Computing on Heterogeneous Architectures, where the efficient use of GPUs has been central. He is now a research scientist at SINTEF, scandinavias largest independent research organization, and works with different aspects of GPU computing.

- **Session(s): 2102 - Evacuate Now? Faster-than-real-time Shallow Water Simulation on GPUs (Tuesday, Sept 21, 17:00)**

Brookwood, Nathan
Research Fellow (Insight 64)

Nathan Brookwood has participated in the information technology industry since the days of the first transistorized computers, and has worked with mainframes, minicomputers, personal computers and semiconductors. Fifteen years ago he moved to Gartner/Dataquest and started tracking the semiconductor market. In 1998 he established Insight 64, where he serves as the research fellow, concentrating on CPUs and GPUs used in general purpose computing applications. Nathan is widely quoted in the general and trade press on topics relating to semiconductor industry. He has degrees from MIT and the Harvard Business School, and resides in Saratoga CA, where he strives to keep his home network with roughly a dozen nodes running smoothly.

- **Session(s): 4004 - Emerging Companies: CEO on Stage featuring Cooliris, empulse GmbH, and Playcast (Wednesday, Sept 22, 16:00)**
- **4005 - Emerging Companies: CEO on Stage featuring Jedox Business Intelligence, Rocketick, and Softkinetic (Wednesday, Sept 22, 17:00**

Brown, Christopher
Partner (Open Data)

Christopher Brown is a Partner at Open Data and has been with the company since January 2005. He has over ten years of experience leading development of machine algorithms in the large data space. He received a B.S. from the University of California at Berkeley, a M.S. from the University of California at Santa Barbara and has done post graduate study at U.C. Berkeley's Haas School of Business and Stanford University.

- **Session(s): 2179 - GPU - An R Library for Native GPU Objects (Tuesday, Sept 21, 16:00)**

Brown, Shawn
Graduate Student (UNC, Chapel Hill)

Shawn Brown received his B.S. in Comp. Sci & Math at BYU in 1992. He worked as a Software Developer at Microsoft from 1994 to 1998. From 1998-2001, he was a software developer for expedia.com, until he was promoted to the Lead Developer in 2001. He then worked as the Developer Manager from 2002 to 2004. Currently, he is a grad student at UNC.

- **Session(s): 2140 - Superfast Nearest Neighbor Searches Using a Minimal kd-tree (Wednesday, Sept 22, 14:00)**

Buck, Ian
Software Director of GPU Computing (NVIDIA)

Ian Buck joined NVIDIA and started the CUDA team six years ago alongside two others, and has had the wonderful pleasure of watching CUDA grow and totally change the world of high performance computing. Before joining NVIDIA, Ian was the development lead on Brook which was the forerunner to generalized computing on GPUs.)

- **Session(s): 2275 - The Evolution of GPUs for General Purpose Computing (Wednesday, Sept 22, 11:00)**

Buckingham, Peter
Tesla Software Manager (NVIDIA)

Peter has had a long history of working in System Software development. He has worked on everything from low-level embedded RTOSes to large scale supercompters and distributed storage systems. Currently Peter leads Tesla Software at Nvidia which is responsible for supporting Data Center and Cluster deployments of GPUs and improving the software ecosystem for large scale Tesla deployments.

- **Session(s): 2225 - Tools for Managing Clusters of NVIDIA GPUs (Tuesday, Sept 21, 17:00)**

Budavari, Tamas
Research Scientist (Johns Hopkins University)

Tamas Budavari is a Research Scientist in the Department of Physics and Astronomy at The Johns Hopkins University, where he focuses on statistical and computational challenges in astroinformatics.

- **Session(s): 2092 - Integrating CUDA into a Large-Scale Commercial Database Management System (Wednesday, Sept 22, 11:00)**

Buisson, Emmanuel
CEO (Numtech)

- **Session(s): 2037 - Numtech & GPGPU, a SME Point of View (Thursday, Sept 23, 09:30)**

Burg, Yoram
President (OptiTex)

With over 20 years of entrepreneurial experience that spanned across Asia, Middle East and North America, Mr. Burg brings with him experience in multiple aspects of running multi-million dollar organizations. Applying experience in Information Technology, Purchasing, Business Planning, Project and Change Management to leading international teams on marketing implementations and support, holding various senior positions in technology and non-technology businesses. Mr. Burg is a Graduate of the MEI (MBA) program by SUT.

- **Session(s): 4010 - Emerging Companies: CEO on Stage featuring NaturalMotion Ltd, OptiTex, and Useful Progress (Thursday, Sept 23, 15:00)**

Bussmann, Michael
Junior Group Leader Computational Radiation Physics (Forschungszentrum Dresden-Rossendorf)

Michael Bussmann studied physics at the Ludwig-Maximilians University and Munich looking for the Higgs boson without finding it, then stayed in Munich for his PHD on laser cooling of relativistic ion beams, this time with success. In 2008 he joined the Laser Particle Radiation Group at the FZD in Dresden. Here he started a project to develop a particle-in-cell algorithm for GPUs. Since August 2010 he is the head of the Computational Radiation Physics Group at the FZD.

- **Session(s): 2090 - Developing Highly Scalable Particle-Mesh Codes for GPUs: A Generic Approach (Tuesday, Sept 21, 15:00)**

Cabezas, Javier
Researcher (Barcelona Supercomputing Center)

Javier Cabezas received a bachelor's degree in Computer Science and a master's degree in Computer Architecture from Universitat Politècnica de Catalunya (UPC). Since 2006, he has been a PhD student in the Computer Architecture Department at UPC. His research is focused on operating system support for heterogeneous massively-parallel computing systems and massively-parallelaccelerators.

- **Session(s): 2226 - Reverse Time Migration with GMAC (Wednesday, Sept 22, 16:00)**

Calleja, Paul
Director of High Performance Computing (University of Cambridge)

Dr. Paul Calleja is Director of High Performance Computing at Cambridge University, where he provides research computing services across all academic disciplines. Dr. Calleja obtained his Ph.D. in computational bio-physics at Bath University. After filling a post-doctoral research position at Birkbeck College, he moved into private industry, where he spearheaded early commercialization of HPC cluster solutions within the U.K. Following six years in the commercial sector, where he led the market transition from proprietary SMP systems to commodity cluster-based solutions. Dr. Calleja returned to academia to lead the formation of a new HPC service at Imperial College, London. From there, he moved to Cambridge University, to form a new HPC service and to direct a major reorganization that has resulted in University-wide HPC capabilities with a novel pay per use cloud computing model. Cambridge University now boasts one of the largest academic supercomputer in the U.K., occupying 20th position among the top 500 when first installed. Dr. Calleja sits on numerous national and international HPC committees and advisory boards, as well being a founding member of the U.K. HPC Special Interest Group.

- **Session(s): 2265 - CUDA Centers of Excellence Super-Session IV (Tuesday, Sept 21, 17:00)**

Campbell, Dan
Research Engineer (Georgia Tech Research Institute)

Dan Campbell is a Princpal Research Engineer at GTRI, and leads its High Performance Computing Branch. His work focuses on improving the ease of use, and increasing the deployment of emerging high performance computing platforms, including GPUs. He currently leads the Applications, Benchmarks, and Metrics team on DARPA's Ubiquitous High Performance Computing program, and is co-chair of the VSIPL (pr: vee-sip-uhl) Forum. His recent work includes the 2007-2009 DARPA exascale requirements studies, automated compiler evaluation, and the GPU VSIPL software library.

- **Session(s): 2126 - Accelerating Signal Processing: Introduction to GPU VSIPL (Thursday, Sept 23, 16:00)**

Capuzzo-Dolcetta, Roberto
Professor (Sapienza Univ. of Roma)

Roberto Capuzzo Dolcetta took his MS in Mathematics and his Phd in Physics in Rome. He has been an associate professor at the University of Roma La Sapienza in Astronomy and Astrophysics since 2000. He is an expert of theoretical and computational astrophysics and is a coordinator of the PhD program in Astronomy. He is also a member of the board of physicisists in the Consiglio Universitario Nazionale.

- **Session(s): 2000 - Gravitational N-body Simulations: How Massive Black Holes Interact with Stellar Systems (Wednesday, Sept 22, 14:00)**

Carmel, Charles
Vice President, Corporate Business Development (Cisco)

Charles Carmel is the Vice President of Corporate Business Development at Cisco. In this role, he leads a team responsible for setting and executing Cisco's worldwide acquisition and venture capital investment strategy. Charles has been a driving force in defining Cisco's acclaimed business development activities which focus on accelerating the company's growth by entering new markets and integrating innovative, new technologies into Cisco's core businesses. Charles has led the development and execution of more than 20 acquisitions totaling over $15 billion during his time at Cisco. These include some of Cisco's largest and most successful transactions such as Linksys and Pure Digital, which propelled Cisco into a leading role for the Consumer technology market; Scientific Atlanta, which positioned Cisco as the leader in the delivery of digital video; and Webex and Tandberg, which established Cisco's leadership in the Collaboration market. Transactions led by Charles have resulted in more than $10 billion in incremental revenue to Cisco since 2002. In addition to his work on acquisitions, Charles has helped accelerate Cisco's role as a leading Corporate Venture Capitalist with responsibility for managing

Cisco's broad based $1.5 billion venture portfolio. He has led investments across many segments and stages which have created strategic and financial returns for Cisco and has been an active participant on the board of directors for a number of portfolio companies. Prior to joining Cisco in 2001, Charles was an investment banker at Goldman Sachs where he was an early member of the Technology Investment Banking group. During his time at Goldman, Charles was involved with over $2.5 billion of financings and M&A transactions across a diverse set of industry leading technology companies. Charles holds an MBA from Stanford's Graduate School of Business and a Bachelors degree from Tufts University..

- **Session(s): 4004 - Emerging Companies: CEO on Stage featuring Cooliris, empulse GmbH, and Playcast (Wednesday, Sept 22, 16:00)**
- **4005 - Emerging Companies: CEO on Stage featuring Jedox Business Intelligence, Rocketick, and Softkinetic (Wednesday, Sept 22, 17:00**

Castonguay, Patrice
PhD Candidate (Stanford University)

Mr. Castonguay is a Ph.D candidate in the Aeronautics and Astronautics department at Stanford University working in the Aerospace Computing Laboratory under the supervision of Professor Antony Jameson. The Aerospace Computing Lab focuses on developing more efficient and robust algorithms for modeling fluid dynamics. Mr. Castonguay is interested in developing efficient high-order methods for the Navier-Stokes equations. Most specifically, his research interests include developing stable and efficient high-order methods for mixed grids.

- **Session(s): 2079 - A Fast, Scalable High-Order Unstructured Compressible Flow Solver (Tuesday, Sept 21, 11:00)**

Catanzaro, Bryan
PhD Candidate (University of California, Berkeley)

Bryan Catanzaro received his BS and MS degrees from Brigham Young University, and is currently a PhD candidate at the University of California, Berkeley. His interests center on programming models for manycore computers, with an applications driven emphasis.

- **Session(s): 2050 - Copperhead: Data-Parallel Python for the GPU (Wednesday, Sept 22, 15:00)**

Cebenoyan, Cem
Senior Manager, Developer Technology (NVIDIA)

Cem manages the Developer Technology teams in North America, China, and Japan. When not in meetings, he spends his time coming up with graphics optimizations and techniques. Before joining NVIDIA, he was a student/research assistant in the Graphics, Visualization, and Usability Lab at Georgia Tech.

- **Session(s): 2157 - DirectX 11 Overview (Pre-Conference Tutorial) (Monday, Sept 20, 13:00)**

Cervantes-Pimentel, Ulises
Senior Kernel Developer (Wolfram Research)

Ulises Cervantes-Pimentel is the Visualization Senior Kernel Developer at Wolfram Research, Inc. He received his Master and PhD at the UIUC in Applied Mathematics and has worked at Wolfram Research improving the plotting and scientific visualization capabilities of Mathematica.

- **Session(s): 2028 - Mathematica for GPU Programming (Tuesday, Sept 21, 14:00)**

Chafi, Hassan
PhD Candidate (Stanford University)

Hassan Chafi is a fourth year PhD candidate in the Electrical Engineering dept at Stanford University. His research interests include computer architectures and parallel programming models.

- **Session(s): 2177 - Simplifying Parallel Programming with Domain Specific Languages (Wednesday, Sept 22, 11:00)**

Chatterjee, Debapriya
Graduate student (University of Michigan)

Debapriya Chatterjee is a Ph.D. candidate working with Prof. Valeria Bertacco in the Electrical Engineering and Computer Science Department at University of Michigan, Ann Arbor. His research focuses on validation and verification solutions for industry-scale digital designs. His solutions entail both novel and adaptable semi-formal verification methods and the use of massively parallel multi-core platforms to boost the performance of key validation applications. Debapriya holds a B. Tech. degree from the Indian Institute of Technology, Kharagpur, India and a MS degree from the University of Michigan; *both degrees are in Computer Science and Engineering.*

- **Session(s): 2306 - Gate-Level Simulation with GP-GPUs (Wednesday, Sept 22, 10:00)**

Chen, Doris
Student (University of Toronto)

Doris received her M.A.Sc and B.A.Sc degrees in Computer Engineering from the University of Waterloo in 2007 and 2005 respectively. She currently works at Altera's Toronto Technology Center on advanced algorithms in the fields of device modeling, CAD optimizations and logic synthesis. She has authored or co-authored 6 papers in FPGA CAD and device reliability.

- **Session(s): 2068 - Parallelizing FPGA Technology Mapping using GPUs (Wednesday, Sept 22, 14:00)**

Chen, Hanning
Research Associate (Northwestern University)

- **Session(s): 2128 - Hybrid Quantum Mechanics/Electrodynamics (QM/ED) Modeling of Solar Cells on a CUDA Cluster (Wednesday, Sept 22, 17:00)**

Cheung, Mark
Physicist (Lockheed Martin Solar & Astrophysics Laboratory)

A senior physicist at Lockheed Martin Solar & A senior physicist at Lockheed Martin Solar & Astrophysics Laboratory, Mark Cheung's research focuses on understanding astrophysical processes occurring in the Sun. He specializes in using radiation magnetohydrodynamics simulations of the solar atmosphere to aid interpretation of solar observations. His scientific work supports a number of NASA-funded solar missions, including the newly launched Solar Dynamics Observatory (SDO), Hinode, the Transition Region and Coronal Explorer (TRACE) and the upcoming Interface Region Imaging Spectrograph (IRIS).

- **Session(s): 2178 - Using GPUs to Track Changes in the Sun (Wednesday, Sept 22, 17:00)**

Chiu, Ting-Wai
Professor (National Taiwan University)

Ting-Wai Chiu is the group leader of TWQCD collaboration in Taiwan. They are using a GPU cluster of more than 200 Nvidia GPUs(GTX480/C2050/S1070/C1060/GTX285) to perform a Monte Carlo simulation of lattice QCD, with 200 Tflops(peak)/36 Tflops(sustained). Their aim is to try to understand the nonperturbative aspects of QCD.

- **Session(s): 2265 - CUDA Centers of Excellence Super-Session IV (Tuesday, Sept 21, 17:00)**

- **2217 - GPU-Based Conjugate Gradient Solvers for Lattice QCD (Wednesday, Sept 22, 16:00)**

Chong, Jike
Principal Software Architect (Parasians, LLC)

Jike Chong is the Founder and Chief Software Architect of Parasians LLC (Parallel Computing Artisans), which specializes in helping clients in compute-intensive industries achieve revolutionary performance in applications directly affecting revenue/cost with GPUs. Jike's prior industry work in parallel computing led to three first-authored US patents at Sun Microsystems and Intel Corporation. He is a Ph.D. researcher at University of California, Berkeley, with B.S. and M.S. in Electrical and Computer Engineering from Carnegie Mellon University.

- **Session(s): 2098 - Enabling On Demand Value-At-Risk for Financial Markets (Thursday, Sept 23, 11:00)**
- **2046 - Efficient Automatic Speech Recognition on the GPU (Thursday, Sept 23, 15:00)**

Chunev, Georgi
Research Assistant (Indiana University)

- **Session(s): 2093 - Computational Photography: Real-Time Plenoptic Rendering (Wednesday, Sept 22, 16:00)**

Civario, Gilles
Head of Capability Computing and Novel Architecture Group (ICHEC)

Gilles is the Head of ICHEC's Capability Computing and Novel Architecture group, and is also involved in the broader aspects of the Centre's mission where his expertise is useful in areas such as highly complicated code installation, debugging or optimisation and hardware evaluation. After completing two Master degrees in Scientific Computing and Algorithms, Gilles joined the R&D team of EDF and was responsible for developing and maintaining nuclear power plant simulation codes, in collaboration with the CEA. From there he worked as a support scientist with CEA/CCRT, one of the largest HPC centres in Europe. For more information, please visit: http://www.ichec.ie/about_us/gilles_civario

- **Session(s): 2086 - GPGPU DL_POLY (Thursday, Sept 23, 16:00)**

Clark, Calvin
Senior Consultant (Microsoft)

Calvin Clark has been at Microsoft for 14 years. He is a Senior Consultant in the Application Development Consulting Group. 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.

- **Session(s): 2147 - GPGPU Development for Windows HPC Server (Tuesday, Sept 21, 15:00)**

Clark, Philip
Reader (Associate Professor) in Particle Physics (University of Edinburgh)

Dr Philip Clark is a reader (associate professor) at the University of Edinburgh. He is the principal investigator for the Edinburgh ATLAS and GridPP particle physics research groups. He is the chairman of the ScotGrid tier-2 compute and data centre. His primary research is in elementary particle physics, but is also interested evolving computer architectures, particularly the advent of many-core and GPGPU devices. He has 672 publications (427 in peer reviewed journals).

- **Session(s): 2135 - Processing Petabytes per Second with the ATLAS experiment at the Large Hadron Collider at CERN (Wednesday, Sept 22, 16:00)**

Clegg, Don
(Supermicro)

- **Session(s): 2293 - Scaling Up and Scaling Out GPUs with Supermicro's Twin™ Architecture (Sponsored by Supermicro) (Wednesday, Sept 22, 11:00)**

Cohen, Jonathan
Senior Research Scientist (NVIDIA Research)

Jonathan Cohen is a Senior Research Scientist with NVIDIA, where he develops methods for using GPUs for scientific computing and physical simulation. Prior to joining NVIDIA, he spent several years working in the Hollywood feature film visual effects industry where he was awarded an Academy Award (Technical Achievement Award) in 2007 for his work on fluid simulation and volumetric modeling for visual effects. He received an undergraduate degree from Brown in Mathematics and Computer Science.

- **Session(s): 2023 - Processing Device Arrays with C++ Metaprogramming (Thursday, Sept 23, 11:00)**
- **2022 - Solving PDEs on Regular Grids with OpenCurrent (Tuesday, Sept 21, 16:00)**

Cole, Brian
Developer (OpenEye Scientific Software)

Brian was an undergraduate at Temple University while doing part-time work at Wyeth. He claimed he did Aikido and after seeing him throw people around it was obvious he had a future at OpenEye. Oh, he does do some programming and many claim he is gifted in that area. At least those who don't want to be thrown around the room. He has a plan to go back for his PhD.

- **Session(s): 2278 - Strategies for Code Encapsulation in GPU Implementations (Thursday, Sept 23, 09:00)**

Collins, Simon
Product Manager (GE Intelligent Platforms)

Simon has been Product Manager for video and graphics products at GE Intelligent Platforms since 1998, during which time he has maintained the company's leading position at the forefront of leading-edge, high-performance commercial technology applied to the rugged defense and aerospace market. Traditional graphics applications have seen the company's products deployed into diverse applications such as cockpit displays in fighter jets, mission computers in helicopters and tanks, and into embedded training systems on naval weapons systems. As the trend towards GPGPU has grown, Simon has defined new products suited to the next generation of Intelligence, Surveillance and Reconnaissance applications, once again taking the performance lead in the rugged marketplace. Prior to taking this role, Simon worked in a number of engineering roles in the nuclear and scientific research industries. He graduated with a B.Sc. in Microelectronics, and an M.Sc. (Eng) in Advanced Manufacturing Technology.

- **Session(s): 2273 - GPUs In the Front Line of our Defenses (Sponsored by GE) (Wednesday, Sept 22, 15:00)**

Corrigan, Andrew
Research Mathematician (Naval Research Laboratory & George Mason University)

Andrew Corrigan is a research mathematician at the Naval Research Laboratory, where he is working on

"What about putting a phone in the plane?"

Laura, age 10

With 3D, your customers are your best designers.

Working in 3D lets you integrate your customers' preferences into your project more easily than ever, even online. Together, you can create, share and experience your ideas - all in 3D. With Dassault Systèmes solutions, your company is empowered by a new, universal language to invent the products of the future.

Discover SolidWorks, CATIA, SIMULIA, DELMIA, ENOVIA and 3DVIA at **www.3ds.com**

DASSAULT SYSTEMES

the GPU implementation of CFD codes and supersonic jet noise reduction. He received his Ph.D. from George Mason University in Computational Mathematics in May 2009. He also performed a post-doc through early 2010 under Prof. Rainald Löhner porting FEFLO to graphics hardware, as well as developing specialized numbering schemes for edge-based unstructured grids on GPUs.

- **Session(s): 2005 - Porting Large-Scale Legacy Fortran Codes (Wednesday, Sept 22, 17:00)**
- **2234 - Unstructured Finite Volume Code on a Cluster with Multiple GPUs per Node (Wednesday, Sept 22, 15:00)**

Cox, Sam
CEO (Milabra)

Sam is an technology entrepreneur and investor, focusing on technology and media ventures. He founded a successful design & software firm in Canada and has consulted for firms in the UK, China, the United States and Canada. He graduated with his MBA from Cass Business School in London, UK with a degree in strategy and completed his undergraduate degree in Art History, Chinese Language and Economics at Queen's University in Canada.

- **Session(s): 4001 - Emerging Companies: CEO on Stage featuring Elemental Technologies, Geomerics, and Milabra (Wednesday, Sept 22, 11:00)**
- **4003 - Emerging Companies Summit Panel: GPUs for Computer Vision (Wednesday, Sept 22, 15:00)**

Crivelli, Luis
Director Solver Development (Dassualt Systems Simulia Corporation)

PhD in Aerospace Engineering. Director Of Solver Development at Dassault Systems Simulia Corporation. 16+ years experience in High Performance Computing and Parallel Computing.

- **Session(s): 2155 - GPGPU in the real world. The ABAQUS experience (Thursday, Sept 23, 14:00)**

Cui, Jingyu
Graduate Student (Stanford University)

Jingyu Cui received his B.E (2005) and M.S. (2008) degree from Tsinghua University, and M.S. (2010) degree from Stanford University. He is currently pursuing his Ph.D. degree working on high speed dynamic 4-dimensional medical imaging. Jingyu has published 8 peer-reviewed papers in conference proceedings as the leading author, 2 journal articles, and a book chapter. He also holds a US patent. Jingyu worked with Microsoft and Google, and made important contributions to several products.

- **Session(s): 2211 - Modern Architecture for Massively Parallel Medical Tomographic Image Reconstruction on a GPU Cluster (Wednesday, Sept 22, 15:00)**

Cui, Xiaohui
Research Scientist (Oak Ridge National Laboratory)

Dr. Xiaohui Cui is the scientist staff of the Computational Sciences & Engineering Division, Oak Ridge National Laboratory of Department of Energy and the adjunct associate professor of University of Louisville in Kentucky. His research interests include swarm intelligence, agent based modeling and simulation, GPU computing, and information retrieval. His research has been reported by MSNBC, New Scientist etc. In 2008 and 2009, he received the Department of Energy Outstanding Mentor Awards.

- **Session(s): 2052 - Power Management Techniques for Heterogeneous Exascale Computing (Tuesday, Sept 21, 16:00)**

Curry, Matthew
A Highly Reliable RAID System Based on GPUs (Sandia National Laboratories and the University of Alabama at Birmingham)

Matthew Curry is a Ph.D. candidate at the University of Alabama at Birmingham. He is a member of the High Performance Computer Laboratory in the Computer and Information Sciences Department under the advisement of Dr. Anthony Skjellum. He is interested in GPU computing, operating systems, and high performance storage.

- **Session(s): 2205 - A Highly Reliable RAID System Based on GPUs (Tuesday, Sept 21, 17:00)**

Dammertz, Holger
PhD Student (Ulm University)

As a PhD student at Ulm University, Germany, my main focus of research is fast Ray Tracing for Global Illumination and related rendering techniques. I am also researching quasi-Monte Carlo methods and parallel algorithms for graphics.

- **Session(s): 2136 - Pseudo Random Number Generators for Massively Parallel Apps (Thursday, Sept 23, 16:00)**

Dasgupta, Aniruddha
Graduate Student (Georgia Institute of Technology)

Aniruddha Dasgupta is currently working towards a Mater's Degree from the department of Electrical and Computer Engineering at Georgia Tech.
His areas of research interests are GPGPU and GPU architecture.

- **Session(s): 2201 - A Case Study of Accelerating Matlab Based Applications using GPUs (Wednesday, Sept 22, 16:00)**

Davidson, Andrew
Graduate Student (University of California, Davis)

Andrew Davidson is a graduate student in the Computer Engineering Department at the University of California, Davis. His research interests include data-parallel algorithms and primitives, numerical methods, and auto-tuning. He is also a developer for the CUDA Parallel Primitives Library (CUDPP) .

- **Session(s): 2085 - Tridiagonal Solvers: Auto-Tuning and Optimizations (Tuesday, Sept 21, 15:00)**

Davis, Nolan
Research Scientist (SAIC)

Nolan R. Davis is a Senior Scientist with SAIC in San Diego. He holds a doctorate in physics from the University of Texas at Dallas, and has spent over 25 years working in physics research, signal and image processing, and high performance computing. He has worked with large corporations and laboratories including SAIC, Lockheed-Martin, the Johns-Hopkins Applied Physics Laboratory, the Naval Research Laboratory, and Walt Disney Feature Animation.

- **Session(s): 2100 - Hybrid GPU/Multicore Solutions for Large Linear Algebra Problems (Thursday, Sept 23, 16:00)**

Dean, Loren
Director of Engineering, MATLAB Products (MathWorks)

Loren Dean is a Director in the MATLAB® development organization. He has responsibility for MathWorks parallel computing products, the Test & Measurement application area and the eProducts and Services organization. Loren has been with MathWorks since 1995. Prior to joining MathWorks, Loren worked for AlliedSignal Aerospace, performing systems analysis and integration for aircraft engines, with extensive use of

MATLAB and Simulink®. Loren has a B.S. and an M.S. in Aeronautical Engineering from Purdue University and an M.B.A. from Northeastern University.

- **Session(s): 2267 - GPU Computing with MATLAB® (Tuesday, Sept 21, 11:00)**

Dean, Tom
Research Scientist (Google Inc.)

Tom Dean is a full-time research scientist at Google in Mountain View, California. From 1993 to 2007 he was Professor of Computer Science and Cognitive and Linguistic Sciences at Brown University. He received his B.A. in mathematics from Virginia Polytechnic Institute & State University in 1982 and his M.Sc. and Ph.D. in computer science from Yale University in 1984 and 1986. His research interests include automated planning and control, computational biology, machine learning, neural modeling, probabilistic inference, robotics and spatial and temporal reasoning. For more information, please visit: http://www.cs.brown.edu/people/tld/pages/bio.html

- **Session(s): 2132 - Accelerating Biologically Inspired Computer Vision Models (Tuesday, Sept 21, 11:00)**
- **4003 - Emerging Companies Summit Panel: GPUs for Computer Vision (Wednesday, Sept 22, 15:00)**

Decrem, Peter
Director, Rates Products (Quantifi)

Peter Decrem heads the Rates Group at Quantifi. As Director, Peter is responsible for managing the product development process of all Rates Solutions within the Quantifi product suite. Peter started in Research and Technology at Bear Stearns and Deutsche Bank. He traded fixed income derivatives, government bonds and agencies for Lehman Brothers and Salomon Brothers. He was responsible for fixed income derivatives trading desk for a number of European banks. Most recently he refocused on technology and specifically concentrated on machine learning and high frequency trading on parallel systems prior to his joining of Quantifi.

- **Session(s): 2040 - Derivatives & Bond Portfolio Valuation in a Hybrid CPU/GPU Environment (Thursday, Sept 23, 14:00)**
- **2297 - Developing CUDA Accelerated .NET Plugins for Microsoft Excel (Tuesday, Sept 21, 17:00)**

Deguy, Sébastien
Founder and CEO (Allegorithmic)

Dr. Sébastien Deguy is the CEO of Allegorithmic, the company behind the Substance procedural textures authoring and rendering system. 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.

- **Session(s): 4002 - Emerging Companies: CEO on Stage featuring Allegorithmic SAS, Bunkspeed, and miGenius (Wednesday, Sept 22, 14:00)**

Del Sordo, Giancarlo
Chief Developer and Product Manager (Acustica Audio)

Giancarlo Del Sordo is the main developer at Acustica Audio. He specialized in Computer Engineering at the University of Pisa. He also worked in the IT department of Intesa Sanpaolo (first Italian bank), in NEXTRA Asset Management and in Cap Gemini Ernst&Young consulting.

- **Session(s): 2076 - Implementing CUDA Audio Networks (Thursday, Sept 23, 09:00)**

Deng, Yangdong
Associate Professor (Tsinghua University)

Yangdong (Steve) Deng received his Ph.D. degree in Electrical and Computer Engineering from Carnegie Mellon University, Pittsburgh, PA, in 2006. He received his ME and BE degrees in Electronic Department from Tsinghua University, Beijing, in 1998 and 1995, respectively.

- **Session(s): 2081 - Morphing a GPU into a Network Processor (Thursday, Sept 23, 15:00)**
- **2264 - CUDA Centers of Excellence Super-Session III (Tuesday, Sept 21, 16:00)**

Dewaele, Ronny
Director Technology Center (Barco)

Ronny Dewaele is responsible for the corporate Technology Center for Networked Visualization in Barco. Ronny and his team focus on exploring new technologies in the domain of network centric video processing.Ronny Dewaele has a Master degree in Computer Science and Applied Mathematics from KU Leuver, Leuven, Belgium. He lives and works in Belgium.

- **Session(s): 2095 - Building High Density Real-Time Video Processing Systems (Thursday, Sept 23, 16:00)**

D'Hondt, Maja
Program Manager (imec)

Maja D'Hondt is program manager at imec, Europe's largest independent research center in nanoelectronics and nano-technology. Her team develops middleware for embedded and high performance systems, such as Barco's GPU-enabled video processing server. Maja holds a PhD from the Vrije Universiteit Brussel in Belgium (2004), then spent some time in Amsterdam on a research project with ASML, and obtained a full research position at INRIA in France.

- **Session(s): 2121 - Maximizing Throughput of Barco's GPU-Enabled Video Processing Server (Thursday, Sept 23, 14:00)**

Diamos, Gregory
PhD Student (Georgia Institute of Technology)

Gregory Diamos is a PhD student at the Georgia Institute of Technology, under the direction of Professor Sudhakar Yalamanchili. He received his B.S. and M.S. in Electrical Engineering from the Georgia Institute of Technology in 2006 and 2008, respectively. His current research interests follow the industry shift from ILP to many core architectures, where the ability to tightly integrate heterogeneous architectures offers the potential for dramatic improvements in efficiency at the cost of increased design complexity.

- **Session(s): 2210 - GPU-Ocelot: An Open Source Debugging and Compilation Framework for CUDA (Thursday, Sept 23, 14:00)**

Dick, Christian
PostGraduate Fellow (Technische Universität München)

Christian Dick received his Diploma in Computer Science from the Technische Universität München in July 2007 with honors. Since August 2007, he has been a postgraduate fellow in the Computer Graphics and Visualization Group at the Technische Universität München under the supervision of Prof. Dr. R. Westermann. His research interests include highly responsive, physics-based simulation of deformable models on high-resolution hierarchical representations, and the visualization of large-scale scientific data sets such as terrain and volume data.

- **Session(s): 2137 - CUDA for Real-Time Multigrid Finite Element Simulation of Soft Tissue Deformations (Wednesday, Sept 22, 14:00)**

Dimitrovici, Dragan
(XENON Systems Pty Ltd)

- **Session(s): 2301 - GPU Cluster Computing: Accelerating Scientific Discovery (Thursday, Sept 23, 09:00)**

Dixon, Matthew
Professor (UC Davis)

Matthew Dixon is a Krener assistant professor in the mathematics department at UC Davis. He received his Ph.D. in applied mathematics from Imperial College (UK) in 2007 and has since held postdoctoral appointments with the Institute for Computational and Mathematical Engineering at Stanford University and the Department of Computer Science at UC Davis. He has also worked as a quantitative risk analyst for a number of leading investment banks and consulted to the Bank for International Settlements.

- **Session(s): 2098 - Enabling On Demand Value-At-Risk for Financial Markets (Thursday, Sept 23, 11:00)**

Domine, Sebastien
Sr. Dir. Developer Tools (NVIDIA)

Sébastien Dominé is the 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 Diplôme d'Ingénieur in Computer Science from EPITA, Paris, France.

- **Session(s): 2151 - Parallel Nsight: Analyzing and Optimizing Massively Parallel Applications [Advanced] (Tuesday, Sept 21, 16:00)**
- **2150 - Parallel Nsight: Debugging Massively Parallel Applications [Advanced] (Tuesday, Sept 21, 14:00)**

Donovan, Scott
System Architect (Citadel Investment Group)

Mr. Donovan has a masters degree in computer science and over 20 years IT experience. Throughout his career he has held positions at exchanges, investment banks, and hedge funds. He is currently a System Architect at Citadel where his main area of focus is accelerating financial models with a combination of grid computing, virtualization, and CUDA / OpenCL.

- **Session(s): 2033 - Integrating GPGPU Accelerated Pricing Models into an Existing Financial Services Infrastructure (Thursday, Sept 23, 09:00)**

Doran, Chris
Founder and Chief Operating Officer (Geomerics)

Dr. Chris Doran is Founder and Chief Operating Officer at Geomerics. He is a leading research scientist with 20 years experience in applied mathematics and theoretical physics, and is the author of a major book on geometry and physics and of over 50 papers. Chris is a regular speaker at major international conferences, including SIGGRAPH, Develop, Nordic Game and Montreal Games Summit. Chris is also Director of Studies in Physics for Sidney Sussex College, Cambridge.

- **Session(s): 4001 - Emerging Companies: CEO on Stage featuring Elemental Technologies, Inc., Geomerics, and Milabra (Wednesday, Sept 22, 11:00)**

Dyken, Chris
Research Scientist (SINTEF)

Christopher Dyken got his Ph.D. in Computational Geometry at the University of Oslo in 2008, where he also lectured computer graphics for four years. He currently holds a position as a research scientist at the Heterogeneous Computing Group at SINTEF ICT Applied Mathematics, focusing on algorithms for heterogeneous architectures and visualization techniques.

- **Session(s): 2020 - GPU-Accelerated Data Expansion for the Marching Cubes Algorithm (Wednesday, Sept 22, 16:00)**

Edgar, Richard
Assistant in Neuroscience (Massachusetts General Hospital, Harvard University)

With a background in theoretical astrophysics, Richard Edgar is now working at Harvard and MGH on a variety of projects in need of GPU acceleration.

- **Session(s): 2001 - Acceleration of the Freesurfer Suite for Neuroimaging Analysis (Thursday, Sept 23, 10:00)**

Enderle, Rob
Analyst (Enderle Group)

Rob is President and Principal Analyst of the Enderle Group, a forward looking emerging technology advisory firm. Recognized as one of the best general Inquiry Analysts in the world, Rob specializes in providing rapid perspectives and suggested tactics and strategies to a large number of clients dealing with rapidly changing global events. Rob lives emerging technology and has a passion for personal technology and market strategy. Rob trained as a news anchor and co-hosted CNET radio during the 90s, has been widely used by both local and national news TV and radio programs, and has been identified as one of the worlds' most influential technology analysts. Currently Rob appears semi-weekly for a tech segment on WSJ radio and writes for ECT (TechNewsWorld, eCommerce Times, Linux World, MacNewsWorld), Dark Reading, Digital Trends, TGDaily, ITBusiness Edge and Datamation. Before founding the Enderle Group in 2003 Rob was the Senior Research Fellow for Forrester Research and the Giga Information Group. While there he ran the eCommerce, Security, and Mobile research practices.

- **Session(s): 4007 - Emerging Companies: CEO on Stage featuring Aqumin, RTT, and Scalable Display Technologies (Thursday, Sept 23, 10:00)**
- **4008 - Emerging Companies: CEO on Stage featuring ICD and Universal Robotics (Thursday, Sept 23, 11:00**

Engsig-Karup, Allan Peter
Assistant Professor, Scientific Computing (Technical University of Denmark)

MSc, PhD in Applied math.Learn more: http://www.imm.dtu.dk/~apek Involved in research related to utilization of GPUs for Scientific Computing. Learn more: http://gpulab.imm.dtu.dk. Research interest in Computational Fluid Dynamics, High-Performance Computing, Coastal Engineering, Scientific Computing, Numerical analysis.

- **Session(s): 2103 - Development of an Efficient GPU-Accelerated Model for Fully Nonlinear Water Waves (Tuesday, Sept 21, 15:00)**

Fahmy, Hany
Director, SI/EMC Engineering (NVIDIA)

- **Session(s): 2080 - Tackling Multi-Gigabit Design Challenges with a Practical Virtual EMI/ESD Lab (Wednesday, Sept 22, 15:00)**

Farber, Robert
Senior Scientist (PNNL)

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

- **Session(s): 2119 - Supercomputing for the Masses: Killer-Apps, Parallel Mappings, Scalability and Application Lifespan (Tuesday, Sept 21, 11:00)**

Fascione, Luca
Senior Research and Development Engineer (Weta Digital)

Luca Fascione is a Senior Research and Development Engineer at Weta Digital. He is currently focused on advanced rendering solutions and the development of innovative and forward-looking film pipelines that harness and expand on the best technologies available around the world. In addition to his work at Weta Digital, Luca spent time as a software engineer at Pixar and Vanguard Animation Studios. He also taught courses on General Computer Graphics and Rendering Languages as a professor at the University of Rome.

- **Session(s): 2305 - PantaRay: Accelerating Out-Of-Core Ray Tracing of Sparsely Sampled Occlusion (Wednesday, Sept 22, 10:00)**

Fassold, Hannes
Scientist (JOANNEUM RESEARCH)

Hannes Fassold finished his study of Technical Mathematics in Graz, Austria, in 2004. Since July, 2004, he has been working as a scientist at JOANNEUM RESEARCH. His major work fields are the development of algorithms for digital film restoration and video quality analysis, using image processing and computer vision methods.

- **Session(s): 2029 - Computer Vision Algorithms for Automating HD Post-Production (Wednesday, Sept 22, 15:00)**

Fatica, Massimiliano
Manager (NVIDIA)

Massimiliano is a manager of the Tesla Performance Group at NVIDIA where he works in the area of GPU computing (high-performance computing and clusters). He holds a laurea in Aeronautical Engineering and a PhD in Theoretical and Applied Mechanics from the University of Rome "La Sapienza".

- **Session(s): 2057 - CUDA-Accelerated LINPACK on Clusters (Tuesday, Sept 21, 14:00)**

Fernandez, Mark
Computer Scientist (Dell)

As a HPC Computer Scientist in Dell's Advanced Systems Group, Dr. Fernandez supports HPC customer/end-user technical efforts at Dell. Dr. Fernandez is responsible for working with customers to capture user requirements and incorporating those requirements into future Dell HPC systems and solution. He also works closely with Dell Engineering throughout the product development cycle.

- **Session(s): 2287 - Internal GPUs on Dedicated x16 Slots - Are They Needed For HPC? (Sponsored by Dell) (Wednesday, Sept 22, 14:00)**

Frauenhofer, Bill
Managing Director (Citigroup Global Markets)

William Frauenhofer is a Managing Director in Citigroup's Technology Group based in Palo Alto, CA. Bill joined the firm in 1996 and worked in the Real Estate & Lodging practice until 2000 when he relocated to San Francisco to join Citigroup's Technology team. Bill is currently the head of Citigroup's global semiconductor team and has worked on a broad base of transactions including M&A, LBOs, and a number of public and private financing transactions, including equity, equity-linked, high yield and investment grade capital. Bill has recently advised clients including: Amkor Technology, Freescale Semiconductor, Atmel, Qimonda, Infineon, Nvidia, Wolfson Microelectronics, NXP, etc. Bill received a M.B.A. with honors from the Stern School of Business at New York University and a B.B.A. in Finance from Loyola College in Maryland. Prior to joining Salomon Brothers, he was an Analyst at The Warwick Group, a boutique investment bank specializing in mid-cap mergers and acquisitions.

- **Session(s): 4009 - Emerging Companies Summit Panel: The "New Normal" For Building Emerging Companies Based On Disruptive Technologies (Thursday, Sept 23, 14:00)**

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 holds a Ph.D. in Electrical and Computer Engineering from the University of Toronto. He currently works at NVIDIA examining computer vision and image processing on graphics hardware.

- **Session(s): 2209 - Accelerating Computer Vision on the Fermi Architecture (Thursday, Sept 23, 14:00)**

Ganesan, Narayan
Research Scientist (University of Delaware)

Dr.Ganesan received his Ph.D from Washington University in St.Louis. His dissertation was on Quantum-Information and Decoherence free Quantum-Computation. His current research is on scientific and High-performance computing on GPUs. His recent work on parallelizing sequentially dependent recurrence computations, technique applied to the popular protein motif-finding problem, delivers unprecedented performance compared to current MPI-GPU-HMMER. He is currently the lead scientist behind the development of an optimized Molecular Dynamics Simulation package based on CHARMM force field. The package, delivers highly-competitive performance and is currently used to study behavior of large lipid-membranes, protein-ligand interactions and multiscale modeling at the University of Delaware.

- **Session(s): 2034 - Reformulating Algorithms for the GPU (Wednesday, Sept 22, 11:00)**
- **2035 - Simulations of Large Membrane Regions (Wednesday, Sept 22, 11:30)**

Gateau, Samuel
Developer Technology Engineer (NVIDIA)

Sam is a member of the Content & Technology Engineer group at NVIDIA, who spends his energy and creativity pushing pixels and teaching high-end, real-time computer graphics to Games, DCC, and CAD developers. Before that, he was enjoying the sun of Toulouse in France working in the virtual reality industry on extravagant visual simulations, navigation systems, showrooms and museum applications.

- **Session(s): 2010 - Implementing Stereoscopic 3D in Your Applications (Pre-Conference Tutorial) (Monday, Sept 20, 16:00)**

Gaudlitz, Daniel
Project Manager (FluiDyna)

Daniel Gaudlitz received his diploma in Aeronautics at the Technical University Dresden (equiv. to Master of Science)in 2003. In 2008, he earned his PhD at the Technical University Munich as a member of the research group for numerical simulations in fluid mechanics led by Prof. N. A. Adams. Since 2009, he has obtained his post-doc at the Institute of Aerodynamics of Technical University Munich and is working as the Project Manager and Manager Research & Development at FluiDyna GmbH.

- **Session(s): 2206 - Accelerated Computational Fluid Dynamics Employing GPUs (Thursday, Sept 23, 09:00)**

Gebbie, Nicholas
Senior Graphics Programmer (Bunkspeed)

- **Session(s): 2074 - Driving a Product from Rasterization to Ray Tracing: The Developer Experience (Tuesday, Sept 21, 15:00)**

Ge, Wei
Professor (Institute of Process Engineering, Chinese Academy of Sciences)

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

- **Session(s): 2263 - CUDA Centers of Excellence Super-Session II (Tuesday, Sept 21, 15:00)**
- **2286 - Towards Peta-Scale Green Computation - Applications of the GPU Supercomputers in the Chinese Academy of Sciences (CAS) (Wednesday, Sept 22, 11:00)**

Gelado, Isaac
Lecturer and Researcher (Universitat Politecnica de Catalunya)

Isaac Gelado is an Assistant Professor at the Computer Architecture Department in Universitat Politecnica de Catalunya at Barcelona. Isaac Gelado holds a Master's degree on Telecommunications Engineering from Universidad de Valladolid, and will get a PhD degree from Universitat Politecnica de Catalunya in July, 2010.

- **Session(s): 2156 - GMAC: Global Memory For Accelerators (Thursday, Sept 23, 09:00)**

Georgiev, Todor
Senior Research Scientist II (Adobe Systems)

Todor Georgiev is a Senior Research Scientist at Adobe Systems, working closely with the Photoshop group. His contributions are often based on transfer of mathematical methods from physics to image processing and vision. Currently he is focusing on developing cameras for radiance capture and interactive plenoptic / lightfield rendering.

- **Session(s): 2093 - Computational Photography: Real-Time Plenoptic Rendering (Wednesday, Sept 22, 16:00)**

Georgii, Joachim
PostDoc (Technische Universität München)

Joachim Georgii is a PostDoc at the computer graphics and visualization group headed by Professor Rüdiger Westermann at the Technische Universität München. In 2007, he received a PhD in computer science at the Technische Universität München. His research interests are simulation and visualization techniques, mainly for deformable bodies, as well as their efficient implementations based on multigrid methods.

- **Session(s): 2137 - CUDA for Real-Time Multigrid Finite Element Simulation of Soft Tissue Deformations (Wednesday, Sept 22, 14:00)**

Gianos, Flip
General Partner (InterWest Partners)

Philip "Flip" Gianos has been part of InterWest's IT team since 1982. With a background in engineering, he has invested in multiple areas of information technology, including semiconductors, computing and networking equipment, and infrastructure and applications software. He is chairman of the board of Xilinx (XLNX), a publicly held company, and is also a board member of several privately held companies, including: Bivio Networks, Brand.net, Convey Computer, and SpectraLinear. Gianos also serves on the advisory board of Storm Ventures II, and is a past president of the Western Association of Venture Capitalists. Prior to joining InterWest, Gianos was with IBM for eight years in engineering management. He managed both chip design and systems integration for several IBM office automation products. Gianos earned his M.B.A. from Harvard University and received his M.S. and B.S. in electrical engineering from Stanford University. He has one international and two U.S. patents.

- **Session(s): 4004 - Emerging Companies: CEO on Stage featuring Cooliris, empulse GmbH, and Playcast Media Systems (Wednesday, Sept 22, 16:00)**
- **4005 - Emerging Companies: CEO on Stage featuring Jedox Business Intelligence, Rocketick, and Softkinetic (Wednesday, Sept 22, 17:00)**

Gokhale, Nachiket
Senior Research Engineer (Weidlinger Associates Inc)

Nachiket Gokhale, Ph.D. is actively involved in multi-disciplinary research and development efforts at Weidlinger Associates, Inc. (WAI). Under a DARPA SBIR, he was involved in the development of a GPU-enabled version of WAI's explicit time domain, commercial finite element software NLFLEX. He is interested in the development of fast GPU enabled algorithms and computer codes for the high-fidelity solution of challenging problems in computational mechanics with an emphasis on transient phenomena in structural mechanics, such as shock and blast; FEM simulation of ultrasound with an emphasis on biomedical imaging and therapy; and the computational design of novel acoustic metamaterials. His project experience includes SBIR and STTR research for DARPA, ONR, and various protective design and structural engineering efforts involving large finite element analyses. He earned his Ph.D . and M.S. both in Mechanical Engineering from Boston University where his work involved the finite element solution of linear and non-linear inverse problems in biomechanical imaging.

- **Session(s): 2061 - Accelerating Explicit FEM Shock & Blast Simulations (Thursday, Sept 23, 10:30)**

Goldsmith, Kevin
Senior Engineering Manager (Adobe Systems, Incorporated)

Kevin Goldsmith is the Manager of the Adobe Image Foundation team. This team created a domain specific language: Pixel Bender; which allows for highly optimized parallel signal-processing computation and a dynamic runtime environment designed to scale from current to future highly-parallel heterogeneous hardware. AIF is currently part of many of Adobe's

flagship applications. Kevin has over 18 years in the computer industry at companies such as Silicon Graphics, Microsoft, IBM Research, (Colossal) Pictures and others.

- **Session(s): 2051 - GPGPU in Commercial Software: Lessons From Three Cycles of the Adobe Creative Suite (Thursday, Sept 23, 11:00)**

Gonzalez, Alberto

Professor (Universidad Politecnica de Valencia)

Alberto Gonzalez was born in Valencia, Spain, in 1968. He received the Ingeniero de Telecomunicacion degree from the Universidad Politecnica de Catalonia, Spain in 1992, and Ph.D degree from de Universidad Politecnica de Valencia (UPV), Spain in 1997. His dissertation was on adaptive filtering for active control applications. From January 1995, he visited the Institute of Sound and Vibration Research, University of Southampton, UK, where he was involved in research on digital signal processing for active control. He is currently heading the Audio and Communications Signal Processing Research Group (www.gtac.upv.es) that belongs to the Institute of Telecommunications and Multimedia Applications (i-TEAM, www.iteam.es). Dr. Gonzalez serves as Professor in digital signal processing and communications at UPV where he heads the Communications Department (www.dcom.upv.es) since April 2004. He has published more than 80 papers in journals and conferences on signal processing and applied acoustics. His current research interests include fast adaptive filtering algorithms and multichannel signal processing for communications and 3D sound reproduction.

- **Session(s): 2116 - Real-time Multichannel Audio Convolution (Thursday, Sept 23, 10:00)**

Gottbrath, Chris

Principal Product Manager (TotalView Technologies, Inc., a Rogue Wave Software company)

Chris Gottbrath is principal product manager for the TotalView Debugger product line at Rogue Wave Softare. His work is focused on making it easier for programmers, scientists and engineers to solve even the most complex bugs and get "back to work." He has pursued this goal in a variety of customer-focused technical roles with the TotalView team over the last seven years. Prior to that, as a graduate student of astrophysics at the University of Arizona in Tucson, he wrote cosmological simulations (with the occasional bug) using C and MPI on a small-scale Beowulf cluster. Chris is a regular contributor to HPC and software development industry conferences worldwide.

- **Session(s): 2299 - Integrating CUDA BLAS with IMSL Fortran (Tuesday, Sept 21, 14:00)**
- **2251 - TotalView Debugger for CUDA (Wednesday, Sept 22, 15:00)**

Govett, Mark

Chief, Advanced Computing Section (NOAA Earth System Research Laboratory)

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

- **Session(s): 2276 - Using GPUs to Run Next-Generation Weather Models (Tuesday, Sept 21, 14:00)**

Gu, Henry

CTO (GIC)

Dr. Henry Gu is the founder of Green International Consulting, a global software development, consulting and outsourcing company. Gu was the GM of Thomson Corporate Research Center in Burbank before founding GIC. Before joining Thomson, Gu was CTO and VP of da Vinci Systems. He led the company's R&D team in developing generations of color corrector for the entertainment industry. In 2001, Gu received the Primetime Emmy Award for Outstanding Achievement in Engineering Development.

- **Session(s): 2043 - Disparity Map Generation (Thursday, Sept 23, 11:00)**

Gupta, Kshitij

Graduate Student Researcher (UC Davis)

Kshitij Gupta is a Ph.D. candidate in the Department of Electrical & Computer Engineering at UC Davis. He is interested in a variety of application domains like audio, image, and video. His primary interests are in exploring novel ways of transforming today's high-performance algorithms onto emerging low-end, low-power, hybrid (CPU/GPU/DSP/ASIP) processors targeted towards mobile and automotive platforms. In his spare time, he likes procrastinating about novel user-interfaces, and hopes to work more actively on it some day. Kshitij received his Masters in Electrical & Computer Engineering from University of Pittsburgh (PA, USA), and his Bachelors in Electronics & Communication Engineering from Osmania University (Hyderabad, India).

- **Session(s): 2175 - Hello GPU: High-Quality, Real-Time Speech Recognition on Embedded GPUs (Thursday, Sept 23, 14:00)**

Gupta, Rohit

Researcher/Teacher (Delft University Of Technology)

Rohit Gupta recently completed his Masters Study at TUDelft (Aug 2010) in Computer Engineering. Rohit's masters thesis subject was in the domain of CFD on GPU computing. Previously, he has worked in the Domain of Embedded Systems for 4 years after completing his bachelors degree in India.

- **Session(s): 2049 - Deflated Preconditioned Conjugate Gradient on the GPU (Wednesday, Sept 22, 14:30)**

Haines, Karen

Professor (WASP/The University of Western Australia)

Dr. Haines completed her PhD in Electrical Engineering at the University of New Mexico. She received her Masters in Engineering at Carnegie Mellon University and her Bachelor of Arts in Mathematics at the University of California, San Diego. Her PhD research efforts have lead to the development of a parallel motion detection algorithm, which is based on the fly's visual processing system. The resulting model is suitable for robotic or computer vision applications. This work relied on distributed parallel programming and advanced scientific visualization methods.

- **Session(s): 2252 - Simulating Housefly Vision Elements Using OpenCL (Wednesday, Sept 22, 16:00)**

Han, Jeff

Founder, Chief Scientist (Perceptive Pixel)

Jeff Han is the founder and chief scientist of Perceptive Pixel A TED speaker in 2006, and named to the Time 100 most influential persons list in 2008, Jeff continues to contribute frequently to the research communities. Jeff's formal training was in electrical engineering and computer science at Cornell University, where he worked on the innovative CU-SeeMe videoconferencing system.

- **4011 - Emerging Companies: CEO on Stage featuring Cinnafilm, Inc., Perceptive Pixel and Total Immersion (Thursday, Sept 23, 16:00)**

Hansen, Charles
Professor (University of Utah)

Charles (Chuck) Hansen is a Professor of Computer Science and an Associate Director of the Scientific Computing and Imaging Institute at the University of Utah. Chuck Hansen has published over 100 peer reviewed journal and conference papers and has been a co-author on three papers recognized with "Best Paper Awards" at the IEEE Visualization Conference (1998, 2001, 2002). He was co-author on the Best Paper at IEEE Pacific Visualization 2010. He was awarded the IEEE Technical Committee on Visualization and Graphics "Technical Achievement Award" in 2005 in recognition of seminal work on tools for understanding large-scale scientific data sets.

- **Session(s): 2264 - CUDA Centers of Excellence Super-Session III (Tuesday, Sept 21, 16:00)**

Hardy, Quentin
National Editor (Forbes Magazine)

Quentin Hardy is National Editor for Forbes Media, responsible for cover stories and features for Forbes magazine, along with stories, a blog and video interviews for the Forbes.com website. Mr. Hardy is a regular on "Forbes on Fox," a weekly business news show on Fox News Channel, as well as shows on CNBC, Bloomberg and other television channels. He hosts numerous panels on technology and business both independently and at Forbes events around the U.S. and overseas.

- **Session(s): 4006 - Fireside Chat with Jen-Hsun Huang - Co-founder & CEO, NVIDIA (Thursday, Sept 23, 09:00)**

Harris, Mark
Senior Developer Technology Engineer (NVIDIA)

Mark Harris is a senior developer technology engineer at NVIDIA, where he works with developers around the world on software for computer graphics and high-performance computing. His research interests include parallel computing, general-purpose computation on GPUs, physically based simulation, real-time rendering, and gastronomy. Mark earned his PhD in computer science from the University of North Carolina at Chapel Hill in 2003. He founded and maintains GPGPU.org, a web site dedicated to general-purpose computation on GPUs.

- **Session(s): 2084 - State of the Art in GPU Data-Parallel Algorithm Primitives (Tuesday, Sept 21, 17:00)**

Harrison, Brian
(NVIDIA)

- **Session(s): 2024 - NVIDIA Acceleration Engines Overview (Pre-Conference Tutorial) (Monday, Sept 20, 13:00)**
- **2308 - Building Cutting-Edge Realtime 3D Applications with NVIDIA SceniX (Wednesday, Sept 22, 10:00)**

Hart, Evan
Software Engineer (NVIDIA)

Evan presently works as a developer technology engineer for NVIDIA. He has worked for over a decade to improve the quality and performance of 3D rendering in applications. He has worked with a diverse set of application domains, including CAD, DCC, visualizaton, and games. Evan received is BS from The Ohio State University.

- **Session(s): 2152 - Using Virtual Texturing to Handle Massive Texture Data (Tuesday, Sept 21, 14:00)**

Härter, Daniel
(University of Freiburg, IMTEK, Laboratory for Process Technology)

From 2000 to 2005, Daniel studied Microsystems Technology at the University of Freiburg, where he wrote his Diploma thesis about miniaturized illumination concepts at the Laboratory of Process Technology.

- **Session(s): 2065 - Massively Accelerating Iterative Gauss-Newton Fitting (Wednesday, Sept 22, 11:00)**

Hayes, David
CEO (ICD)

David is the Chief Executive Officer of ICD and has a wealth of experience in the mobile and consumer electronics markets. David founded Velocity Mobile - an ICD collaboration with Inventec in 2007. Prior to ICD, David was Chief Executive Officer of A Living Picture (ALP), a company he formed to develop "Momento," an advanced, digital picture frame technology. ALP was later acquired by i-mate in December 2006 and David became the Chief Technology Officer of i-mate following the acquisition. Prior to ALP and i-mate, David was the founder and Chief Executive Officer of DAT plc, one of the early pioneers of over-the-air device management for Windows Mobile phones. Here, David forged strong links within the Windows Mobile community, from ODMs through to operators and retailers.

- **Session(s): 4008 - Emerging Companies: CEO on Stage featuring ICD and Universal Robotics (Thursday, Sept 23, 11:00)**

Herbert, Garrett
Partner, M&A Transaction Services (Deloitte & Touche LLP)

Garrett Herbert is a partner and leads the Silicon Valley M&A Transaction Services practice in San Jose, CA and is the national leader for M&A Transaction Services for the Telecom, Media, and Technology Industry group. Garrett has over 18 years of professional experience that last 13 years has been dedicated to M&A. He has extensive transaction experience in advising financial and strategic buyers on due diligence, accounting structuring and financial reporting aspects of transactions in technology, semiconductors, and software transactions both domestically and internationally. In addition to his M&A experience with Deloitte, Garrett has M&A experience as an investment professional in industry with Mentmore Holdings Corporation (a private equity group), Stellex Technologies (wireless communications equipment), and Register.com (NASDAQ: RCOM) where he was responsible for target evaluation, due diligence, divestitures, and post-transaction integration.

- **Session(s): 4009 - Emerging Companies Summit Panel: The "New Normal" For Building Emerging Companies Based On Disruptive Technologies (Thursday, Sept 23, 14:00)**

Herbst, Jeff
Vice President of Business Development (NVIDIA)

Jeff is the Vice President of Business Development at NVIDIA Corporation, the world leader in visual computing technologies (and inventor of the GPU). 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.

- **Session(s): 4000 - Emerging Companies Summit Opening Address (Wednesday, Sept 22, 10:00)**
- **4001 - Emerging Companies: CEO on Stage featuring Elemental Technologies, Geomerics, and Milabra (Wednesday, Sept 22, 11:00)**
- **4002 - Emerging Companies: CEO on Stage featuring Allegorithmic SAS, Bunkspeed, and miGenius (Wednesday, Sept 22, 14:00)**
- **4004 - Emerging Companies: CEO on Stage featuring Cooliris, empulse GmbH, and Playcast (Wednesday, Sept 22, 16:00)**
- **4005 - Emerging Companies: CEO on Stage featuring Jedox Business Intelligence, Rocketick, and Softkinetic (Wednesday, Sept 22, 17:00**
- **4007 - Emerging Companies: CEO on Stage featuring Aqumin, RTT, and Scalable Display Technologies (Thursday, Sept 23, 10:00)**
- **4008 - Emerging Companies: CEO on Stage featuring ICD and Universal Robotics (Thursday, Sept 23, 11:00)**
- **4009 - Emerging Companies Summit Panel: The "New Normal" For Building Emerging Companies Based On Disruptive Technologies (Thursday, Sept 23, 14:00)**
- **4010 - Emerging Companies: CEO on Stage featuring NaturalMotion Ltd, OptiTex, and Useful Progress (Thursday, Sept 23, 15:00)**
- **4011 - Emerging Companies: CEO on Stage featuring Cinnafilm, Inc., Perceptive Pixel and Total Immersion (Thursday, Sept 23, 16:00)**

Hill, Chris

Principle Research Scientist (M.I.T.)

Chris Hill is a computational scientist from MIT who specializes in modeling planetary fluid dynamics. With collaborators he has been developing fluid models of atmosphere and ocean processes for 20 years. He is a lead developer of the open-source M.I.T General Circulation Model (http://mitgcm.org) and has been exploring applications of accelerators for several years. With colleagues, he is developing a GPU oriented accelerator library for geosciences.

- **Session(s): 2167 - Designing a Geoscience Accelerator Library Accessible from High Level Languages (Wednesday, Sept 22, 17:00)**

Hoang-Trong, Tuan

PhD Student (George Mason University)

Tuan got his B.Eng. in Computer Science and Engineering from HoChiMinh City University in Vietnam in 2005. His M.Eng at Chonnam National University (South Korea) was in Computer Engineering; where he conducted research in artificial neural network, protein-spot maching in 2-dimensional gel electrophoresis from 2006-2008. Since 2008, he's a PhD student at George Mason University, Department of Bioinformatics and Computational Biology. His current research interests are in calcium signalling, building cardiac cell model using high-performance computing with GPU technology.

- **Session(s): 2172 - Unveiling Cellular & Molecular Events of Cardiac Arrhythmias (Tuesday, Sept 21, 11:00)**

Hoberock, Jared

Research Scientist (NVIDIA)

Jared Hoberock joined NVIDIA Research in October 2008. His current research interests include high performance ray tracing and parallel programming models. Jared has a contributed to both OptiX, NVIDIA's high performance ray tracing API, and Thrust, an open source library of high-level parallel primitives. Jared received a bachelor's degree in computer engineering from the University of Missouri at Columbia and a Ph.D in computer science from the University of Illinois at Urbana-Champaign. Jared is a two-time recipient of the NVIDIA Graduate Research Fellowship.

- **Session(s): 2220 - Thrust by Example: Advanced Features and Techniques (Thursday, Sept 23, 14:00)**

Hoeg, Steve

(Adobe)

- **Session(s): 2224 - GPU Acceleration in Adobe Creative Tools (Tuesday, Sept 21, 15:00)**

Huang, Jen-Hsun

CEO & President (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.

- **Session(s): 4006 - Fireside Chat with Jen-Hsun Huang - Co-founder & CEO, NVIDIA (Thursday, Sept 23, 09:00)**

Hummel, Michael

Managing Director (empulse GmbH)

Michael studied electronics and computer science at Cranfield University, University of Hertforshire and FHT-Esslingen. He received several awards for outstanding achievements. In 1994 he joined Accenture and worked as a process and technology consultant. As youngest manager of Germany he left Accenture in 1999 and continued his career as project and program manager for complex business and technology projects. In 2007 he founded "empulse" togehter with a former colleague as a professional service and software development company.

- **Session(s): 4004 - Emerging Companies: CEO on Stage featuring Cooliris, empulse GmbH, and Playcast Media Systems (Wednesday, Sept 22, 16:00))**

Humphrey, John

Senior Engineer (EM Photonics, Inc)

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 an FDTD solver based on OpenGL methods. Since then, he has worked on accelerated algorithms in a variety of fields, including linear algebra solvers and computational fluid dynamics engines.

- **Session(s): 2153 - CULA - A Hybrid GPU Linear Algebra Package (Thursday, Sept 23, 15:00)**
- **2154 - The Impact of Data Movement on GPU Performance (Wednesday, Sept 22, 16:00)**

Hwu, Wen-mei
Professor (University of Illinois, Urbana-Champaign)

Wen-mei W. Hwu is a Professor of ECE at the University of Illinois at Urbana-Champaign. He received the ACM Maurice Wilkes Award, the ACM Grace Murray Hopper Award, and the ISCA Most Influential Paper Award. He is a fellow of IEEE and ACM and leads the GSRC Concurrent Systems Theme. He directs the UIUC CUDA Center of Excellence. Dr. Hwu also received his Ph.D. degree in Computer Science from UC Berkeley.

- **Session(s): 2264 - CUDA Centers of Excellence Super-Session III (Tuesday, Sept 21, 16:00)**
- **2249 - New Programming Tools GPU Computing (Wednesday, Sept 22, 10:00)**

Iribe, Brendan
President (Scaleform)

Brendan Iribe co-founded Scaleform and established the company as the #1 video game user interface (UI) and video codec provider. Brendan pioneers all aspects of product research, development, and promotion at Scaleform. Under his leadership, Scaleform GFx has been adopted by most commercial 3D engines (UE3, CryEngine, Gamebryo) and licensed for use in over 600 titles in less than 4 years, including hit games from 19 of the top 20 worldwide video game publishers.

- **Session(s): 2241 - Standing Out: Implementing a Great Stereo UI (Thursday, Sept 23, 14:00))**

Iles, Andrew
Software Director (NVIDIA)

- **Session(s): 2225- Tools for Managing Clusters of NVIDIA GPUs (Tuesday, Sept 21, 17:00)**

Ismert, Ryan
Director of Engineering (Sportvision, Inc.)

Ryan has been building augmented reality systems for broadcast TV at Sportvision for 7 years. He currently leads a team focused on disrupting the current state of camera tracking and broadcast rendering by leveraging the power of multiple GPUs.

- **Session(s): 2123 - Enabling Augmented Reality with GPU Computing (Thursday, Sept 23, 15:00)**

Iyer, Kumar
Product Manager (NVIDIA)

Kumar Iyer is a Product Manager of Developer Tools at NVIDIA, where he works on the most advanced GPU development tools in the world. Prior to his work at NVIDIA, Kumar worked on PC and console games at Electronic Arts, and in research at the USC Institute for Creative Technologies. Kumar holds a B.S. in Computer Science from UCLA, and a MBA from the UCLA Anderson School of Management.

- **Session(s): 2245 - Parallel Nsight for Microsoft Visual Studio (Pre-Conference Tutorial) (Monday, Sept 20, 16:00)**
- **2149 - Overview of Parallel Nsight for Visual Studio (Thursday, Sept 23, 10:00)**
- **2149 - Overview of Parallel Nsight for Visual Studio (Tuesday, Sept 21, 11:30)**

Jamison, Andrew
CEO (Scalable Display Technologies)

Andrew Jamison is an experienced executive with four successful start-ups to date. Most recently Andrew was vice president of marketing for software maker Patron Systems. Prior to that, he was vice president of sales for Entelagent Software Corporation and ViewTech Corporation. He founded GroupNet, Inc., a PictureTel Corporation reseller. Jamison spent ten years with PictureTel Corporation where he was the fourth employee of the company and also a founding member of the European management team of PictureTel International LTD. During his years at PictureTel, revenue grew to over $200M and market value reached more than $1.0B. He received his undergraduate business education from Northeastern University and conducted post-graduate studies in finance at Fairfield University.

- **Session(s): 4007 - Emerging Companies: CEO on Stage featuring Aqumin, RTT, and Scalable Display Technologies (Thursday, Sept 23, 10:00)**

Jargstorff, Frank
Software Engineer (NVIDIA)

Frank Jargstorff is a software engineer leading NVIDIA's Performance Primitives effort (NPP). Frank received his degree in computer science in 1997 from the University of Tübingen, Germany.

- **Session(s): 2216 - CUDA Libraries Open House (Wednesday, Sept 22, 11:00)**

Jensen, Eric
Partner, Business Department Chair (Cooley LLP)

Eric C. Jensen is a business partner in the Cooley Palo Alto office. Mr. Jensen is head of the Firm's Business department and a member of the Management Committee. Mr. Jensen has been with the Firm since 1988 and a partner since 1994. Mr. Jensen practices securities and general corporate law, with an emphasis on the representation of emerging and public software, semiconductor, internet, and other information technology companies. He also has extensive experience representing venture capital funds and underwriters. He has counseled clients in the areas of corporate formations, venture financings, public offerings of equity and debt, mergers and acquisitions, joint venture, licensing and related strategic transactions, employee incentive matters and SEC reporting and compliance. Mr. Jensen has been included as one of The Best Lawyers in America in 2006 - 2011 and named as one of Northern California's "Super Lawyers" in 2007 - 2010. Mr. Jensen has also been ranked as a leading lawyer in Investment Funds: Venture Capital in Chambers USA 2010 edition.

- **Session(s): 4009 - Emerging Companies Summit Panel: The "New Normal" For Building Emerging Companies Based On Disruptive Technologies (Thursday, Sept 23, 14:00)**

Jeong, Byungil
Visualization Scientist (TACC / UT-Austin)

Byungil Jeong is a visualization scientist with the Texas Advanced Computing Center at the University of Texas at Austin and a primary Scalable Adaptive Graphics Environment (SAGE) architect. His research interests include scalable parallel graphics architecture, collaborative remote visualization, large-scale data visualization, and high-resolution display systems. Jeong has a PhD in computer science from the University of Illinois at Chicago.

- **Session(s): 2144 - Large-Scale Visualization Using A GPU Cluster (Wednesday, Sept 22, 16:00)**

Jeong, Won-Ki
Research Scientist (Harvard University)

Dr. Jeong is a research scientist at the Harvard Center for Brain Science (CBS). His research interests include

image processing, scientific visualization, and GPGPU in the field of biomedical image analysis. He received a Ph.D. Degree in Computer Science from the University of Utah in 2008, and was a member of the Scientific Computing and Imaging (SCI) institute at Utah. He received a NVIDIA Fellowship in 2007. He is currently a professional member of ACM.

- **Session(s): 2139 - Interactive Histology of Large-Scale Biomedical Image Stacks (Wednesday, Sept 22, 14:00)**

Juckeland, Guido
Senior System Engineer (HPC), Leader Hardware Accelerator Group (TU Dresden - ZIH)

Guido Juckeland received his M.Sc. from TU Dresden in Information System Technology. He is responsible for the operation and design of HPC resources of TU Dresden. Currently, he is working on a Ph.D. thesis "Performance Analysis for Hardware Accelerators".

- **Session(s): 2090 - Developing Highly Scalable Particle-Mesh Codes for GPUs: A Generic Approach (Tuesday, Sept 21, 15:00)**
- **2089 - Analyzing CUDA Accelerated Application Performance at 20 PFLOP/s (Wednesday, Sept 22, 17:00)**

Kapasi, Ujval
CUDA Platform SW (NVIDIA)

- **Session(s): 2216 - CUDA Libraries Open House (Wednesday, Sept 22, 11:00)**

Kaplan, Michael
Vice President of Strategic Development (mental images/ NVIDIA)

30 years of experience and contributions in the 3D graphics industry; masters degree from the Cornell University Program of Computer Graphics in 1980; inventor of spatial-subdivision raytracing, the first object-oriented 3D scene graph, and other technologie; Michael Kaplan is currently the VP Strategic Development at mental images/NVIDIA. He is responsible for the iray CUDA accelerated photorealistic rendering technology project.

- **Session(s): 2013 - iray - GPUs and the Photorealistic Rendering Revolution (Tuesday, Sept 21, 14:00)**

Kerr, Andrew
PhD Student (Georgia Institute of Technology)

- **Session(s): 2210 - GPU-Ocelot: An Open Source Debugging and Compilation Framework for CUDA (Thursday, Sept 23, 14:00)**

Khadtare, Mahesh
Member, Technical Staff (Computational Research Laboratories, Pune, INDIA.)

- **Session(s): 2298 - Accelerated Image Quality Assessment using Structural Similarity (Thursday, Sept 23, 11:30)**

Kilgard, Mark
Principal System Software Engineer (NVIDIA)

- **Session(s): 2127 - OpenGL (Pre-Conference Tutorial) (Monday, Sept 20, 16:00)**

Kim, Hyesoon
Assistant Professor (Georgia Tech)

Hyesoon Kim is an Assistant professor in the School of Computer Science at the Georgia Institute of Technology. Her research interests include high-performance energy-efficient heterogeneous architectures, programmer-compiler-microarchitecture interaction especially for CPU/GPU systems. She received an MS and a Ph.D in computer engineering at The University of Texas at Austin.

- **Session(s): 2164 - Analytical Performance Models to Improve the Efficiency of GPU Computing (Wednesday, Sept 22, 14:00)**

Kirsanov, Danil
Scientist (ANSYS)

- **Session(s): 2066 - Accelerating System Level Signal Integrity Simulation (Thursday, Sept 23, 16:30)**

Kloeckner, Andreas
Courant Instructor (Courant Institute, NYU)

Andreas recently completed his PhD in applied mathematics with Jan Hesthaven at Brown University, working on various aspects of high-order finite element methods. In September 2010, he will be joining the Courant Institute of Mathematical Sciences at New York University to work on problems in computational electromagnetics with Leslie Greengard. He is the main author of the PyCUDA and PyOpenCL GPU computation packages.

- **Session(s): 2041 - PyCUDA: Even Simpler GPU Programming with Python (Wednesday, Sept 22, 14:00)**

Kohlmeyer, Axel
Associate Director (Institute for Computational Molecular Science, Temple University)

Axel Kohlmeyer is the Associate Director of the Institute for Computational Molecular Science and the Associate Professor of Chemistry and Computer Science at Temple University. He earned his PhD in Theoretical Chemistry at the University of Ulm. His main interests are making creative use of scientific computing tools to advance understanding of processes at the molecular and atomistic level, and making these tools more capable and accessible.

- **Session(s): 2168 - Interactive Molecular Dynamics for Nanomechanical and Nanochemical Experiments (Wednesday, Sept 22, 10:00)**

Korf, Dave
Scalable Computing & Infrastructure Organization, Marketing (HP)

Mr Korf has over 15 years of engineering expierence with the last 20 plus years in various senior marekting, product management and partner management positions. Accelerators, Parnters and competitive analysis are currently some of his focus areas.

- **Session(s): 2233 - Solving Your GPU Computing Needs (Sponsored by HP) (Tuesday, Sept 21, 14:00)**

Kramer, Thomas
VP Product Management (MainConcept)

Thomas is VP Product Management at MainConcept. He has more than 8 years of experience in Video and Audio Technologies ranging from actual Codec Programming and Pre-Sales activities to Codec Product Management. His current responsibility is for SDK business as well end-user Applications. In his previous businesses he has worked for archiving, bank and webservice oriented clients. He has skills in education and training for various PC based tools and applications. Thomas has studied Business Informatics in Cologne, Germany.

- **Session(s): 2048 - H.264/AVC Video Encoding with CUDA and OpenCL (Thursday, Sept 23, 09:00)**

Krishna, Murali
Junior Research Associate (Infosys Technologies Limited)

Murali Krishna graduated with a Masters degree from IIIT Bangalore and has been working as a Junior Research Associate with SETLabs, the R&D arm of Infosys Technologies. Murali has worked in areas of Parallelism in workflows, automatic data parallelism extraction in legacy applications. Now, Marali's research is focused on porting data parallel applications to GPUs.

- **Session(s): 2120 - High Performance Complex Event Processing on GPGPU (Wednesday, Sept 22, 14:00)**

Krishnamurthy, Adarsh
Student (University of California Berkeley)

Adarsh Krishnamurthy is a PhD Candidate in the department of Mechanical Engineering at U.C. Berkeley. His research interests include Computer Aided Design (CAD), solid modeling, GPU algorithms, computational geometry, and ultrasonic non-destructive testing. He received his Bachelors and Masters in Mechanical Engineering from Indian Institute of Technology, Madras.

- **Session(s): 2171 - Parallel Algorithms for Interactive Mechanical CAD (Thursday, Sept 23, 11:00)**

Kunz, Holger
Director, Workstation Software Development (NVIDIA)

Holger Kunz work has been in the area of professional visualization covering real-time visualization, photo-realistic visualization, multi-GPU rendering, shading languages and computer vision. He received his Dipl. Inform. at the University of Erlangen Nürnberg.

- **Session(s): 2024 - NVIDIA Acceleration Engines Overview (Pre-Conference Tutorial) (Monday, Sept 20, 13:00)**

Lanza, Drew
Partner (Morgenthaler)

Drew, based in Menlo Park, CA, joined Morgenthaler in 2000 and became a Partner in 2001. Drew focuses on cleantech, semiconductors and systems. He is currently a Director of Cortina Systems, Overture Networks, OmniPV, Unity Semiconductor, Autonet Mobile, SiPort, ZeroG, and R2 Semiconductor. Drew spent 15 years in senior operating positions in the telecommunications industry starting companies in both the components and the systems sectors of that industry. Drew was a founder and VP of Engineering at E/O Networks where he helped to design and produce a long reach rural fiber optic telephony system. Drew started his optical telecommunications career in 1986 at Raynet, a pioneering company in the development of fiber to the home technologies. Drew's many roles at Raynet included VP of Marketing and VP of International Development. Drew was the founding CEO of Lightwave Microsystems, a leader in the design and manufacture of high volume optical integrated circuits. Drew graduated magna cum laude from Harvard with an MBA in 1987. He received his BSEE & MSEE degrees from Stanford in 1979.

- **Session(s): 4001 - Emerging Companies: CEO on Stage featuring Elemental Technologies, Inc., Geomerics, and Milabra (Wednesday, Sept 22, 11:00)**
- **4002 - Emerging Companies: CEO on Stage featuring Allegorithmic SAS, Bunkspeed, and miGenius (Wednesday, Sept 22, 14:00)**

Lauer, Tobias
Researcher (University of Freiburg)

Tobias Lauer received his PhD in computer science from the University of Freiburg in 2007. His current research focus is on efficient algorithms and data structures for multidimensional aggregation for Online Analytical Processing (OLAP). His work is funded by the German Research Foundation (DFG) and is carried out in collaboration with Jedox Business Intelligence, a company specializing in Business Intelligence software.

- **Session(s): 2237 - Accelerating Business Intelligence Applications with Fast Multidimensional Aggregation (Wednesday, Sept 22, 15:00)**

Lazebnik, Roee
Director of Product Development (Siemens Healthcare)

Dr. Roee Lazebnik's professional experience and training consist of clinical radiology, biomedical engineering, information technology, software development, and healthcare business. He is the author of numerous published manuscripts, textbook chapters, and conference proceedings on many topics in medical imaging.

- **Session(s): 2169 - Real-time Volumetric Medical Ultrasound Applications for GPU Computing (Wednesday, Sept 22, 10:00)**

Le Grand, Scott
Principal Engineer (NVIDIA)

Scott is a principal engineer on the CUDA platform team at NVIDIA with a B.S. in biology from Siena College and a Ph.D. in biochemistry from the Pennsylvania State University. Scott developed Genesis, the first molecular modeling system for home computers, Folderol, the first distributed computing project targeting the protein folding problem, and BattleSphere, a 3D space shooter for the Atari Jaguar. More recently, he ported the Folding@Home and AMBER molecular modeling codebases to CUDA.

- **Session(s): 2218 - Redesigning Molecular Dynamics for GPUs and GPU Clusters (Wednesday, Sept 22, 15:00)**

Leback, Brent
Engineering Manager (The Portland Group)

Brent Leback is an Engineering Manager for PGI. He has worked in various positions over the last 26 years in HPC customer support, math library development, applications engineering and consulting at QTC, Axian, PGI and STMicroelectronics.

- **Session(s): 2143 - CUDA Fortran Programming for NVIDIA GPUs (Wednesday, Sept 22, 15:30)**

Lecomber, David
CTO (Allinea Software)

David Lecomber is one of the founders of Allinea and leads the research and development team behind Allinea DDT, the world's most scalable parallel debugger.

- **Session(s): 2039 - GPU Debugging with Allinea DDT (Wednesday, Sept 22, 11:00)**

Lee, HyoukJoong
PhD Student (Stanford University)

HyoukJoong Lee is a PhD student in electrical engineering at Stanford University. His research interests include parallel systems architecture and programming models. He has a BS degree from Seoul National University.

- **Session(s): 2177 - Simplifying Parallel Programming with Domain Specific Languages (Wednesday, Sept 22, 11:00)**

Lee, John
(Appro)

John K. Lee joined Appro in 2001, and is responsible

for leading Appro's hardware product development engineering team. In addition, Mr. Lee leads the company's Project Management team that is responsible for deploying Appro's complex cluster solutions. He has served as the Program Executive for some of Appro's most important cluster projects such as 2006 Peloton Project as well as 2007 TLCC Cluster Project. Prior to his role at Appro, Mr. Lee served in both Sales and Service Management capacities at multiple storage and telecom companies.

- **Session(s): 2270 - Appro's GPU Computing Solutions (Tuesday, Sept 21, 15:00)**

Lefebvre, Matthieu
PhD Student (ONERA)

Matthieu Lefebvre is a PhD student at ONERA, the french aerospace lab, and at Université Paris 13. He is working on accelerating CFD simulations on GPU.

- **Session(s): 2045 - Roe-Pike Scheme for 2D Euler Equations (Wednesday, Sept 22, 14:00)**

Lever, Ben
Senior Research Engineer (NICTA)

Ben Lever is a senior research engineer at NICTA currently developing new methodologies and frameworks for describing computer vision algorithms that can target heterogeneous, highly-parallel platforms. Prior to NICTA, Ben was a hardware design engineer at Canon Research before joining Synopsys as a software engineer for developing real-time simulation models of embedded processors.

- **Session(s): 2173 - Enabling Large-Scale CCTV Face Recognition (Thursday, Sept 23, 11:00)**

Lewin, Dan'l
Corporate Vice President, Strategic and Emerging Business Development (Microsoft)

Dan'l Lewin is responsible for leading Microsoft's global engagement with startups and venture capitalists. In addition, Lewin has executive, site, and citizenship responsibility for the company's operations in the Silicon Valley, based in Mountain View, California, which currently employ 2,500 people and supports business relationships with industry partners in Silicon Valley. Lewin's business development teams focus on supporting the software startup and entrepreneur ecosystem developing on the Microsoft platform while helping foster and grow local software economies worldwide. Through the Microsoft BizSpark Program, and the Microsoft Innovation Center Program, the groups help accelerate startup success in more than 100 countries.

- **Session(s): 4001 - Emerging Companies: CEO on Stage featuring Elemental Technologies, Inc., Geomerics, and Milabra (Wednesday, Sept 22, 11:00)**
- **4002 - Emerging Companies: CEO on Stage featuring Allegorithmic SAS, Bunkspeed, and miGenius (Wednesday, Sept 22, 14:00)**

Li, Hongjian
Graduate Student (The Chinese University of Hong Kong)

Hongjian Li is currently working on a M.Phil. degree under the supervisions of Prof. Kwong Sak Leung and Prof. Man Hon Wong. His major research interests are GPU applications in bioinformatics, particularly computer-aided drug design by means of high throughput in silico virtual screening via structure-based ligand-protein docking. He is developing new software that utilizes the computational horsepower of graphics processors with a purpose of accelerating the pipeline of drug discovery.

- **Session(s): 2088 - Nucleotide String Matching Using CUDA-Accelerated Agrep (Thursday, Sept 23, 16:00)**

Lichtenbelt, Barthold
Sr. OpenGL Manager (NVIDIA)

Barthold Lichtenbelt is a Sr. Manager of the OpenGL core driver team at NVIDIA. He is also the Chair of the OpenGL ARB Khronos Working Group.

- **Session(s): 2127 - OpenGL (Pre-Conference Tutorial) (Monday, Sept 20, 16:00)**

Lin, Chun-Yuan
Assistant Professor (Department of CSIE, Chang Gung University)

Chun-Yuan Lin joined the Department of Computer Science and Information Engineering at Chang Gung University as an assistant professor. His research interests are in the areas of parallel and distributed computing, parallel algorithms, algorithm analysis, information retrieve, proteomics, and bioinformatics.

- **Session(s): 2105 - CUDA-FRESCO: An Efficient Algorithm for Mapping Short Reads (Thursday, Sept 23, 15:00)**

Linderman, Michael
Engineering Research Associate (Stanford University)

Michael Linderman is an Engineering Research Associate in the Computer Systems Laboratory at Stanford University. His research focuses on using graphics processing units (GPUs) and other heterogeneous computer systems to accelerate computational systems biology and other data- and compute-intensive applications. Michael earned a Ph.D. and MS in Electrical Engineering from Stanford University in 2009 and 2006 respectively, and B.S. from Harvey Mudd College in 2003.

- **Session(s): 2030 - High-Throughput Cell Signaling Network Learning with GPUs (Thursday, Sept 23, 09:00)**

Loddoch, Alex
Research Scientist (Chevron)

Alex received an MSc in Physics in 2001 and a PhD in Geophysics in 2007 from University in Muenster, Germany, working on topics including geophysical fluid dynamics, parallel computing and data compression. In 2007 he joined Chevron as a Research Scientist, working in the area of High Performance Computing and particularly GPU computing.

- **Session(s): 2174 - Reverse Time Migration on GPUs (Wednesday, Sept 22, 15:00)**

Löhner, Rainald
Professor (George Mason University)

Rainald Lohner is the head of the CFD center at the department of computational and data sciences of George Mason University in Fairfax, VA, in the outskirts of Washington, D.C. He received a MSc in Mechanical Engineering from the Technische Universitaet Braunschweig, Germany, as well as a PhD and DSc in Civil Engineering from the University College of Swansea, Wales, where he studied under Profs. Ken Morgan and Olgierd Zienkiewicz. His areas of interest include numerical methods, solvers, grid generation, parallel computing, visualization, pre-processing, fluid-structure interaction as well as shape and process optimization. His codes and methods have been applied in many fields, including aerodynamics or airplanes, cars and trains, hydrodynamics of ships, submarines and UAVs, shock-structure interaction, dispersion analysis in urban areas and haemodynamics of vascular

meet the TITANS

Ciara TITAN + NVIDIA technology

With the TITAN Series, get GPU computing with greater flexibility and tailor your cluster to your specific HPC workload. Ciara GPU TITAN systems offer a rich array of data center-ready solutions including the Tesla™ S1070 or the C2050, from rack-mount systems to blades and with the networking, storage and redundancy features you want. Our TITAN Clusters contribute to the GPU world through our scalable solutions (1U, 2U, 4U) with up to 80 TeraFlops in a rack at your disposal to solve today's most demanding challenges.

For more information on our solutions, visit us on booth xxx during the show or online at www.ciaratech.com.

CIARA
Speed | Performance | Passion for Innovation

GPU TECHNOLOGY CONFERENCE
Sept. 20-23 2010 - San Jose, CA

NVIDIA PARTNERFORCE TESLA PREFERRED PROVIDER

W: ciara-tech.com - T: 1-866-789-7225 - E: nvidia@ciara-tech.com

diseases. He is the author of more than 600 articles covering the fields enumerated above, as well as a textbook on Applied CFD Techniques.

Session(s): 2005 - Porting Large-Scale Legacy Fortran Codes (Wednesday, Sept 22, 17:00)

Loop, Charles
Senior Researcher (Microsoft Research)

Charles Loop is a Senior Researcher in the Microsoft Research Graphics Group. He has worked extensively in the areas of curve and surface modeling and rendering, including work on n-sided patches, smooth patch complexes, as well as GPU algorithms for rendering vector art. Charles is best known for the triangle mesh subdivision algorithm that bears his name. He is currently working on data parallel algorithms for REYES style rendering and accelerating raytracing of surface primatives.

Session(s): 2129 - Hardware Subdivision and Tessellation of Catmull-Clark Surfaces (Tuesday, Sept 21, 16:00)

Lorach, Tristan
Computergraphics Engineer (NVIDIA)

Tristan Lorach has worked on many realtime interactive events all over the world. Tristan is now working at NVIDIA in the developer technical relations department (a.k.a Devtech), participating on a variety of projects in relation with NVIDIA partners. At the same time, he is also contributing to R&D writing demos for new GPU Chips.

Session(s): 2056 - Next-Generation Rendering with CgFX (Tuesday, Sept 21, 16:00)

Ltaief, Hatem
Sr. Research Associate (University of Tennessee)

Hatem Ltaief received the MSc degree from the school of engineering at the University of Claude Bernard Lyon I, France, the MSc in applied mathematics at the University of Houston and the PhD degree in computer science from the University of Houston. He is a research associate in the Innovative Computing Laboratory in the Department of Electrical Engineering and Computer Science at the University of Tennessee, Knoxville. His research interests include parallel algorithms, specifically in the area of numerical linear algebra, and also parallel programming models and performance optimization for parallel architectures spanning distributed and shared memory systems, as well as next generation multi-core and many-core processors.

Session(s): 2138 - Faster, Cheaper, Better – Hybridization of Linear Algebra for GPUs (Thursday, Sept 23, 09:00)

Lu, Peter
Post-Doctoral Research Fellow (Harvard University)

Peter J. Lu is a post-doctoral research fellow in physics at Harvard University; he focuses on the physics of attractive colloids, integrating high-performance imaging analysis. He conducts experiments aboard the International Space Station, and has also published his discoveries of modern quasicrystal geometry in medieval Islamic architectural tilings; the first precision compound machines, from ancient China; the first use of diamond, in prehistoric China; and the first natural quasicrystalline mineral. For more infomation, see http://www.peterlu.org

Session(s): 2242 - Swarming Bacteria and Diffusing Particles: High-Throughput Analysis of Microscopic 3D Motion (Wednesday, Sept 22, 17:00)

Lucas, Bob
Computational Sciences Division Director (University of Southern California)

Bob Lucas is the Computational Sciences Division Director at the University of Southern California's Information Sciences Institute. He has been developing parallel linear solvers since 1985.

Session(s): 2240 - Accelerating LS-DYNA with MPI, OpenMP, and CUDA (Thursday, Sept 23, 14:30)

Lumsdaine, Andrew
Professor (Indiana University)

Andrew Lumsdaine is a professor in the School of Informatics & Computing at Indiana University, and an Associate Director of the Digital Science Center and Director of the Open Systems Lab at the Pervasive Technology Institute. Lumsdaine received his Ph.D. from MIT in 1992, and from 1992 through 2001, he was a faculty member in the Department of Computer Science and Engineering at the University of Notre Dame. His research interests include computational science and engineering, parallel and distributed computing, software engineering, generic programming, mathematical software, and numerical analysis. Lumsdaine is a member of ACM, IEEE, and SIAM, as well as the MPI Forum, the BLAS technical forum and the ISO C++ standards committee. In 1995, he received the Career Development Award from the National Science Foundation.

Session(s): 2093 - Computational Photography: Real-Time Plenoptic Rendering (Wednesday, Sept 22, 16:00)

Lunn, Philip
CEO (Bunkspeed)

Philip Lunn is the visionary founder of Bunkspeed. He has brings his passion for computer graphics to democratize the creation of photographic quality 3D imagery and animation to every Bunkspeed product. Simplicity without compromise enables explosive growth of new users and new untapped markets. Mr. Lunn has over 20 years of technology and entrepreneurial experience and holds a Bachelor of Science degree in Mechanical Engineering from the University of Arizona.

Session(s): 4002 - Emerging Companies: CEO on Stage featuring Allegorithmic SAS, Bunkspeed, and miGenius (Wednesday, Sept 22, 14:00)

Mallick, Dr. Sudeep
Principle Research Scientist (Infosys)

Session(s): 2120 - High Performance Complex Event Processing on GPGPU (Wednesday, Sept 22, 14:00)

Malcolm, James
VP of Engineering (AccelerEyes)

James Malcolm is VP of Engineering at AccelerEyes and a co-founder. He holds degrees in Mathematics (BS), Computer Science (BS, MS), and Electrical Engineering (MS) from Georgia Institute of Technology.

Session(s): 2271 - Compose CUDA Masterpieces! Write better, Leverage More (Thursday, Sept 23, 16:00)

Marbach, Jonathan
Director of Software Architecture and Engineering (TerraSpark Geosciences, LLC)

Jon is currently working as Director of Software Architecture and Engineering at TerraSpark Geosciences, makers of the breakthrough 3D Seismic Interpretation package Insight Earth. He specializes in OpenGL and Virtual Reality, and received his PhD in Computer Science from the University of Colorado at Boulder in 2009. His PhD dissertaion investigates techniques for supporting stereoscopic 3D for simulataneous participants in a virtual reality environment.

- **Session(s): 2107 - Accelerating Stereographic and Multi-View Images Using Layered Rendering (Thursday, Sept 23, 15:00)**

Masaie, Issei
Chief GPU Engineer (Prometech Software, Inc.)

He received his masters degree in Quantum Engineering and System Science from the University of Tokyo. He joined Prometech Software, Inc. in 2007. He currently heads the development of the GPU accelerated version of Particleworks.

- **Session(s): 2106 - Particleworks: Particle-based CAE Software on Multi-GPU (Thursday, Sept 23, 11:30)**

Mason, Chris
Product Manager (Acceleware)

- **Session(s): 2208 - Acceleration of SIMULIA's Abaqus Solver on NVIDIA GPUs (Thursday, Sept 23, 15:30)**

Mastrobuono Battisti, Alessandra
PhD Student (Sapienza- University of Rome)

Alessandra Mastrobuono Battisti is a second year PhD student in Astronomy at the University of Rome La Sapienza. She received her Bachelor Degree in Physics and Astrophysics from the University La Sapienza in 2006. She also has a Masters Degree in Astronomy Astrophysics from the same University.
Her PhD program concerns the study of the dynamical evolution of N-Body systems.
At this scope, she realized a high performing code, NBSymple that runs on composite architecture.

- **Session(s): 2000 - Gravitational N-body Simulations: How Massive Black Holes Interact with Stellar Systems (Wednesday, Sept 22, 14:00)**

Matsuoka, Satoshi
Professor (Tokyo Institute of Technology)

Satoshi Matsuoka received his Ph. D. from the University of Tokyo in 1993. He became a full Professor at the Global Scientific Information and Computing Center (GSIC) of Tokyo Institute of Technology (Tokyo Tech / Titech) in April 2001, leading the Research Infrastructure Division Solving Environment Group of the Titech campus. He has won several awards including the Sakai award for research excellence from the Information Processing Society of Japan in 1999, and recently received the JSPS Prize from the Japan Society for Promotion of Science in 2006 from his Royal Highness Prince Akishinomiya.

- **Session(s): 2265 - CUDA Centers of Excellence Super-Session IV (Tuesday, Sept 21, 17:00)**
- **2280 - TSUBAME2.0 Experience (Wednesday, Sept 22, 10:00)**

Maurer, Lance
CEO (Cinnafilm, Inc.)

Lance Maurer is the founder, president and CEO of Cinnafilm, Inc. – an American engineering company dedicated to global leadership in image optimization using innovative and affordable parallel-processing methods. Prior to launching Cinnafilm, Maurer spent ten years, primarily with Goodrich Aerospace, designing, analyzing and testing technology used in the world's most advanced spacecraft and launch vehicles; clients include: NASA, Boeing, Ball, Kodak and Lockheed-Martin. He would eventually become an expert in thermal, structural and materials design for extreme environmental applications, honing regimented engineering philosophy in the true "failure-is-not-an-option" defense industry. During his tenure as an aerospace engineer he also pursued his life passion of storytelling and filmmaking on nights and weekends, finding the time to write, direct and produce numerous independent films (one feature film, two shorts and a music video). It was during these parallel journeys that his "a-ha" moment occurred and Lance's entrepreneurial spirit ignited a new path for him to follow in the field of image processing.

- **Session(s): 4011 - Emerging Companies: CEO on Stage featuring Cinnafilm, Perceptive Pixel, and Total Immersion (Thursday, Sept 23, 16:00)**

McAllister, Dave
Development Lead, OptiX Development (NVIDIA)

David is the lead developer for the next release of OptiX. Prior to joining the OptiX group two years ago David was a GPU architect at Nvidia for eight years, working on the design of all GPU families from GeForce 3 through Fermi. David holds a Ph.D. in Computer Science from the University of North Carolina at Chapel Hill.

- **Session(s): 2261 - Introduction to GPU Ray Tracing with NVIDIA OptiX (Pre-Conference Tutorial) (Monday, Sept 20, 14:30)**

McMains, Sara
Associate Professor (University of California Berkeley)

Sara McMains is an Associate Professor of Mechanical Engineering at UC Berkeley. Her research interests include geometric design for manufacturing feedback, solid modeling, CAD/CAM, GPU algorithms, computer aided process planning, layered manufacturing, computer graphics, visualization, virtual prototyping, and virtual reality. She received her A.B. from Harvard and her M.S. and Ph.D. from UC Berkeley, all in Computer Science. She is the recipient of Best Paper Awards from Usenix (1995) and ASME DETC (2000), a Best Poster and a Best Paper Award from the ACM Solid and Physical Modeling Symposium (2007, 2008 -- 2nd place), and the NSF CAREER Award (2005).

- **Session(s): 2171 - Parallel Algorithms for Interactive Mechanical CAD (Thursday, Sept 23, 11:00)**

Meister, Benoit
Senior Engineer (Reservoir Labs)

Benoit Meister received his BSc in Physics and his PhD in computer science from Strasbourg University, in automatic program parallelization and optimization using a polyhedral model of computation loops. After a post-doc in Verimag Grenoble, Benoit has joined Reservoir Labs where he contributed to the development and management of R-Stream, an advanced auto-parallelizing compiler based on extensions of the polyhedral optimization techniques to date. At the moment, R-Stream successfully targets 6 radically different architectures and parallelizes a broad range of applications.

- **Session(s): 2202 - A Programming Model and Tool for Automatic High-Performance C to CUDA Mapping (Thursday, Sept 23, 09:00)**

Menon, Shashi
R&D Manager (Schlumberger)

- **Session(s): 2141 - Moving the Frontier of Oil and Gas Exploration and Production with GPUs (Wednesday, Sept 22, 10:00)**

Meredith, Jeremy
Computer Scientist (Oak Ridge National Laboratory)

- **Session(s): 2089 - Analyzing CUDA Accelerated Application Performance at 20 PFLOP/s (Wednesday, Sept 22, 17:00)**

Merrill, Duane
Ph.D Candidate (University of Virginia)

Duane Merrill is a Ph. D. candidate at the University of Virginia, Department of Computer Science. His advisor is Professor Andrew Grimshaw. Before graduate school, he was a software developer for Avaki Corporation, specializing in grid computing middleware. His current research interests lay in parallel and high-performance computing, specifically in regard to programming models and algorithmic primitives for GPGPU, stream, and many-core architectures. Much of his prior academic work has involved concurrent systems in one form or another, including grid and distributed computing; virtual machines and hypervisor technologies; operating systems and meta-systems; and security architecture and protocols.

- **Session(s): 2296 - CUDA Optimization for Ninjas: A Case Study of High-Performance Sorting (Wednesday, Sept 22, 15:00)**

Micikevicius, Paulius
Developer Technology Engineer (NVIDIA)

Paulius Micikevicius is a Developer Technology Engineer at NVIDIA with a focus on parallel computing and performance. 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 and a B.S. in Computer Science from Midwestern State University.

- **Session(s): 2012 - Analysis-Driven Performance Optimization (Thursday, Sept 23, 15:00)**
- **2011 - Fundamental Performance Optimizations for GPUs (Wednesday, Sept 22, 17:00)**

Miller, Phillip
Director, Workstation Software Product Management (NVIDIA)

Phillip Miller is an accomplished software product manager with 16 years experience guiding industry leading solutions from companies such as Autodesk and Adobe. He is also a registered architect, bringing real-world experience in using tools and directing projects to software creation. At NVIDIA, Mr. Miller directs much of the professional middleware produced to enable software developers to quickly leverage the potential of the GPU within the applications they produce.

- **Session(s): 2024 - NVIDIA Acceleration Engines Overview (Pre-Conference Tutorial) (Monday, Sept 20, 13:00)**
- **2261 - Introduction to GPU Ray Tracing with NVIDIA OptiX (Pre-Conference Tutorial) (Monday, Sept 20, 14:30)**

Mitchell, Kenny
Research Lead (Black Rock Studio)

Kenny is the research lead at Black Rock, Disney Interactive Studios. His Ph.D. introduced the use of real-time 3D for information visualization on consumer hardware, including a novel recursive perspective projection technique. Over the past ten years, he has shipped games using high-end graphics technologies including voxels, PN patches, displacement mapping, and clipmaps. In between shipping games for the Harry Potter franchise and now racing games with Split/Second, he is also involved in developing new intellectual properties.

- **Session(s): 2047 - Bridging Ray and Raster Processing on GPUs (Tuesday, Sept 21, 11:00)**

Morrison, Michael
(NVIDIA)

- **Session(s): 2308 - Building Cutting-Edge Realtime 3D Applications with NVIDIA SceniX (Wednesday, Sept 22, 10:00)**

Mooney, Al (Adobe)

- **Session(s): 2224 - GPU Acceleration in Adobe Creative Tools (Tuesday, Sept 21, 15:00)**

Morton, Scott
Geophysical Advisor (Hess Corporation)

Scott Allyn Morton received his B.A. in physics & math from Gustavus Adolphus College and his Ph.D. in astrophysics from University of Illinois at Urbana-Champaign. He has 25 years of experience in computational and theoretical physics distributed between academia, the computer industry and the petroleum industry. Scott has worked at NCSA (National Center for Supercomputing Applications), Shell, Thinking Machines, Cray Research and SGI (Silicon Graphics Inc). Scott manages the geophysical technology development group for Hess Corporation and is responsible for monitoring, testing, adopting and developing new geophysical and computational technologies.

- **Session(s): 2059 - Industrial Seismic Imaging on GPUs (Wednesday, Sept 22, 11:00)**

Moulik, Supratik
(University of Pennsylvania)

Dr. Moulik is a cardiovascular imaging fellow at the University of Pennsylvania. Combining an engineering degree from Carnegie Mellon University with 10 year of graduate medical education, Dr. Moulik is a unique blend of physician and programmer. The breadth of training has allowed him to develop GPU computing algorithms for the medical imaging community which are both intuitive and robust.

- **Session(s): 2036 - Algorithms for Automated Segmentation of Medical Imaging Studies Utilizing CUDA (Tuesday, Sept 21, 16:00)**

Mroue, Abdul
Post-Doc Fellow (CITA, Univ. Of Toronto)

Abdul Mroue is a PostDoc Fellow at the Canadian Institute for Theoretical Astrophysics, University of Toronto. He received his PhD degree in Physics from Cornell University in 2009. His research work focuses on solving general relativity numerically.

- **Session(s): 2108 - Binary Black Holes Simulations using CUDA (Wednesday, Sept 22, 16:00)**

Mrsic-Flogel, Janko
CTO (Mirriad)

Janko has over 15 years experience in Internet and Mobile Internet and was a pioneer of interactive television. He has achieved high-yielding returns to shareholders in technology company sales and is the founder of several successful IT and Telecoms ventures. Prior to joining MirriAd he was Managing Director of Dynamical Systems Research, Technical Director at Digital Mobility, and Director of Applied Technology at BBC Vecta, the BBC's Venturing arm. Janko is an acknowledged speaker, lecturer, author of scientific publications and inventor of patent applications in areas of Computing, Neural Networks, Telecommunications, Interactive Television and Electrical Engineering. Janko holds a PhD in Electrical Engineering from Imperial College.

- **Session(s): 4003 - Emerging Companies Summit Panel: GPU for Computer Vision (Wednesday, Sept 22, 15:00)**

Mueller, Frank
Associate Professor (North Carolina State University)

- **Session(s): 2272 - GStream: A General-Purpose Data Streaming Framework on GPUs (Thursday, Sept 23, 09:00)**
- **2026 - MatCloud: Accelerating Matrix Math GPU Operations with SaaS (Tuesday, Sept 21, 17:00)**

Murphy, K.C.
Vice President of Marketing (NextIO)

K.C. is a semiconductor veteran with more than 25 years of experience, and time spent with some of the industry's best known companies: AMD, Cadence and Broadcom. Prior to NextIO, K.C. was an Executive Officer of Broadcom and the VP/GM of Broadcom's RF and Advanced Mixed Signal Group. K.C. joined Broadcom through Broadcom's acquisition of Pivotal Technologies, where he was President and CEO. Before Pivotal, K.C. was Executive Vice President, Strategic Business Group and Corporate Strategy, at Cadence Design Systems. K.C. joined Cadence after many years at AMD where his last position was Vice President, Corporate Strategy and PC Systems.

- **Session(s): 2247 - Reconfiguring a Pool of GPUs on The Fly (Sponsored by NextIO) (Tuesday, Sept 21, 16:00)**

Narayanan, Babu
Lab Manager (GE Global Research)

Lab Manager, Medical Image Analysis Lab, GE Research, JFWTC, Bangalore

- **Session(s): 2094 - Nearly Instantaneous Reconstruction for MRIs (Tuesday, Sept 21, 14:00)**

Nash, Steve
Applied Engineer (NVIDIA)

Steve is an Applied Engineer in the Professional Solutions Group at Nvidia, focussing on scalable visualization solutions.

- **Session(s): 2010 - Implementing Stereoscopic 3D in Your Applications (Pre-Conference Tutorial) (Monday, Sept 20, 16:00)**
- **2071 - Large Scale Visualization Soup (Wednesday, Sept 22, 11:00)**

Naumov, Maxim
Software Engineer (NVIDIA)

Maxim Naumov's expertise is in the area of parallel numerical linear algebra. In particular, he has worked on parallel iterative linear systems and eigenvalue solvers. He received his Ph.D. in Computer Science (with specialization in Computational Science and Engineering) in 2009 and his B.Sc. in Computer Science and Mathematics in 2003, all from Purdue University – West Lafayette. He currently works in NVIDIA CUDA Platform team developing parallel numerical algorithms for Graphical Processing Units (GPUs). He has previously worked in the Intel Corporation Microprocessor Technology Lab and Computational Software Lab, and received a 2008-09 Intel Foundation Ph.D. Fellowship.

- **Session(s): 2070 - CUSPARSE Library: A Set of Basic Linear Algebra Subroutines for Sparse Matrices (Thursday, Sept 23, 11:00)**

Navratil, Paul
Visualization Scientist (Texas Advanced Computing Center)

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

- **Session(s): 2144 - Large-Scale Visualization Using A GPU Cluster (Wednesday, Sept 22, 16:00)**

Negrut, Dan
Assistant Professor (University of Wisconsin)

Dan received his Mechanical Engineering Ph.D. from the University of Iowa in 1998. He spent six years in the software industry. In 2004, he served as an Adjunct Professor in the Department of Mathematics at the University of Michigan. He spent 2005 as a Mathematics and Computer Science Visiting Scientist at Argonne National Laboratory. Dan joined the University of Wisconsin in 2006. His interests are in Computational Science and he leads the Simulation-Based Engineering Lab.

- **Session(s): 2231 - Driving on Mars, Redux: System Level Simulation of Dynamic Systems (Wednesday, Sept 22, 10:00)**

Nessim, Maurice
(Schlumberger)

- **Session(s): 2141 - Moving the Frontier of Oil and Gas Exploration and Production with GPUs (Wednesday, Sept 22, 10:00)**

Nichols, David
Research Director (Schlumberger)

Dave Nichols has worked in the seismic imaging industry for 25 years. Much of his work has focused on the intersection of new HPC technologies and new imaging algorithms. After working in seismic data processing and software development for 5 years, he returned to academia to study for a PhD at Stanford University. He has since worked for Schlumberger in a number of roles ranging from managing seismic research at the long term research labs to managing the process for introducing new technologies.

- **Session(s): 2142 - Complex Geophysical Imaging Algorithms Enabled by GPU technology (Wednesday, Sept 22, 14:00)**

Nicoletti, Bruno
CTO (The Foundry)

Bruno Nicoletti has worked in visual effects since graduating with a degree in Computer Science and Mathematics from Sydney University in 1987. He has worked in TV and Film production at companies such as Rushes, The Computer Film Company (now Framestore) and Animal Logic, as well as developing software commercially at Animal Logic, Softimage and Discreet Logic (now both part of Autodesk). He has developed 2D image processing software as well as 3D animation, rendering and modelling tools. In 1996 he founded The Foundry to develop visual effects plug-ins and oversaw it's initial growth. The Foundy has since grown and develops a range of applications for visual effects. Now the company's CTO, he acts as senior engineer for the company and is overseeing the effort to move The Foundry's software to a new image processing framework that can exploit CPUs and GPUs to yield dramatic speed improvements.

- **Session(s): 2125 - Developing GPU Enabled Visual Effects For Film And Video (Wednesday, Sept 22, 14:00)**

Nienhaus, Marc
Sen. Graphics Software Engineer (mental images GmbH)

Marc Nienhaus is the leading engineer of mental images subsurface data visualization solution. He received a master in mathematics with a minor in computer science from the University of Muenster and a PhD in computer science from the Hasso Plattner Institute at the University of Potsdam. In 2006, Marc researched as a post-doc at the Northwestern University and led a research project at the University of Potsdam before joining mental images R&D department in 2007. His research interest include scalable rendering and visualization techniques, GPU-based rendering and computing, and photorealistic graphics. He has published various papers on GPU-based real-time rendering, non-photorealistic rendering, and depiction techniques for symbolizing dynamics.

- **Session(s): 2014 - Scalable Subsurface Data Visualization Framework (Wednesday, Sept 22, 17:00)**

Obenschain, Keith
Computer Scientist (Naval Research Lab Code 6440)

Mr. Keith Obenschain graduated with a BS in Computer Science from the University of Illinois at Urbana-Champaignand and has been employed since graduation at the Laboratory for Computational Physics and Fluid Dynamics at NRL. Since 2002, Mr. Obenschain has been the lead software engineer for NRL's CT-Analyst project, responsible for the overall software engineering effort on CT-Analyst, including the architecture, performance enhancements, and the actual implementation.

- **Session(s): 2234 - Unstructured Finite Volume Code on a Cluster with Multiple GPUs per Node (Wednesday, Sept 22, 15:00)**

O'Brien, Kevan
(NVIDIA)

Digital filmmaking is on the cusp of another revelation, revolutions where Kevan has either helped be a catalyst or a torch bearer burning the books old rules. In his 23 year history in filmmaking, computer technology has evolved at an alarming rate enabling consumers and professionals to harness the power of desktop workstations. With the advent of CUDA's mark in the HD and film making, the next logical step is Stereo 3D on the desktop and Kevan's unique vision and insight will bring relevance to any audience as he is still a commissioned filmmaker in his own right. Kevan currently works in the Quadro group at NVIDIA hoping the IRS won't catch up with him too quickly as he's spent most of his salary on large cups of Starbucks coffee.

- **Session(s): 2279 - Working Man's Guide to 3D Video Editing (Thursday, Sept 23, 16:00)**
- **2222 - Working Man's Guide to 3D Video Editing (Tuesday, Sept 21, 14:00)**

Obukhov, Anton
Developer Technology Engineer (NVIDIA)

Anton Obukhov is a Developer Technology Engineer at NVIDIA Corporation since 2008. His field of interests include computer vision, video encoding, and multimedia processing. 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.

- **Session(s): 2075 - GPU-Accelerated Video Encoding (Thursday, Sept 23, 11:00)**

Oliker, Aaron
Managing Partner/Director 3D Technology (BioDigital)

Aaron is the 3D Technology Director and Managing Partner at BioDigital Systems. His efforts have led the development new simulation products that are currently implemented at virtual surgery centers for companies and institutions around the country. Aaron's past work with the Smile Train organization led to revolutionary new ways to empower physicians in developing nations through 3D virtual training. Aaron is a research assistant professor of Educational Informatics at New York University School of Medicine.

- **Session(s): 2146 - Virtual Surgery (Wednesday, Sept 22, 11:00)**

Ordureau, Sylvain
CEO (UsefulProgress)

The way of Sylvain Ordureau both academic and professionnal has enabled him to acquire skills in 4 keys : scientific research, computing, communications and finance. His business consist in the creation of models that test scientific hypotheses for living or inert matter. With his experience as a partner in a communication group, he has acquired control of the management and financing of projects, primarily in the service of outside companies and since Feb2003 in UsefulProgress.

- **Session(s): 4010 - Emerging Companies: CEO on Stage featuring NaturalMotion Ltd, OptiTex, and Useful Progress (Thursday, Sept 23, 15:00)**

Palamadai, Ekanathan
Research & Development Engineer (ANSYS)

Working on R&D in Nexxim Circuit Simulator

- **Session(s): 2066 - Accelerating System Level Signal Integrity Simulation (Thursday, Sept 23, 16:30)**

Pande, Vijay
Director, Folding@home Distributed Computing Project (Stanford University)

Prof. Pande is currently an Associate Professor of Chemistry and (by courtesy) of Structural Biology and of Computer Science at Stanford University. Prof. Pande received a BA in Physics from Princeton University in 1992 and PhD in physics from MIT in 1995. Prof. Pande's current research centers on the development and application of novel grid computing simulation techniques to address problems in chemical biology. In particular, he has pioneered novel distributed computing methodology to break fundamental barriers in the simulation of kinetics and thermodynamics of proteins and nucleic acids.

- **Session(s): 2007 - Folding@home: Petaflops on the Cheap Today; Exaflops Soon? (Thursday, Sept 23, 11:00)**

Pappas, Jack
Co-founder, CEO (TidePowerd)

Mathematician, entrepreneur, "yankee". Put another way: I grew up in Princeton, NJ, then attended the University of Alabama to study Mathematics. Coincidentally, this is also where I met Nick, TidePowerd's co-founder, CFO, and resident Frenchman. Born from my love of math and hatred of C programming, TidePowerd was accepted as the first-ever participant in Red Gate Software's "Springboard" startup incubator. Since then, we've built our first tool, GPU.NET, which makes GPU computing easier than ever!

- **Session(s): 2294 - GPU.NET with TidePowerd (Wednesday, Sept 22, 17:00)**

Patel, Sandeep
Assistant Professor (University of Delaware)

- **Session(s): 2035 - Simulations of Large Membrane Regions (Wednesday, Sept 22, 11:30)**

Patney, Anjul
Graduate Student (University of California, Davis)

Anjul is a third year PhD student in the Department of Electrical and Computer Engineering at University of California, Davis. He works under the guidance of Prof. John Owens in the area of graphics and computer architecture. Anjul is interested in pursuing hardware and software challenges in the design of programmable rendering architectures. Before UC Davis, he received his B.Tech. in Electrical Engineering from Indian Institute of Technology, Delhi in 2007.

- **Session(s): 2162 - Real-time Reyes: Programmable Rendering on Graphics Processors (Wednesday, Sept 22, 17:00)**

Peddie, Jon
President (Jon Peddie Research)

Dr. Jon Peddie is one of the pioneers of the graphics industry. 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, is the author of several books on computer graphics, and was named one of the most influential analysts, he is frequently quoted in trade and business publications, and contributes articles to numerous publications.

- **Session(s): 4001 - Emerging Companies: CEO on Stage featuring Elemental Technologies, Inc., Geomerics, and Milabra (Wednesday, Sept 22, 11:00)**
- **4002 - Emerging Companies: CEO on Stage featuring Allegorithmic SAS, Bunkspeed, and miGenius (Wednesday, Sept 22, 14:00)**
- **4003 - Emerging Companies Summit Panel: GPUs for Computer Vision (Wednesday, Sept 22, 15:00)**

Pedersen, Chris
Market Development Manager (NVIDIA)

Chris studied electrical engineering at BYU, but has spent most of his career as an intrapranuer — helping start new businesses in large companies. He began his career at Hewlett Packard where he helped start businesses in communications testing, video servers, consumer PC's, handheld devices and digital entertainment products and services. He's also created startup businesses and worked as an industry analyst / consultant. Today he works with ISVs to develop compelling mobile application that run on NVIDIA Tegra-powered devices.

- **Session(s): 2019 - GPU-Accelerated Internet Technologies & Trends (Tuesday, Sept 21, 14:00)**

Penner, Eric
Research Associate (Hotchkiss Brain Institute, University of Calgary, Canada)

Eric Penner is a rendering engineer at Electronic Arts and research associate at the Imaging Informatics lab at the University of Calgary. For the last few years Eric has worked on both a creative R&D team as well as shipped several game titles at EA. Before that he worked as a consultant while pursuing his M.Sc. in real-time GPU volume rendering. Eric enjoys developing real-time algorithms in traditional rendering pipelines and volume renderers. Some of his latest work includes two chapters in the upcoming book GPU Pro 2, on skin shading and pixel shader amoritization.

- **Session(s): 2235 - Advanced Medical Volume Rendering and Segmentation on the GPU (Tuesday, Sept 21, 15:00)**

Peterfreund, Natan
CTO (Playcast Media Systems)

Dr. Natan Peterfreund is a world-renowned expert with more than 20 years of research experience in video and image processing technologies. Prior to founding Playcast, Dr. Peterfreund was Chief Scientist, Video Technologies in the DSP Group [NASDAQ: DSPG]. While serving as a principal scientist in Harmonic [NASDAQ: HLIT], he was one of the authors of the H.264 compression standard. Dr. Peterfreund holds a D.Sc degree in EE from the Technion Israel Institute of Technology.

- **Session(s): 4004 - Emerging Companies: CEO on Stage featuring Cooliris, empulse GmbH, and Playcast Media Systems (Wednesday, Sept 22, 16:00)**

Peters, Alan
CTO (Universal Robotics, Inc.)

Richard Alan Peters II is an Associate Professor of Electrical Engineering at Vanderbilt University and the Chief Technology Officer of Universal Robotics, Inc (UR). During the years 2000-2008 he was a visiting researcher on the NASA-JSC Robonaut Project. For Robonaut, he developed a multimodal short-term memory system and a sensory-motor coordination (SMC) control system that learns tasks from teleoperation. This technology led to the development of UR's adaptive control systems that learn SMC to enable industrial robots to work in uncertain, dynamic environments. He is a Phi Beta Kappa graduate of Oberlin College (Ohio) where he received an A. B. in Mathematics (May 1979). He received an M.S. (1985) and a Ph.D. (1988) in Electrical Engineering from the University of Arizona, where he was a fellow of the American Electronics Association. He is the author of more than 70 scientific papers and holds four US patents. His research interests include sensory-guided robotics, computer vision, image processing, embedded systems, and mathematics.

- **Session(s): 2091 - The GPU in the Reactive Control of Industrial Robots (Thursday, Sept 23, 16:00)**

Peters, Amanda
PhD Candidate (Harvard University)

Amanda received a M.S. degree in Computer Science from Harvard University and a bachelor's degree from Duke University in both physics and computer science. She spent three years working at IBM on the Blue Gene supercomputers where her job responsibilities included applications porting and optimizing as well as system performance analysis. She is currently pursuing a PhD in Applied Physics at Harvard University with a primary research focus on computational fluid dynamics.

- **Session(s): 2163 - Leveraging GPUs for Evolutionary Game Theory (Wednesday, Sept 22, 10:00)**

Peters, David
Founder and CEO (Universal Robotics)

David Peters has overseen operations at Universal Robotics since its inception, raising capital, supervising

NVIDIA® PARALLEL NSIGHT™
POWER OF GPU COMPUTING
SIMPLICITY OF VISUAL STUDIO

GTC ATTENDEES! Visit the Parallel Nsight™ Lounge by Microsoft from 10am to 8pm each day for hands-on learning and instruction on Parallel Nsight, the industry's first development environment for GPU-accelerated applications. Get your questions answered on Parallel Nsight, Visual Studio 2010, Windows HPC Server 2008 and CUDA C/C++ development.

NVIDIA Parallel Nsight 1.5 for GPGPU Development

NVIDIA® Parallel Nsight™ software is the industry's first development environment for massively parallel computing integrated into Microsoft Visual Studio, the world's most popular development environment. Parallel Nsight is a powerful plug-in that allows programmers to develop for both GPUs and CPUs within Microsoft Visual Studio.

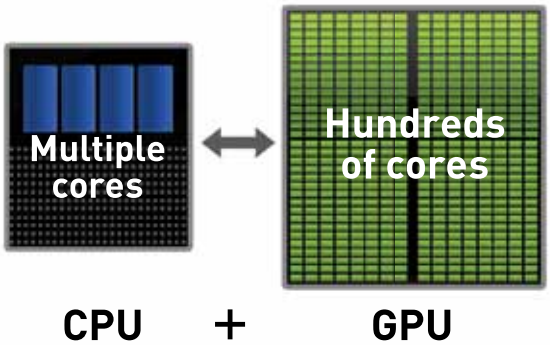

MASSIVELY PARALLEL COMPUTING

NVIDIA's Parallel Nsight in combination with Visual Studio 2010 provides developers with a highly productive, visual environment for building the best parallel capable applications available.

The power of GPU Computing is delivering game changing performance increases to the Medical, Finance, Energy, Consumer and Research fields. For over 1,000 customer stories visit www.nvidia.com/cuda

NEW FOR PARALLEL NSIGHT 1.5

- Microsoft Visual Studio 2010 Support
- Tesla Compute Cluster (TCC) Support
- Single System Compute Debugging

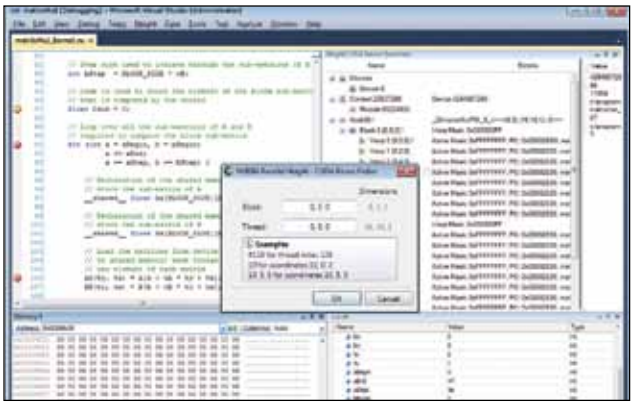

DEBUGGER

- Debug compute kernels directly on GPU hardware
- Examine thousands of threads executing in parallel using the familiar Locals, Watch, Memory and Breakpoints windows in Visual Studio
- View GPU memory directly using the standard Memory windows in Visual Studio
- Use conditional breakpoints to quickly identify and correct errors in massively parallel code
- Identify memory access violations using the CUDA C/C++ Memory Checker

ANALYZER (PROFESSIONAL VERSION ONLY)

- Capture CPU and GPU level events, including: API calls, kernel launches, memory transfers and custom application annotations
- Single correlated timeline displays all captured events
- Timeline inspection tools allow for the examination of workload dependencies
- Filter and sort captured events using specialized reporting views
- Profile CUDA kernels using GPU performance counters
- *NEW* DirectCompute profiling

To learn more about Parallel Nsight, go to **www.nvidia.com/ParallelNsight**

NVIDIA

Visual Studio PARTNER

engineering and leading marketing initiatives. Under David's leadership, Universal has developed innovative products that surpass current technologies while building a business ecosystem, which includes channel partner, Yaskawa, the leading producer of industrial robots. In the entertainment industry as a producer for 17 years, David facilitated the development and production of feature films, managing complex financing and intellectual property affairs. David has a history of completing complex projects within budget, resulting in profits to investors, and garnering status as a Completion Bond approved executive.

- **Session(s): 4008 - Emerging Companies: CEO on Stage featuring ICD and Universal Robotics (Thursday, Sept 23, 11:00)**

Pfister, Hanspeter
Professor (Harvard University)

Hanspeter Pfister is Gordon McKay Professor of the Practice in the School of Engineering and Applied Sciences at Harvard University. His research lies at the intersection of visualization, computer graphics, and computer vision. Dr. Pfister has a Ph.D. in Computer Science from the State University of New York at Stony Brook and an M.S. in Electrical Engineering from the Swiss Federal Institute of Technology, Zurich, Switzerland.

- **Session(s): 2262 - CUDA Centers of Excellence Super-Session I (Tuesday, Sept 21, 14:00)**
- **2281 - Domain-Specific Languages (Wednesday, Sept 22, 15:00)**

Phillips, Everett
Applied Engineer - GPU Computing (NVIDIA)

is an Engineer in the Tesla Performance Group at NVIDIA where he works in the area of GPU Computing, HPC, and Clusters. He holds an M.S. in Mechanical and Aeronautical Engineering from the University of California, Davis.

- **Session(s): 2057 - CUDA-Accelerated LINPACK on Clusters (Tuesday, Sept 21, 14:00)**

Phillips, James
Senior Research Programmer (University of Illinois)

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. 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.

- **Session(s): 2054 - NAMD, CUDA, and Clusters: Taking GPU Molecular Dynamics Beyond the Desktop (Thursday, Sept 23, 14:00)**

Phung, Huynh
*Research Engineer (A*STAR Institute of High Performance Computing)*

Huynh Phung received a Ph.D. degree in Computer Science from School of Computing, National University of Singapore. Currently, Huynh is working as a research engineer at A*STAR Institute of High Performance Computing, Singapore.

- **Session(s): 2109 - Migration of a Complete 3D Poisson Solver from Legacy Fortran to CUDA (Wednesday, Sept 22, 10:30)**

Pien, Homer
Director of the Laboratory for Medical Imaging and Computations (MGH / HMS)

Homer Pien is the Director of the Laboratory for Medical Imaging and Computations, in the Department of Radiology at Massachusetts General Hospital and Harvard Medical School. Dr. Pien's research focuses on the use of computations to enable new clinical imaging applications, for cardiovascular disease, trauma, and cancer.

- **Session(s): 2282 - GPU-Enabled Biomedical Imaging (Wednesday, Sept 22, 17:00)**

Pierce, Dan'l
Partner (Access Analytics Int'l, LLC)

Dan'l Pierce is a principal partner of Access Analytics Int'l, LLC. Dan'l has also served as Director of Business and Market Development for Cray, WW Lead HPC Scientist for Intel, Senior Vice President at Washington Mutual, and in technical and management roles at The Boeing Company. Dan'l received his PhD in Applied Mathematics from North Carolina State University, and his MS and BS degrees in Mathematics and Computer Science from Northern Illinois University.

- **Session(s): 2213 - BCSLIB-GPU: Significant Performance Gains for CAE (Thursday, Sept 23, 15:00)**

Pimentel, Ken
Director, Media & Entertainment (Autodesk)

Ken Pimentel is responsible for product strategy for 3ds Max, Showcase and other solutions in Autodesk's Media and Entertainment Division. Prior to Autodesk, Ken worked at Engineering Animation and helped major manufacturers incorporate advanced visualization techniques to reduce the need for physical prototypes. Before EAI, Ken helped start Sense8, one of the first VR companies. Ken holds a BS EE degree from UC Davis and is coauthor of "Virtual Reality: Through the New Looking Glass".

- **Session(s): 2165 - Rendering Revolution (Tuesday, Sept 21, 11:00)**

Pinskiy, Dmitriy
Sr Software Engineer (Walt Disney Animation Studios)

Senior Software engineer at Walt Disney Animation Studious; designed and developed CG animation tools for feature animation movies Chicken Little, Meet the Robinsons, Bolt, Tangled. Prior to that, working at Alias|Wavefront (now Autodesk) on Maya, an award winning software, implemented various animation deformers as well as modeling / sculpting tools. Published a number of computer graphics research papers (as recent as Eurographics 2010) and presented SIGGRAPH talks.

- **Session(s): 2284 - GPU implementation of Collision-Based Deformation (Wednesday, Sept 22, 17:30)**

Pinto, Nicolas
PhD Student (MIT)

Nicolas Pinto is a PhD Student in Computational Neuroscience at MIT. He is currently a member of the DiCarlo Lab at MIT, the Sinha Lab for Vision Research at MIT, and the Cox Visual Neuroscience Group at Harvard/Rowland. His research interests lie at the intersection of Brain and Computer Sciences.

- **Session(s): 2204 - Bridging GPU Computing and Neuroscience to Build Large-Scale Face Recognition on Facebook. (Wednesday, Sept 22, 14:00)**

- **2176 - Easy GPU Meta-programming: A Case Study in Biologically-Inspired Computer Vision (Thursday, Sept 23, 10:00)**

Pissoort, Davy
Professor (KHBO-FMEC)

Davy Pissoort was born in 1978. He received theM.S. and Ph.D. degrees in electrical engineering from Ghent University, Ghent, Belgium, in 2001 and 2005, respectively. From October 2005 to October 2006, he was a Postdoctoral Researcher with the Fund for Scientific Research-Flanders (Belgium) (FWO-Vlaanderen) at the Department of Information Technology, Ghent University. In November 2006, he joined the Eesof-EDA Department, Agilent Technolgies, Ghent, Belgium, as a Research Engineer. Since August 2009, he is with the KHBO, Belgium where he is also head of the Flanders' Mechatronics Engineering Centre. His current research interests include the development of fast and efficiënt electromagnetic simulation methods, electromagnetic compatibility, as well as the analysis and testing of the mechanical and thermal reliability of electronic modules.

- **Session(s): 2080 - Tackling Multi-Gigabit Design Challenges with a Practical Virtual EMI/ESD Lab (Wednesday, Sept 22, 15:00)**

Prevrhal, Sven
Staff Research Scientist (Philips)

PhD in Physics from Technical University Vienna, Austria 1997: work on medical imaging of Osteoporosis Faculty Member University of California, San Francisco, until 2009: work on Computed Tomography image reconstruction. Now with Philips Medical Imaging, San Jose: work on Positron Emission Tomography image reconstruction.

- **Session(s): 2211 - Modern Architecture for Massively Parallel Medical Tomographic Image Reconstruction on a GPU Cluster (Wednesday, Sept 22, 15:00)**

Price, Daniel
Engineer (EM Photonics, Inc.)

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.

- **Session(s): 2154 - The Impact of Data Movement on GPU Performance (Wednesday, Sept 22, 16:00)**

Pryor, Gallagher
VP of Research (AccelerEyes)

Gallagher Pryor is VP of Research at AccelerEyes and a cofounder. He is the inventor of GPU computing in MATLAB. He holds degrees in Computer Science from Georgia Institute of Technology.

- **Session(s): 2268 - Think Data-Parallel! Building Data-Parallel Code with M (Tuesday, Sept 21, 15:30)**

Ramey, Will
Sr. Product Manager, GPU Computing (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.

- **Session(s): 2004 - Languages, APIs and Development Tools for GPU Computing (Pre-Conference Tutorial) (Monday, Sept 20, 13:00)**

Rasmusson, Allan
PhD Candidate (University of Aarhus (NVIDIA intern))

Works within the field of Quantitative Tissue Analysis, primarily with optimizing the process of counting and measuring cells in microscopic images of histological tissue sections using 3D visualization and medical imaging computations.

- **Session(s): 2021 - Efficient Volume Segmentation on the GPU (Wednesday, Sept 22, 17:00)**

Raue, Kristian
CEO and Founder (Jedox Business Intelligence)

2002 - today: CEO & Founder, Jedox Business Intelligence
1991 - 2000: CEO & Founder, IntelliCube AG
1989 - 1991: Management Consultant
1982 - 1989: Technical University of Darmstadt (double major in engineering and business administration)

- **Session(s): 4005 - Emerging Companies: CEO on Stage featuring Jedox Business Intelligence, Rocketick, and Softkinetic (Wednesday, Sept 22, 17:00)**

Reid, Ian
Chief Commercial Officer (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.

- **Session(s): 2063 - Banking on Monte Carlo... and Beyond (Thursday, Sept 23, 15:30)**

Reil, Torsten
CEO (NaturalMotion Ltd)

Torsten Reil is co-founder and CEO of NaturalMotion. He holds a BA in Biology from Oxford University and 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's Zoology department, from which he spun off NaturalMotion. Torsten has been named amongst MIT's TR100 global top innovators, Next-Gen's 25 People in the Games Industry, and Develop magazine's 25 Game Changers.

- **Session(s): 4010 - Emerging Companies: CEO on Stage featuring NaturalMotion Ltd, OptiTex, and Useful Progress (Thursday, Sept 23, 15:00)**

Reyna, Nabor
Graduate Student(Rice University)

Nabor Reyna Jr. is a third year graduate student in the Computational and Applied Mathematics Department at Rice University. Reyna is currently finishing up his Masters work in Compressive Sensing (CS). His interests consists of high performance computing, image processing as well as face recognition. Reyna is also a member in organizations such as ACM, SIAM and

SACNAS.

- **Session(s): 2300 - High-Performance Compressive Sensing using Jacket (Wednesday, Sept 22, 10:30)**

Richardson, Alan
Graduate Student (M.I.T)

BA in Theoretical Physics from Trinity College Dublin, MSc in High Performance Computing from EPCC at the University of Edinburgh. Currently pursuing a PhD in Geophysics at MIT.

- **Session(s): 2167 - Designing a Geoscience Accelerator Library Accessible from High Level Languages (Wednesday, Sept 22, 17:00)**

Rios, Joseph
Research Aerospace Engineer (NASA)

Joseph is a researcher at NASA Ames Research Center. His work focuses on computational issues associated with Traffic Flow Management in the National Airspace System. He completed his B.A. in Mathematics from UC Santa Cruz before teaching high school in the Peace Corps and in California. Upon returning to school, he received his M.S. in computer science from Cal State Hayward and his Ph.D. in Computer Engineering from UC Santa Cruz.

- **Session(s): 2214 - Faster Simulations of the National Airspace System (Tuesday, Sept 21, 11:00)**

Roberts, Mike
Research Assistant (Hotchkiss Brain Institute, University of Calgary, Canada)

Mike Roberts is currently pursuing his MSc in Computer Science at the University of Calgary in Canada. Mike is also researching interactive level set segmentation of large medical data sets. Recently, he published a paper on his research at the ACM SIGGRAPH/Eurographics Conference on High Performance Graphics 2010.He also did an internship at NVIDIA during the summer of 2009 and helped to develop the world's first IDE-integrated native GPU debugger entitled NVIDIA Parallel Nsight.

- **Session(s): 2235 - Advanced Medical Volume Rendering and Segmentation on the GPU (Tuesday, Sept 21, 15:00)**
- **2236 - A Work-Efficient GPU Algorithm for Level Set Segmentation (Thursday, Sept 23, 09:00))**

Robison, Austin
Lead Developer, OptiX integration (NVIDIA)

Austin Robison is a Research Scientist at NVIDIA working as part of the OptiX team on GPU ray tracing. His research interests include high performance ray tracing, physically-based rendering and hybrid rendering. Austin holds a B.S. in Computer Science from the University of Chicago and an M.S. in Computer Science from the University of Utah.

- **Session(s): 2250 - GPU Ray Tracing Exposed: Under the Hood of the NVIDIA OptiX Ray Tracing Engine (Tuesday, Sept 21, 17:00)**

Rogan, Aaron
Research Scientist and System Adminstrator (Neva Ridge Technologies)

I currently work for Neva Ridge Technologies where I specialize in radar image processing. Over the past year and half I have transitioned legacy code to run on a GPU which has resulted in performance improvements anywhere from 20x to 400x depending ont he algorithm. My interests are in GPGPU programming, image processing, algorithm development and numerical simulations.

- **Session(s): 2003 - Using CUDA to Accelerate Radar Image Processing (Thursday, Sept 23, 15:00)**

Rollin, Philippe
Applied Engineer (NVIDIA)

Philippe Rollin is currently a Graphics Application Engineer in the Professional Solution Group at NVIDIA, where he is investigating new tools and technologies to bring the best out of their products. In the past, Philippe was the 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.

- **Session(s): 2227 - OpenGL 4.0 Tessellation for Professional Applications (Tuesday, Sept 21, 15:00)**

Rossbach, Christopher
Researcher (Microsoft Research)

Chris Rossbach is a researcher at Microsoft Research Silicon Valley. He received his Ph.D. in Computer Science from the University of Texas at Austin, and his B.S. in Computer Systems Engineering from Stanford University. His research focuses on operating systems and architecture, and emphasizes the development of better tools for managing and taking advantage of concurrency.

- **Session(s): 2124 - Operating System Abstractions for GPU Programming (Thursday, Sept 23, 10:00)**

Rubin, Eri
Senior CUDA R&D developer (OptiTex)

Eri Rubin is a Senior CUDA R&D Developer. From 2006 to present, he is the senior R&D Developer and lead developer on 3 products: OptiTex 2D/3D, CAD/CAM, Fashion & Textile. He is also the head of GPGPU projects, porting an implicit cloth solver engine to CUDA; 3D engine architecture work; developing CG Tools for Clothes Design within a 3D interface; and developing Maya, MentalRay projects and plugins.

- **Session(s): 2246 - The challenges of integrating CUDA engines into an existing package, yet not sinking the boat (Wednesday, Sept 22, 14:00)**

Ruge, Thomas
Software Manager (NVIDIA)

- **Session(s): 2024 - NVIDIA Acceleration Engines Overview (Pre-Conference Tutorial) (Monday, Sept 20, 13:00_**

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.

- **Session(s): 2015 - Efficient Tridiagonal Solvers for ADI methods and Fluid Simulation (Tuesday, Sept 21, 14:00)**

Salian, Satish
Manager CUDA Debugger Tools (NVIDIA)

Satish Salian is the Manager for CUDA debugger 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 before moving to CUDA was responsible for the development of

NVIDIA's Graphics tools and NVAPI SDK. Satish received his Bachelors degree in Computer Engineering from University of Pune, India.

Session(s): 2002 - CUDA Debugging on Linux and MacOS with cuda-gdb (Thursday, Sept 23, 10:00)

Sanders, Jason

Senior Software Engineer (NVIDIA)

Jason Sanders is a senior software engineer in the CUDA Platform Group at NVIDIA. While at NVIDIA, he helped develop early releases of CUDA system software and co-wrote the book CUDA by Example. Jason received his M.S. in computer science from the University of California Berkeley where he published research in GPU computing, and he holds a B.S.E. in electrical engineering from Princeton University. Prior to joining NVIDIA, he previously held positions at ATI Technologies, Apple, and Novell. When he's not coding or writing books, Jason is typically in the gym, playing soccer, or shooting photos.

Session(s): 2131 - Introduction to CUDA C (Pre-Conference Tutorial) (Monday, Sept 20, 14:30)

Sarzo, Rudy

Principal (SMI)

Rudy has been a professional recording and performing artist worldwide for over 20 years. As a member of Ozzy Osbourne's band, from March 1981 to September 1982, Rudy toured the world in support of the "Blizzard of Ozz" and "Diary Of a Madman" records. His bass playing can be heard on Ozzy's multimillion selling CD "Tribute" and "Speak of the Devil" CD and DVD.

Session(s): 2279 - Working Man's Guide to 3D Video Editing (Thursday, Sept 23, 16:00)

Scherl, Holger

Computer Scientist (Siemens AG)

Holger is a computer scientist in the field of medical image processing. He pursues his PhD studies at the University of Erlangen-Nuremberg, Germany, specializing in hardware-accelerated cone-beam CT reconstruction. Since 2007 he is a system architect in R&D at Siemens Healthcare and focuses on the development and implementation of cone-beam CT reconstruction algorithms on graphics hardware.

Session(s): 2096 - High-Speed CT Reconstruction in Medical Diagnosis & Industrial NDT Applications (Tuesday, Sept 21, 11:00)

Schneider, Jens

Postdoctoral Fellow (King Abdullah University of Science and Technology)

Dr. Jens Schneider received his MA from RWTH Aachen, Germany in 2003 and his doctorate from Technische Universitaet Muenchen in 2009. He is currently a postdoctorate fellow at the King Abdullah University of Science and Technology (KAUST) where he works in the Geometric Modeling and Scientific Visualization Center. His research interests include GPU-based algorithms, Computer Graphics, Scientific Visualization, and GPU-friendly Data Compression.

Session(s): 2139 - Interactive Histology of Large-Scale Biomedical Image Stacks (Wednesday, Sept 22, 14:00)

Schuh, Andrew

Project Manager (University of Illinois)

Session(s): 2249 - New Programming Tools GPU Computing (Wednesday, Sept 22, 10:00)

Schulten, Klaus

Professor (University of Illinois, Urbana-Champaign)

Klaus Schulten received his Ph.D. from Harvard University in 1974. He is Swanlund Professor of Physics and is also affiliated with the Department of Chemistry as well as with the Center for Biophysics and Computational Biology. Professor Schulten is a full-time faculty member in the Beckman Institute and directs the Theoretical and Computational Biophysics Group at the University of Illinois Urbana-Champaign, IL. Honors and awards: Award in Computational Biology 2008; Humboldt Award of the German Humboldt Foundation (2004); University of Illinois Scholar (1996); Fellow of the American Physical Society (1993); Nernst Prize of the Physical Chemistry Society of Germany (1981).

Session(s): 1002 – Day 2 Keynote with Dr. Klaus Schulten, University of Illinois at Urbana-Champaign

Sheehan, Andrew T.

Managing Director (Sutter Hill Ventures)

Andy Sheehan focuses his investments on internet software, services and digital media companies. Andy currently is a director of Buzz Media, Inc., Global Liquid Markets, LLC, Grain Communications Group, Inc., and Yext, Inc. His prior directorships have included, @ Road (acquired by Trimble), AllBusiness.com, BakBone Software, Datran Media, Intermix Media and Myspace (acquired by News Corp.) and ReachLocal. Andy joined the firm in 2007 from VantagePoint Venture Partners, where he was a managing director. Previously, he worked at Alex. Brown & Sons and ABS Capital Partners. Andy received his BA from Dartmouth College with a degree in English. He earned his MBA in 1985 from the Wharton School.

Session(s): 4009 - Emerging Companies Summit Panel: The "New Normal" For Building Emerging Companies Based On Disruptive Technologies (Thursday, Sept 23, 14:00)

Shoemaker, Austin

CTO and Co-Founder (Cooliris)

Austin is CTO and co-founder of Cooliris. Austin was a master's degree student in Computer Science at Stanford University specializing in artificial intelligence, and 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.

Session(s): 4004 - Emerging Companies: CEO on Stage featuring Cooliris, empulse GmbH, and Playcast Media Systems (Wednesday, Sept 22, 16:00)

Silva, Claudio

Professor (University of Utah)

Claudio T. Silva is a professor of computer science and a faculty member of the Scientific Computing and Imaging (SCI) Institute at the University of Utah. He coauthored more than 150 technical papers and eight U.S. patents, primarily in visualization, geometric computing, and related areas. He received IBM Faculty Awards in 2005, 2006, and 2007, and best paper awards at IEEE Visualization 2007 and IEEE Shape Modeling International 2008. He is a member of the ACM, Eurographics, and IEEE.

Session(s): 2248 - Parallel Processing on GPUs at the University of Utah (Wednesday, Sept 22, 14:00)

Sinclair, Matt

Research Assistant (UW-Madison)

Matt Sinclair is a Ph.D. student at the University of Wisconsin-Madison in the Electrical and Computer Engineering Department. He received his B.S. degrees

in Computer Engineering and Computer Science with Honors from the University of Wisconsin-Madison in 2009. His interests lie in processor microarchitecture and high performance computing.

- **Session(s): 2044 - GRASSY: Leveraging GPU Texture Units for Asteroseismic Data Analysis (Wednesday, Sept 22, 15:00)**

Snepvangers, Jeroen
President and CEO (RTT USA)

Jeroen Snepvangers is President and CEO of RTT USA, Inc., a subsidiary of RTT AG, a public company trading at the Frankfurt Stock Exchange. RTT is the largest real-time 3D computer graphics software and CGI animation services provider to the Automotive and Aircraft industries. The company serves its customers in industrial design and digital (3D CGI) marketing. Prior to joining RTT, Jeroen was a Management Consultant at Urban Science Applications Inc., a technology-based management consultancy for the automotive and financial industries, focusing on optimizing retail networks. He obtained his MBA at IMD, Switzerland in 2002. He has a Bachelors Degree in Applied Mathematics from Warwick University, UK. He has an IB, International Baccalaureate from St. Clare's College, Oxford, UK

- **Session(s): 4007 - Emerging Companies: CEO on Stage featuring Aqumin, RTT, and Scalable Display Technologies (Thursday, Sept 23, 10:00)**

Solano, Lizandro
(Iowa State University)

Lizandro Solano-Quinde received his MSc. in Electrical Engineering in 2006 at Iowa State University; currently he is a Ph.D. candidate in Computer Engineering at Iowa State University under the advisement of Dr. Brett Bode and Dr. Arun Somani. He is affiliated to the Scalable Computing Laboratory at the Ames Laboratory, Department of Energy. His research interests are in the fields of High Performance Computing, Computer Architecture and Fault Tolerance.

- **Session(s): 2292 - Implementation of High-Order Adaptive CFD Methods on GPUs (Thursday, Sept 23, 10:30)**

Somani, Arun
Anson Marston Professor (Iowa State University)

Arun K. Somani is currently Anson Marston Distinguished Professor and Jerry R. Junkins Endowed Chair Professor of Electrical and Computer Engineering at Iowa State University where he first served as David C. Nicholas Professor during 1997-2002. He earned his MSEE and PhD degrees in electrical engineering from the McGill University, Montreal, Canada, in 1983 and 1985, respectively. He worked as Scientific Officer for Govt. of India, New Delhi from 1974 to 1982 and as a faculty member at the University of Washington, Seattle, WA from 1985 to 1997 in electrical engineering and computer science and engineering departments where he was promoted to the level of Professor in September 1995.

- **Session(s): 2292 - Implementation of High-Order Adaptive CFD Methods on GPUs (Thursday, Sept 23, 10:30)**

Southard, Dale
Senior Solution Architect (NVIDIA)

As a Senior Solution Architect with NVIDIA, Dale assists in the integration of GPUs in HPC and Cloud computing environments. He was previously with Lawrence Livermore National Lab as the architect of their visualization hardware solutions and part of the systems team that deployed and managed their large HPC machines.

- **Session(s): 2017 - Lessons Learned Deploying the World's First GPU-Based Petaflop System (Tuesday, Sept 21, 15:00)**

Spatz, Pierre
Head of Quantitative Research (Murex SAS)

Pierre joined Murex in 1989 and has a master degree in computer science and applied mathematics from ENSIMAG. After various leading positions in the Murex software development team, Pierre launched the Murex Analytics initiative in 2002.

- **Session(s): 2032 - Practical Methods Beyond Monte Carlo in Finance (Thursday, Sept 23, 10:00)**

Srinivasan, Savitha
Corporate Venture Partner (IBM)

Savitha Srinivasan is a Corporate Venture Partner in IBM's Venture Capital Group in Corporate Strategy where she develops strategic relationships with venture capitalists and their portfolio companies to leverage external innovation for mutual strategic advantage. She has nearly 20 years of experience at IBM in leadership roles addressing the strategic priorities of IBM's Services businesses. She leads the strategic development of IBM's Services venture ecosystem, with each of the Global Technology Services business units – Strategic Outsourcing, Integrated Technology Services and Managed Business Process Services by early identification of companies, fostering pilots, partnerships and contributing to M&A pipeline.

- **Session(s): 4007 - Emerging Companies: CEO on Stage featuring Aqumin, RTT, and Scalable Display Technologies (Thursday, Sept 23, 10:00)**
- **4008 - Emerging Companies: CEO on Stage featuring ICD and Universal Robotics (Thursday, Sept 23, 11:00)**

Stam, Joe
Sr. Applications Engineer (NVIDIA)

JOSEPH STAM is a Senior Applications Engineer for NVIDIA Corporation. He has a focus on computer vision, video, and image processing applications of Graphics Processors for both professional and embedded products. Prior to joining NVIDIA in 2007, he worked in the automotive industry for 12 years on research and development of imaging hardware and computer vision algorithms for vehicle based vision products. Joe received a B.S. degree in Engineering Physics & Computer Science from Hope College in Holland, Michigan and an M.S. degree in Electrical Engineering from Michigan State University. He is an inventor on 82 U.S. patents and several foreign patents, many of which relate to computer vision software and imaging hardware technologies. Joe resides in Holland Michigan with his wife and four children.

- **Session(s): 2215 - Extending OpenCV with GPU Acceleration (Thursday, Sept 23, 10:00)**
- **4003 - Emerging Companies Summit Panel: GPUs for Computer Vision (Wednesday, Sept 22, 15:00)**

Stephan, Philippe
CTO (RMS)

Philippe Stephan is the Chief Technology Officer of RMS. Prior to RMS, he directed product development as the CTO of San Francisco based Moody's KMV, the award winning credit risk analytics vendor. Philippe has also held senior management positions at Internet startups and built derivatives risk systems for SocGen and CA Lazard Financial Product Bank in London. Philippe started his career as a key contributor to the

Supercomputer Muscle.
Briefcase Size.

NEXT COMPUTING

PNY Professional Graphics Solutions NVIDIA QUADRO

NextComputing and PNY Technologies present a portable development and demonstration platform for Fermi-based NVIDIA® Quadro® accelerated applications

See us in the PNY Pavilion
or visit us at NextComputing.com

development of the Eiffel programming language in the early 1990s.

- **Session(s): 2077 - Catastrophic Risk Management: Fast and Flexible with GPU Analytics (Wednesday, Sept 22, 17:00)**

Stich, Timo
Developer Technology Engineer (NVIDIA)

Timo Stich is a Developer Technology Engineer for NVIDIA Corporation. His focus is on image processing and compute applications of Graphics Processors. Prior to joining NVIDIA, he was a member of the 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.

- **Session(s): 2087 - Fast High-Quality Panorama Stitching (Thursday, Sept 23, 14:00)**

Stier, Jochen
Founder (Geist Software Labs)

Jochen Stier holds a Ph.D. in Computer Science from the University of Victoria. His research interest are real time 3D visualization and high performance computing. Prior to his Ph.D. studies, Jochen spent several years as a software engineer and architect in the industrial automation and telecommunication industry. In 2007, Jochen founded Geist Software Labs to market OpenCL Studio and provide solutions for high performance computing and 3D visualization.

- **Session(s): 2148 - Rapid Prototyping and Visualization with OpenCL Studio (Tuesday, Sept 21, 15:00)**

Stone, Christopher
Research Scientist (Intelligent Light)

Dr. Stone received his PhD from the Georgia Institute of Technology's School of Aerospace Engineering in 2003. He research areas include computational fluid dynamics and high performance parallel computing.

- **Session(s): 2110 - Acceleration of a Novel Rotorcraft Wake Simulation (Thursday, Sept 23, 10:00)**

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.

- **Session(s): 2073 - High Performance Molecular Simulation, Visualization, and Analysis on GPUs (Wednesday, Sept 22, 16:00)**

Strzodka, Robert
Senior Researcher (Max Planck Institut Informatik)

Robert Strzodka has been the head of the research group, Integrative Scientific Computing, at the Max Planck Institut Informatik since 2007. His research focuses on efficient interactions of mathematic, algorithmic and architectural aspects in heterogeneous high performance computing. Previously, Robert was a visiting assistant professor at the Stanford University and until 2005, a postdoc at the caesar research center in Bonn. He received his doctorate in numerical mathematics from the University of Duisburg in 2004.

- **Session(s): 2038 - The Best of Both Worlds: Flexible Data Structures for Heterogeneous Computing (Wednesday, Sept 22, 14:00)**

Surati, Rajeev
President (Scalable Display Technologies)

Cofounder and President Scalable Display Technologies, founder and CTO Flash Communications sold to Microsoft as a basis for their msn messenger services, founder and president photo.net worlds largest amateur photographer community, sold to NameMedia. Education: SB, SM, Ph.D. MIT in Electrical Engineering and Computer Science. DOE Computational Fellow 1995.

- **Session(s): 2134 - Ultra High Resolution Displays and Interactive Eyepoint Using CUDA (Wednesday, Sept 22, 10:00)**

Tai, Bill
General Partner (Charles River Ventures)

Bill Tai joined CRV in 2002 to lead the firm's West coast practice. He has funded companies as a venture capitalist since 1991. 16 startups he has backed became publicly listed companies, among them 8x8 Inc., Award Software, Network Peripherals, Microtune, Terayon, and Transmeta. His focus areas include disruptive enabling technology, digital media and web based services. His current portfolio includes Glyde - a next generation commerce platform company, Maxthon - a web browser that has been downloaded over 430 million times, Nantero – a carbon nanotube memory company, and Scribd – a social publishing company with over 50M monthly users, among others. He is a 'hand on' early stage VC, having served as Chairman of IPInfusion (acquired by Access Corp), founded as CEO iAsiaWorks (IPO in 2000) and founded as President TRADEBEAM, a leader in SAAS based Global Trade Management (acquired by CDC Software). Prior to his VC career, he positioned public offerings of Adaptec, Atmel, Cirrus Logic, Dallas Semiconductor, Exar, and Zilog while establishing the semiconductor research practice at Alex. Brown and Sons.

- **Session(s): 4010 - Emerging Companies: CEO on Stage featuring NaturalMotion Ltd, OptiTex, and Useful Progress (Thursday, Sept 23, 15:00)**
- **4011 - Emerging Companies: CEO on Stage featuring Cinnafilm, Inc., Perceptive Pixel, and Total Immersion (Thursday, Sept 23, 16:00)**

Tal, Uri
CEO (Rocketick)

Uri Tal has 12 years of experience in management, design and implementation of hardware acceleration technologies. Before founding Rocketick, Uri Managed a large R&D team that developed FPGA-based acceleration solutions in the intelligence corps of the Israeli Defense Forces. And was a system architect in Siliquent / Broadcom. B.Sc. (Summa Cum Laude) and M.Sc. in Electrical Engineering ,Technion.

- **Session(s): 4005 - Emerging Companies: CEO on Stage featuringJedox Business Intelligence, Rocketick, and Softkinetic (Wednesday, Sept 22, 17:00)**

Tarui, Kento
Researcher (AquaCast Corporation)

Kento Tarui received his Masters of Engineering from Tokyo Institute of Technology in 2005. He then earned his Ph.D.in Engineering from Tokyo Institute of Technology in 2008. Currently, he is a researcher at the Aquacast

Corporation.

- **Session(s): 2114 - Cascaded HOG on GPU (Thursday, Sept 23, 16:00)**

Taufer, Michela
Assistant Professor (University of Delaware)

My research interests include new algorithms and architectures for resource- and time-expensive applications in computational chemistry, physics, and biology; effective migration of large scale simulations to global computing systems based on public resources; and performance analysis, -modeling and -optimization of large scale simulations on heterogeneous, distributed systems.

- **Session(s): 2034 - Reformulating Algorithms for the GPU (Wednesday, Sept 22, 11:00)**
- **2035 - Simulations of Large Membrane Regions (Wednesday, Sept 22, 11:30)**

Taylor, John
Science and Business Leader (CSIRO)

Dr John Taylor is the Science and Business leader for Computational and Simulation Sciences at CSIRO, Australia. Dr Taylor has written more than 140 articles and books on computational and simulation science and its application to such areas as climate change, global biogeochemical cycles, air quality, water resources and environmental policy. His research has addressed local and global scientific and environmental policy issues. He has extensive experience developing applications on high performance computers. At CSIRO, Dr Taylor has led the development and deployment of Australia's and one of the world's first CPU+GPU clusters.

- **Session(s): 2301 - GPU Cluster Computing: Accelerating Scientific Discovery (Thursday, Sept 23, 09:00)**

Thamm, Tom-Michael
VP Products (mental images GmbH)

Mr. Thamm is the Vice President for Products at mental images and is responsible for the product management of all products, in particular of mental ray, mental mill, and RealityServer. In addition, he is managing and coordinating the customer support in cooperation with Engineering. Mr. Thamm joined mental images in 1989. He has led several key projects of mental images, such as the definition of the extended OBJ format and the integration of mental ray into many of the major CAD systems. He has studied Mathematics and subsequently developed free-form surface algorithms and various 3D file formats.

- **Session(s): 2014 - Scalable Subsurface Data Visualization Framework (Wednesday, Sept 22, 17:00)**

Thomason, Lee
Principal Scientist (Adobe Systems)

Lee Thomason is a principal scientist and Flash Player architect for Adobe Systems. He prototypes GPU technology and leads the GPU development for the Flash Player.

- **Session(s): 2060 - GPUs in a Flash: Mapping the Flash Animated Software Vector Rendering Model to the GPU (Tuesday, Sept 21, 17:00)**

Thrun, Sebastian
Professor / Distinguished Engineer (Stanford University / Google)

Sebastian Thrun is a professor of computer science and electrical engineering at Stanford, where he directs the Stanford AI Lab. He is also a distinguished engineer at Google. Thrun's team won the DARPA Grand Challenge, a US-Government sponsored autonomous robot race that took place in 2005. Thrun also pioneered the scientific field of probabilistic robotics, and he co-invented Google Street View. In recognition of his contributions, Thrun has been elected into the US National Academy of Engineering and the German Academy of Sciences. He is an elected fellow of the AAAI, ECCAI, and WTN. Popular Science included Thrun in their "Brilliant Ten," Forbes Magazine in their "E-Gang" members, Scientific American in their list of 50 world technology and policy leaders, and Fortune selected him as one of the 50 smartest people in tech. Wired Magazine awarded Thrun's robot Stanley the top spot in the most influential robots of all times. The robot is now part of a permanent exhibition in the Smithsonian Museum of American History. Thrun has authored 11 books and over 300 scientific articles.

- **Session(s): 1003 – Closing Keynote with Dr. Sebastian Thrun, Stanford**

Toelke, Jonas
Chief Computational Software Development (Ingrain)

Dr. Tölke's work is focused on the design, implementation and verification of simulation engines for multi-phase flow on Supercomputers and other high performance computing hardware. He holds a PhD in engineering from Technische Universität München and a "Venia Legendi" (Habilitation) in computational engineering in fluid mechanics from Technische Universität Braunschweig. Before joining Ingrain, he provided consulting services in the field of multi-phase simulations and worked as director for research software engineering at Exa Corp.

- **Session(s): 2170 - Lattice Boltzmann Multi-Phase Simulations in Porous Media using GPUs (Wednesday, Sept 22, 15:00)**

Tombroff, Michel
CEO (Softkinetic)

Prior to his CEO role at Softkinetic Michel spent 17 years in the software industry in the start-up, pre-IPO and public stages. He spent 8 years at TIBCO Software, where his last position was VP Sales, and 7 years at Chorus Systems (acquired by Sun Microsystems), where his last role was Head of Engineering. He received a B.S. EE from University of Brussels and a MSc Computer Science from University of California, Santa Barbara.

- **Session(s): 4005 - Emerging Companies: CEO on Stage featuringJedox Business Intelligence, Rocketick, and Softkinetic (Wednesday, Sept 22, 17:00)**

Tomov, Stan
(University of Tennessee)

Stanimire (Stan) Tomov is Research Scientist at ICL and Adjunct Assistant Professor in EECS at UTK. He received Ph.D. in Mathematics from TAMU in 2002 and held positions at LLNL and BNL. Stan's research interests are in parallel algorithms, numerical analysis, and high-performance scientific computing. He co-leads the UTK's CCOE on the development of MAGMA, a new generation of linear algebra libraries, extending the sequential LAPACK-style algorithms, for highly parallel, GPU and multicore heterogeneous architectures.

- **Session(s): 2138 - Faster, Cheaper, Better – Hybridization of Linear Algebra for GPUs (Thursday, Sept 23, 09:00)**
- **2263 - CUDA Centers of Excellence Super-Session II (Tuesday, Sept 21, 15:00)**

Townsend, Richard
Assistant Professor (University of Wisconsin-Madison)

Richard Townsend is a computational astrophysicist

at the University of Wisconsin-Madison, interested in the rotation, oscillations, magnetic fields and outflows of massive, luminous stars. His recent research has focused on investigating how GPU computing can be brought to bear on the steep data analysis challenges faced in Asteroseismology. Initial projects showing particular promise include GPU-accelerated period searching in non-uniformly sampled data, and fast spectrum interpolation by leveraging the untapped capabilites of GPU texturing units.

- **Session(s): 2082 - CU-LSP: GPU-based Spectral Analysis of Unevenly Sampled Data (Wednesday, Sept 22, 10:00)**

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.

- **Session(s): 2159 - Programming the NVIDIA Digital Video Pipeline with Direct3D (Pre-Conference Tutorial) (Monday, Sept 20, 14:30)**
- **2158 - Programming the NVIDIA Digital Video Pipeline with OpenGL (Pre-Conference Tutorial) (Monday, Sept 20, 13:00)**
- **2161 - NVIDIA Quadro Digital Video Pipeline Overview (Tuesday, Sept 21, 16:00)**

Tsingos, Nicolas

Senior Staff Engineer (Dolby Laboratories)

Nicolas Tsingos leads interactive audio research at Dolby Laboratories. He is actively involved in the development of next-generation features for Axon, Dolby's in-game voice solution as well as spatial audio solutions for consumer and cinema applications.

- **Session(s): 2042 - Interactive 3D Audio Rendering Systems (Thursday, Sept 23, 11:00)**

Tzeng, Stanley

Graduate Student (University of California, Davis)

Stanley is a 3rd year PhD student at the University of California, Davis working with Professor John Owens. Over the summer he is working at Microsoft Research in Redmond with the Extreme Computing Group. His research involves GPGPU algorithms, task parallel scheduling and alternative rendering pipelines on the GPU. In his free time he loves to go explore different places to eat and go on puzzle hunts.

- **Session(s): 2162 - Real-time Reyes: Programmable Rendering on Graphics Processors (Wednesday, Sept 22, 17:00)**

Uzzan, Bruno

CEO & Founder (Total Immersion)

Bruno 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.

- **Session(s): 4011 - Emerging Companies: CEO on Stage featuring Cinnafilm, Inc., Perceptive Pixel, and Total Immersion (Thursday, Sept 23, 16:00)**

Valich, Theo

President (Bright Side Network Inc)

Theo Valich founded Bright Side Network and its subsidiaries, developing products such as next-generation automotive user interface, CPU/GPGPU computational benchmark, proprietary web engine and providing high-end 4K resolution videos in fully digital video production studio, utilizing latest GPU technology developments. Prior to founding Bright Side Network, Valich served as GPGPU technology analyst at JPR, CTO at Provox, Senior Editor at TG Daily and Tom's Hardware, as well as Contributing Editor on The Inquirer.

- **Session(s): 2303 - Using Tegra to Solve The Electric Car Power Dilemma (Tuesday, Sept 21, 14:00)**
- **2304 - Harnessing the GPU to Accelerate Automotive Development (Tuesday, Sept 21, 17:00)**

Vandermersch, Philippe

Senior Software Engineer (NVIDIA)

Philippe Vandermersch joined the CUDA Platform group in 2009, leading the development of the CUBLAS and CUSPARSE Libraries. Before that, Philippe was a Senior Video architect, working on the NVIDIA Multi-Standard Decoder solution (MSDEC). Prior to joining NVIDIA, Philippe has worked as an Embedded engineer at Equator Technologies and as a DSP engineer at Siemens ICN. He is the inventor of two patents, and holds a Master degree from the Institut National des Telecommunications in France.

- **Session(s): 2216 - CUDA Libraries Open House (Wednesday, Sept 22, 11:00)**

Varah, Sean

CEO (MotionDSP Inc.)

Dr. Sean Varah is the CEO of MotionDSP, having founded the company in 2005. Previous to MotionDSP, he co-founded Q Media Partners, and served as Director of Consumer Technology Investments at Sony Music's 550 Digital Media Ventures. He also 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.

- **Session(s): 2027 - GPU-Based Image Processing in Military Applications (Thursday, Sept 23, 09:00)**

Varhol, Peter

HPC Editor (Desktop Engineering Magazine)

Peter Varhol is an industry veteran with over twenty years experience as a technology journalist, software developer, product manager, and university processor. Currently he is HPC editor for Desktop Engineering magazine, and also principal at industry consulting firm Technology Strategy Research. He has graduate degrees in computer science, applied mathematics, and psychology.

- **Session(s): 2130 - GPU Computing and a Revolution in Design Engineering (Tuesday, Sept 21, 11:00)**

Varshney, Amitabh

Professor (University of Maryland)

Amitabh Varshney is a Professor of Computer Science at the University of Maryland at College Park where he directs the NVIDIA CUDA Center of Excellence. His interests include GPU-based heterogeneous parallel computing for computational biology, nano assembly, plasma physics, climate modeling and several other applications. Varshney received a NSF CAREER Award in 1995 and the IEEE Visualization Technical Achievement Award in 2004. He is a Fellow of IEEE.

- **Session(s): 2263 - CUDA Centers of Excellence Super-Session II (Tuesday, Sept 21, 15:00)**

Venkataraman, Shalini
Applied Engineer (NVIDIA)

Shalini Venkataraman is an applied engineer withNVIDIA'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.

Session(s): 2009 - 4D Visualization and Analysis of Flow (Tuesday, Sept 21, 17:00)

Vermes, Domokos
Associate Professor (Worcester Polytechnic Insitute)

Domokos Vermes received his MS in Electrical Engineering at the Technische Universität Dresden. He then earned his Ph.D. in Mathematics at the University of Szeged and Doctorate in Mathematics at the Hungarian Academy of Sciences. Currently, he is the Associate Professor of Mathematics and the Founding Director of Financial Mathematics Graduate Program and Laboratory at the Worcester Polytechnic Institute. In the past, he attended the University of Washington and Brown University. He specializes in: optimization under uncertainty, optimal control of Stochastic Proc.,computer assisted medical decision making, quantitative finance, portfolio optimization and risk management, and high-performance data analysis.

Session(s): 2111 - Using R for High-Performance Data Analysis (Tuesday, Sept 21, 16:30)

Vetter, Jeffrey
Professor / Distinguished R&D Staff Member (Georgia Tech / Oak Ridge National Lab)

Jeffrey Vetter, Ph.D., has a joint appointment between Oak Ridge National Laboratory (ORNL) and the Georgia Institute of Technology (GT). At ORNL, Vetter is a Distinguished R&D Staff Member, and the founding group leader of the Future Technologies Group in the Computer Science and Mathematics Division. At GT, Vetter is a Joint Professor in the Computational Science and Engineering School, where he serves as the Principal Investigator of the NSF Track 2D Experimental Computing Facility, named Keeneland, for large scale heterogeneous computing using graphics processors, and of the NVIDIA CUDA Center of Excellence.

Session(s): 2262 - CUDA Centers of Excellence Super-Session I (Tuesday, Sept 21, 14:00)

Vidal, Antonio M.
Professor (Universidad Politecnica de Valencia)

Antonio M. Vidal receives his Ph.D. degree in Computer Science in 1990, from the Universidad Politecnica de Valencia, Spain, where he is currently a full professor. He coordinates the project "High Performance Computing on Current Architectures for Problems of Multiple Signal Processing", financed by the Generalitat Valenciana in the frame of PROMETEO Program for research groups of excellence. His main areas of interest include parallel computing with applications in numerical linear algebra and signal processing.

Session(s): 2116 - Real-time Multichannel Audio Convolution (Thursday, Sept 23, 10:00)

Vo, Huy
Research Assistant (University of Utah)

Huy T. Vo is a PhD student at the SCI Institute, University of Utah, under the supervision of Professor Claudio T. Silva. His main research interests are in High Performance Computing and Visualization Systems. Huy is currently working on HyperFlow, a parallel streaming framework for large-scale visualization where data-flows can be executed efficiently on clusters of machines with multiple CPUs and GPUs.

Session(s): 2248 - Parallel Processing on GPUs at the University of Utah (Wednesday, Sept 22, 14:00)

Volkov, Vasily
PhD Student (UC Berkeley)

Vasily Volkov has contributed to substantial performance improvements in CUBLAS and CUFFT and has received NVIDIA Graduate Fellowship in 2008. He is currently a Ph.D. candidate at UC Berkeley.

Session(s): 2238 - Better Performance at Lower Occupancy (Wednesday, Sept 22, 15:00)

Vuik, Kees
Professor (Delft University Of Technology)

Session(s): 2049 - Deflated Preconditioned Conjugate Gradient on the GPU (Wednesday, Sept 22, 14:30)

Vukicevic, Vladimir
Principal Engineer (Mozilla Corporation)

Vladimir is a principal engineer at Mozilla, where he works on core browser technology. He is involved in adding new capabilities to the web platform for use by both web content and the Firefox browser, and focuses on improving the rich media and graphics capabilities of the web. His early experiments with 3D in HTML canvas led to the WebGL standard.

Session(s): 2113 - WebGL: Bringing 3D to the Web (Tuesday, Sept 21, 15:00)

Wade, Will
Business Alliance Manager (HP)

Will Wade is the Business Alliance Manager for Hewlett-Packard Company's Workstation Global Business Unit. His responsibilities include working with Intel, AMD, and NVIDIA to develop programs and solutions that benefit workstation users. Will joined Hewlett-Packard in 1997 as an engineer in Test & Measurement. In 1999 he transitioned to the Environmental Test Center for the workstation business and later moved in to a technical marketing role working with software partners for workstation applications. He became a Workstation Product Manager in 2003 where he led the workstation graphics strategy, and later the entry workstation business.

Session(s): 2233 - Solving Your GPU Computing Needs (Sponsored by HP) (Tuesday, Sept 21, 14:00)

Walker, Ross
Research Professor (San Diego Supercomputer Center)

Ross Walker is a Research Professor at the San Diego Supercomputer Center, an Adjunct Professor in the Chemistry and Biochemistry at UCSA and an NVIDIA CUDA Fellow. He runs the Walker Molecular Dynamics (MD) Lab leading a team that develops advanced techniques for MD Simulations supporting simulations improving drug and biocatalyst design. His work includes improved Quantum Mechanical, Molecular Mechanical models and the development of a GPU accelerated version of the AMBER Molecular Dynamics engine

PMEMD.

- **Session(s): 2269 - Bringing GPUs to Mainstream Molecular Dynamics Packages (Thursday, Sept 23, 10:00)**

Wang, Long
Associate Professor (Super Computing Center, Institute of Computer Network Information of CAS)

Dr. Long Wang, he got his Ph.D in Computational Mathematics from AMSS, CAS(Chinese Academy of Sciences), in 2004. Then, he went to department of scientific & engineering computing of Peking University for postdoc from 2004 to 2006. Now, he is associate professor in Super Computing Center, Computer Network Information Center of CAS. His research interests include AMR algorithm, large scale GPU computing and high performance computing software. He implemented parallel galatic wind code using 8192 cores on DeepComp 7000 supercomputing machine in 2008. This summer, he held the first international GPU workshop in Harbin, China (See: gpu-smp.sccas.cn).

- **Session(s): 2286 - Towards Peta-Scale Green Computation - Applications of the GPU Supercomputers in the Chinese Academy of Sciences (CAS) (Wednesday, Sept 22, 11:00)**

Wang, Peng
Developer Technology Engineer (NVIDIA)

Peng Wang is a member of the Developer Technology group at NVIDIA, where he develops algorithms for GPU computing. Dr. Wang received his Ph.D. in Computational Physics from Stanford University, where his primary research was the development of multi-physics codes for computational fluid dynamical simulations of astrophysical turbulence. Dr. Wang's education also includes a M.S. in Physics and a B.S. in Scientific Computing from Nankai University, Tianjin, P.R. China.

- **Session(s): 2008 - OpenCL Optimization (Thursday, Sept 23, 14:00)**
- **2006 - Short-Range Molecular Dynamics on GPU (Wednesday, Sept 22, 14:00)**

Wang, Xiaowei
(Institute of Process Engineering)

- **Session(s): 2286 - Towards Peta-Scale Green Computation - Applications of the GPU Supercomputers in the Chinese Academy of Sciences (CAS) (Wednesday, Sept 22, 11:00)**

Wang, Z.J.
Professor (Iowa State University)

Z.J. Wang, Professor of Aerospace Engineering and Director of Computational Fluid Dynamics (CFD) Center at the Iowa State University (ISU), received his Ph.D. in Aerospace Engineering from the University of Glasgow in 1990. His research areas include adaptive high-order methods for the Navier-Stokes equations, algorithm and flow solver development for structured and unstructured, overset and adaptive Cartesian grids, computational aeroacoustics and electromagnetics, parallel computing on CPUs and GPUs, geometry modeling and grid generation.

- **Session(s): 2292 - Implementation of High-Order Adaptive CFD Methods on GPUs (Thursday, Sept 23, 10:30)**

Warburton, Timothy
Associate Professor (Rice University)

Tim Warburton specializes in devising new algorithms for solving partial differential equations. He is a leader in the development of discontinuous Galerkin methods for electromagnetics simulations. He co-authored the first major text on discontinuous Galerkin methods, published by Springer in 2008. He is currently an associate professor in the department of Computational and Applied Mathematics at Rice University.

- **Session(s): 2078 - Shockingly fast and accurate CFD simulations (Wednesday, Sept 22, 11:00)**

Warren, Stephen
Snr Linux Software Engineer (NVIDIA)

Stephen Warren is a Software Engineer at NVIDIA, working on VDPAU (Video Decode and Presentation API for Unix) and related portions of the Linux graphics driver.

- **Session(s): 2016 - VDPAU: PureVideo on Unix (Thursday, Sept 23, 15:00)**

Washbrook, Andy
Postdoctoral Research Assistant (University of Edinburgh)

Dr. Andrew Washbrook is a physicist programmer based at the University of Edinburgh working for the GridPP collaboration. His previous physics research investigated evidence for supersymmetric particle production at CERN and he has also been a technical account manager for a leading enterprise open source software company. His current research interests include investigating emerging computing methods that can be used to improve the future software framework of the ATLAS experiment.

- **Session(s): 2135 - Processing Petabytes per Second with the ATLAS experiment at the Large Hadron Collider at CERN (Wednesday, Sept 22, 16:00)**

Weber, Jason
Internet Explorer Performance Lead (Microsoft)

Jason Weber is an engineering lead on the Microsoft Internet Explorer team. Jason is focused on ensuring that Internet Explorer 9 is ready for the performance demands of HTML5 applications, including hardware accelerated graphics and compiled javascript. Jason has been with Microsoft for thirteen years. Before joining the Internet Explorer team in 2008, Jason worked on projects ranging from Microsoft Office to Visual Studio, and was a member of Chairman Bill Gates technical staff.

- **Session(s): 2274 - Harnessing the Power of the GPU in Internet Explorer 9 (Tuesday, Sept 21, 16:00)**

Weiskopf, Paul
Sr. Vice President, Corporate Development (Adobe)

As senior vice president of corporate development, Paul Weiskopf leads activities related to Adobe's strategic planning, alliances, mergers and acquisitions, venture investments and new business initiatives. Weiskopf has held this role since May 2008, when he was promoted from his former role as Adobe's vice president of corporate development. Since joining Adobe in 2005, Weiskopf has managed Adobe's activities in the areas of corporate strategy, mergers and acquisitions, and venture investing, helping Adobe's product businesses identify market expansion opportunities and evaluate options for building, buying and partnering to deliver the most complete solutions to customers. He played an instrumental role in the $1.8 billion acquisition of Omniture, Inc., the $3.4 billion acquisition of Macromedia, Inc., and a number of smaller acquisitions and venture investments.

- **Session(s): 4010 - Emerging Companies: CEO on Stage featuring NaturalMotion Ltd, OptiTex,**

and Useful Progress (Thursday, Sept 23, 15:00)

- 4011 - Emerging Companies: CEO on Stage featuring Cinnafilm, Inc., Perceptive Pixel, and Total Immersion (Thursday, Sept 23, 16:00)

Whitehead, Nathan
CUDA Software Engineer (NVIDIA)

Nathan Whitehead works on the CUDA Platform team. He holds a PhD in Computer Science from the University of California, Santa Cruz.

- Session(s): 2216 - CUDA Libraries Open House (Wednesday, Sept 22, 11:00)

Williams, David M.
PhD Candidate (Stanford University)

Mr. Williams is a Ph. D. candidate in the Aero/Astro department at Stanford University under the advisement of Professor Antony Jameson. In particular, Mr. Williams is interested in developing efficient higher-order solvers capable of handling real world applications. He focuses on characterizing fluid flow over complex geometries under viscous, unsteady, and compressible conditions.

- Session(s): 2079 - A Fast, Scalable High-Order Unstructured Compressible Flow Solver (Tuesday, Sept 21, 11:00)

Williams, Ian
Director PSG Applied Engineering (NVIDIA)

Ian Williams is currently the Director of Applied Engineering within NVIDIA's Professional Solutions Group, where he has worked since 2001. He holds a BSc in Engineering Science and Technology from Loughborough University (UK) as well as an MBA from Pepperdine University (USA). He is a Chartered Mechanical Engineer with the Institute of Mechanical Engineers (UK) and has been awarded 7 patents. He is also Chairman SPEC/GPC committee which is part of the Standard Performance Evaluation Corporation.

- Session(s): 2279 - Working Man's Guide to 3D Video Editing (Thursday, Sept 23, 16:00)
- 2222 - Working Man's Guide to 3D Video Editing (Tuesday, Sept 21, 14:00)

Wilson, Adam
Postdoctoral Fellow (University of Cincinnati)

Adam Wilson received his Ph.D. in Biomedical Engineering at the University of Wisconsin in the field of Neuroprosthetics. His research has focused on developing brain-computer interfaces using implantable electrodes in humans, and using GPU processing to process high-bandwidth, high-channel-count data real-time. His work has been featured internationally, including NPR, CNN, and Time magazine (for being named one of the top 10 inventions of 2009), and was named to Popular Science's Brilliant 10 Class of 2009.

- Session(s): 2122 - Using GPUs for Real-Time Brain-Computer Interfaces (Wednesday, Sept 22, 15:00)

Wilton, Richard
Research Scientist (The Johns Hopkins University)

Richard Wilton obtained his BS and MD degrees from UCLA. He is interested in astronomical and genomics research computation at the multi-terabyte and petabyte scale.

- Session(s): 2115 - Modified Smith-Waterman-Gotoh Algorithm for CUDA Implementation (Thursday, Sept 23, 14:00)
- 2092 - Integrating CUDA into a Large-Scale Commercial Database Management System (Wednesday, Sept 22, 11:00)

Winarsky, Norman
VP Ventures, Licensing, and Strategic Programs (SRI International)

Norman Winarsky is SRI's Vice President of Ventures, Licensing, and Strategic Programs. As such he is responsible for creating SRI's highest value venture and license opportunities. He is the creator and founder of SRI's venture process, including venture and license incubation, seed funding, the EIR program, and Venture Capital engagement. He chairs SRI's Commercialization Board and the nVention Board, a partnership with the venture capital community that develops early-stage investment opportunities.

- Session(s): 4007 - Emerging Companies: CEO on Stage featuring Aqumin, RTT, and Scalable Display Technologies (Thursday, Sept 23, 10:00)
- 4008 - Emerging Companies: CEO on Stage featuring ICD and Universal Robotics (Thursday, Sept 23, 11:00)

Witchel, Emmett
Professor (University of Texas at Austin)

- Session(s): 2124 - Operating System Abstractions for GPU Programming (Thursday, Sept 23, 10:00)

Woolley, Cliff
CUDA Developer Technology Engineer (NVIDIA)

Cliff Woolley is a developer technology engineer at NVIDIA focused on enabling high-performance computing on GPUs. He completed his Master of Computer Science degree at the University of Virginia in 2003, where his research group was among the earliest in academia to investigate the use of GPUs for general purpose computing.

- Session(s): 2018 - OpenCL on the GPU (Pre-Conference Tutorial) (Monday, Sept 20, 16:00)

Wu, Ren
Senior Scientist (HP Labs)

Dr. Ren Wu is a Senior Research Scientist at HP Labs, Palo Alto. His research interests include data-intensive high-performance computing, massively parallel algorithms and computational intelligence. In recent years he has been focusing on GPU acceleration of large scale analytics, and is well known for his work on GPU-accelerated clustering algorithms.

- Session(s): 2069 - GPU-Accelerated Business Intelligence Analytics (Wednesday, Sept 22, 16:00)

Wu, Xing
Research Assistant (North Carolina State University)

Yongpeng Zhang is a graduate student in Computer Science at the North Carolina State University. His research interests are data-intensive programming models, cloud computing and GPGPU. He received his Bachelor's and Master's degree in Beihang University and Drexel University, both in Electrical Engineering.

- Session(s): 2272 - GStream: A General-Purpose Data Streaming Framework on GPUs (Thursday, Sept 23, 09:00)
- 2026 - MatCloud: Accelerating Matrix Math GPU Operations with SaaS (Tuesday, Sept 21, 17:00)

Yaacovi, Yoram
CTO and General Manager, Technologies (Microsoft Israel, R&D Center)

Yoram is the CTO and General Manager of Technologies of the Microsoft Israel Development Center (ILDC). His responsibilities include the ILDC Innovation labs – a greenhouse for breakthrough research and development as well as new technologies and business groups,

academia and industry outreach and the technology connection with the Microsoft corporate headquarters. Yoram started his career in 1983 at Elbit Computers where he programmed an innovative real time data entry and display system for the Air Force. In 1984 he joined Intel where he worked for 9 years in engineering, consulting and training roles including work with leading companies in Israel and leading a Unix and X-Windows development team at Intel in Santa Clara.

- **Session(s): 4003 - Emerging Companies Summit Panel: GPU for Computer Vision (Wednesday, Sept 22, 15:00)**

Yalamanchili, Sudhakar
Professor (Georgia Institute of Technology)

Sudhakar Yalamanchili earned his PhD degree in Electrical and Computer Engineering in 1984 from the University of Texas at Austin. After several years in industry, he joined the faculty of the School of Electrical and Computer Engineering at the Georgia Institute of Technology where he is a Joseph M. Pettit Professor of Computer Engineering. His research interests are in productivity tools for heterogeneous architectures/systems and modeling and simulation of high performance many core architectures focusing on performance, energy, and thermal characterization.

- **Session(s): 2210 - GPU-Ocelot: An Open Source Debugging and Compilation Framework for CUDA (Thursday, Sept 23, 14:00)**

Yang, Yi
PhD Student (North Carolina State University)

Yi Yang is the third year Ph.D student in Department of Electrical and Computer Engineering at North Carolina State University. He received master degree from Chinese Academy of Sciences in 2005 and Bachelor degree from University of Science and Technology of China in 2002. His research interests include High-performance computing, General Purpose Computation on Graphics Processors, code generation and optimization.

- **Session(s): 2067 - Experiences with Code Optimizations for High Performance GPGPU Programs (Tuesday, Sept 21, 16:00)**

Young, Eric
Manager of Developer Technology Profesional and Consumer Applications (NVIDIA)

Eric Young manages the developer technology group responsible for professional and consumer developers. He has graduated from Cornell University in 1995 with a Master of Computer Engineering and University of Michigan in 1994 with a Bachelors in Electrical Engineering.

- **Session(s): 2260 - DirectCompute (Pre-Conference Tutorial) (Monday, Sept 20, 14:30)**

Young, Paul
(Adobe)

- **Session(s): 2224 - GPU Acceleration in Adobe Creative Tools (Tuesday, Sept 21, 15:00)**

Zanella, Fabrizio
Systems Manager (CST of America)

Fabrizio Zanella has over 15 years of experience working on Signal Integrity characterization of high speed digital systems. Mr. Zanella has worked for several companies, including Teradyne and EMC Corporation. Currently, he is the Systems Manager at CST of America, a worldwide provider of full wave electromagnetic software. In this role, he leads the high performance computing effort, advising customers on improving overall peformance of CST tools.

- **Session(s): 2133 - 3D Full Wave EM Simulations Accelerated by GPU Computing (Thursday, Sept 23, 16:00)**

Zaspel, Peter
(University of Bonn)

Peter Zaspelis a research assistant at the Institute for Numerical Simulation of the University of Bonn, Germany. He studied Computer Science and is now working on his PhD. His research topics are computational fluid dynamics, general-purpose computations on graphics hardware, and visualization.

- **Session(s): 2083 - GPU Accelerated Solver for the 3D Two-phase Incompressible Navier-Stokes Equations (Wednesday, Sept 22, 16:00)**

Zeitlin, Michael
CEO (Aqumin)

Michael began his career at Texaco as a Senior Scientist in 1980 and became a Texaco Fellow in 1997. He received the Carnegie Mellon and American Management Institute Award for Innovation in Information Technology in 1998. In 1999 his work was honored with a permanent position in the Archives of the Smithsonian Institution. He founded Magic Earth, LLC. in 2000 and grew operations to profitability in three months. Magic Earth was acquired by Halliburton that same year for $100 million.

- **Session(s): 4007 - Emerging Companies: CEO on Stage featuring Aqumin, RTT, and Scalable Display Technologies (Thursday, Sept 23, 10:00)**

Zhang, Yao
Graduate Student (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.

- **Session(s): 2085 - Tridiagonal Solvers: Auto-Tuning and Optimizations (Tuesday, Sept 21, 15:00)**

Zhang, Yubo
PhD Student (UC Davis)

Yubo Zhang is a PhD student supervised by Prof. Kwan-Liu Ma at the department of Computer Science, UC Davis. His research interests include numerical methods, computer graphics and visualization.

- **Session(s): 2145 - Photo Editing on the GPU with MuseMage (Thursday, Sept 23, 09:00)**

Zhang, Yunquan
Professor (Institute of Software, CAS)

Prof. Yun-quan Zhang is the Associate Director of the Parallel Computing Laboratory, Institute of Software, Chinese Academy of Sciences in Beijing, China. He received his PhD degree in computer software and theory from the same institute in 2000 and has worked at the Institute as a research scientist since then. His major research interests are in the areas of high performance parallel computing, with particular emphasis on the design of large scale parallel computation modes and numerical libraries, and large system performance modeling and evaluation. He has published over 90 papers and trained over 20 master and Ph.D. students.

- **Session(s): 2286 - Towards Peta-Scale Green Computation - Applications of the GPU Supercomputers in the Chinese Academy of Sciences (CAS) (Wednesday, Sept 22, 11:00)**

Zhao, Kaiyong
Graduate Student (HKBU)

I received my B.Eng. degree in the Aircraft Design and Technology from Beijing Institute of Technology (BIT), Beijing, P. R. China, in 2005. Then worked in CCUR two years. I am currently an MPhil student in the Department of Computer Science, Hong Kong Baptist University.

- **Session(s): 2145 - Photo Editing on the GPU with MuseMage (Thursday, Sept 23, 09:00)**

Zhou, Huiyang
Associate Professor (North Carolina State University)

Huiyang Zhou is currently an associate professor in the Department of Electrical and Computer Engineering at North Carolina State University. His research focuses on high performance microarchitecture, low-power design, architecture support for system dependability, and GPU Computing. He is a recipient of NSF CAREER award and a senior member of the IEEE.

- **Session(s): 2067 - Experiences with Code Optimizations for High Performance GPGPU Programs (Tuesday, Sept 21, 16:00)**

Ziegler, Gernot
Developer Technology (Compute) (NVIDIA)

Gernot Ziegler (MSc/civ.ing.) is an Austrian engineer with an MSc degree in Computer Science and Engineering from Linköping University, Sweden. He pursued his PhD studies at the Max-Planck-Institute for Informatics in Saarbrücken, Germany, where he specialized in GPU algorithms for computer vision and data-parallel algorithms for spatial data structures. As a member of NVIDIA's DevTech-Compute team, Gernot now consults in high performance computing on graphics hardware.

- **Session(s): 2020 - GPU-Accelerated Data Expansion for the Marching Cubes Algorithm (Wednesday, Sept 22, 16:00)**
- **2021 - Efficient Volume Segmentation on the GPU (Wednesday, Sept 22, 17:00)**

Zigon, Robert
Sr Staff Development Engineer (Beckman Coulter)

Bob is the Software Technical Lead for Flow Cytometry analysis products within Beckman Coulter. His interests include high performance computing, numerical analysis and information retrieval theory.

- **Session(s): 2055 - Application of Fermi GPU to Flow Cytometry and Cancer Detection (Thursday, Sept 23, 10:00)**

A Powerful Platform for Amazing Performance

Performance. To get it right, you need a foundry with an **Open Innovation Platform™** and process technologies that provides the flexibility to expertly choreograph your success. To get it right, you need TSMC.

Whether your designs are built on mainstream or highly advanced processes, TSMC ensures your products achieve maximum value and performance.

Product Differentiation. Increased functionality and better system performance drive product value. So you need a foundry partner who keeps your products at their innovative best. TSMC's robust platform provides the options you need to increase functionality, maximize system performance and ultimately differentiate your products.

Faster Time-to-Market. Early market entry means more product revenue. TSMC's DFM-driven design initiatives, libraries and IP programs, together with leading EDA suppliers and manufacturing data-driven PDKs, shorten your yield ramp. That gets you to market in a fraction of the time it takes your competition.

Investment Optimization. Every design is an investment. Function integration and die size reduction help drive your margins. It's simple, but not easy. We continuously improve our process technologies so you get your designs produced right the first time. Because that's what it takes to choreograph a technical and business success.

Find out how TSMC can drive your most important innovations with a powerful platform to create amazing performance. Visit **www.tsmc.com**

tsmc®

Open Innovation Platform™

EXHIBIT HALL

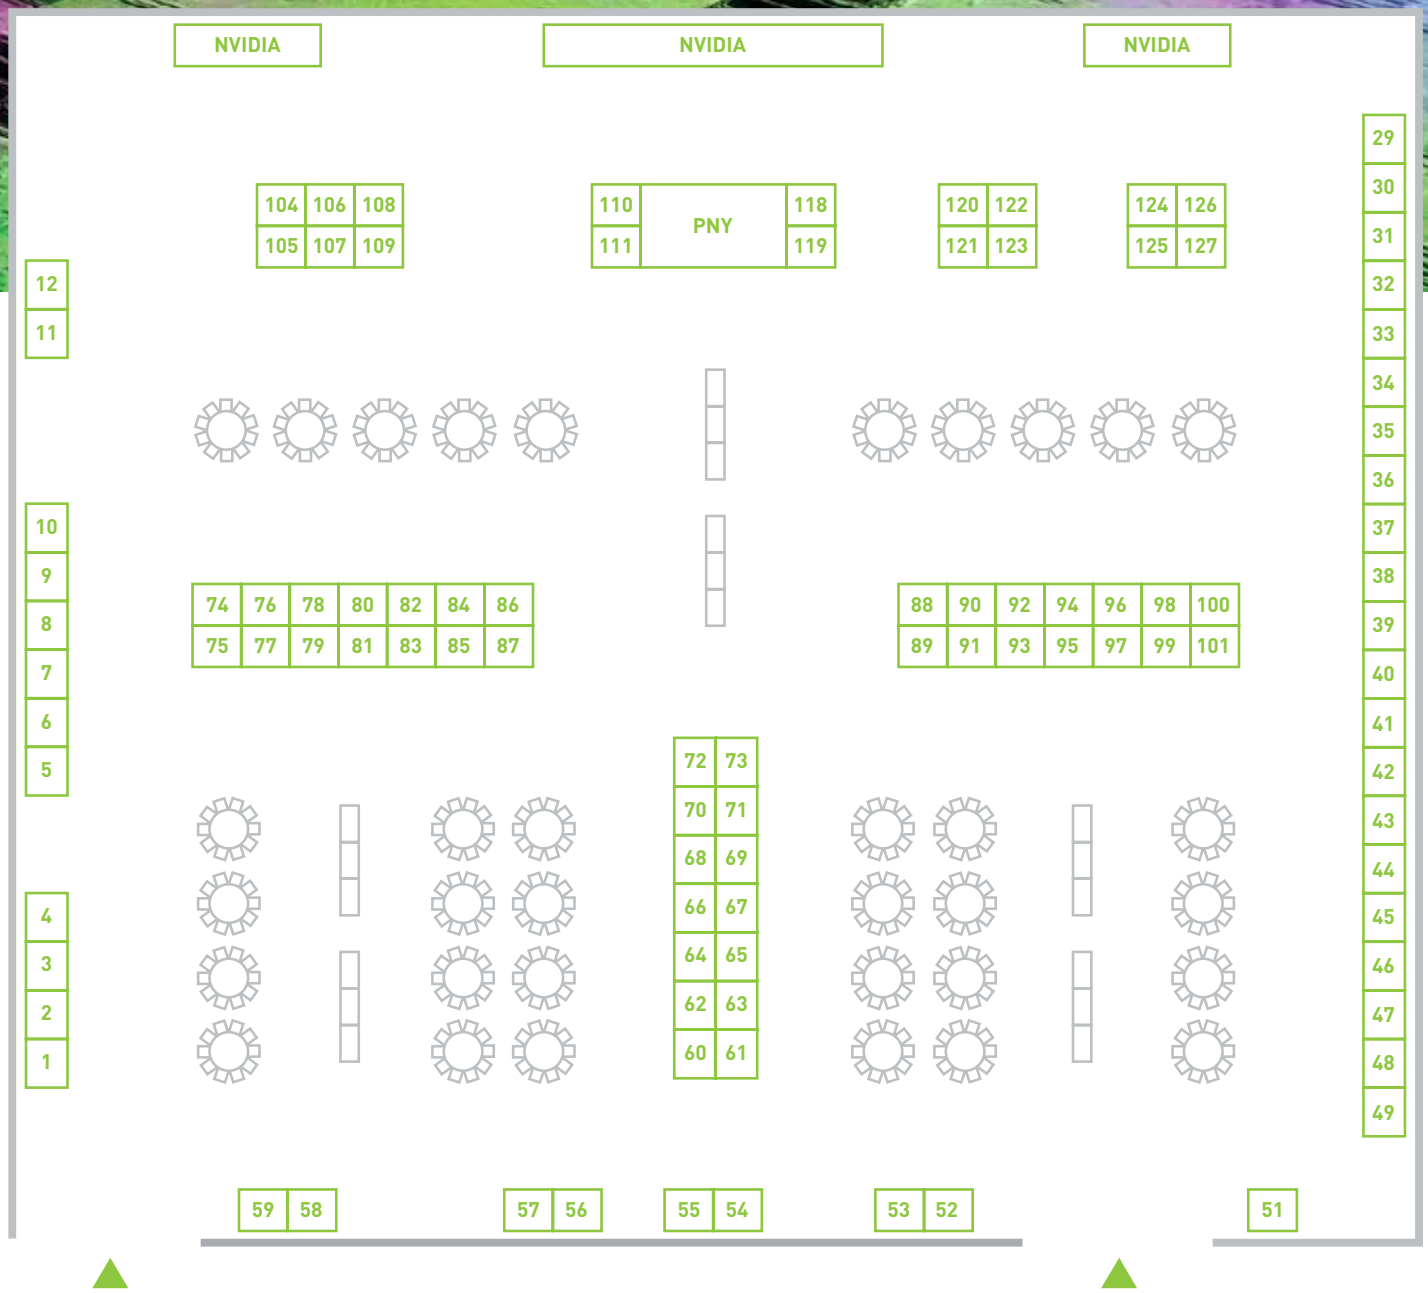

HALL 1

CONCOURSE AND RESEARCH POSTERS

EXHIBITORS	BOOTH #
3DreamTeam	8
3DTV Solutions	1
AccelerEyes	56
Acceleware	52
Ace Computers	76
Allegorithmic	6
AMAX Information Technologies	89/91
Appro International, Inc.	75
Aspen Systems	68
Binatix Inc	42
BioDigital	98
BOXX Technologies	110
Bright Computing	10
Bunkspeed	104
CIARA Technologies	55
Cinnafilm Inc	109
Cirrascale Corporation	99
CodeSourcery	38
Colfax International	51
Cooliris	37
Creative Consultants LLC	29/30
Cubix Corporation	125
CyberLink Corporation	70
Dassault Systemes	93
Dell	106/108
Discretix	123
EM Photonics	54
embodee	77
Empulse	39
Exxact Corporation	95
Filter Foundry	83
GE Fanuc Intelligent Platforms	101
Geomerics	33
Giada	65
GraphStream Incorporated	118
HP	60/62/64
HPC Project / Wild Systems	5
Hynix	79/81
IBM	127
Israel Economic Mission	121
James River Technical	57
Jedox Business Intelligence	40
JMR Elctronics, Inc	111
Koi Computers	84
Los Alamos	2

EXHIBITORS	BOOTH #
Mathworks	66
Mazda Technologies	41
Mellanox Technologies	49
mental images	74
Micoy	Concourse
Microsoft	87
Microway Inc	105/107
miGenius Limited	32
Mingleverse	82
Morgan Kaufmann	44
NaturalMotion LTD	94
NextComputing	119
NextIO	61/63
NVIDIA	Backwall
NVIDIA Foundation	126
OptiTex USA Inc	120
PEER 1 Hosting	100
Penguin Computing	43
Perceptive Pixel	92
PhaseSpace Inc	69
Phototour	45
Platform Computing	72
Playcast	122
PNY	110-119
Polywell	67
Prometech Software, Inc.	4
PSSC Labs	59
RealityFrontier	3
Reservoir Labs	86
RTT USA, Inc	9
Scalable Display Technologies	47
ScaleForm Corporation	35
SGI	85
Stonetrip	7
Softkinetic	48
Supermicro	71/73
Synnex Corp	97
Tech-X Corporation	34
The Portland Group	36
Tide Powerd Ltd.	78
T-Platform	46
Trinity Racing Concepts, LLC	11/12
Universal Robotics	88/90
Useful Progress	58
VisiSonic Corporation	80
Wolfram Research	31

PLATINUM SPONSORS

Adobe

Adobe revolutionizes how the world engages with ideas and information. Our award-winning software and technologies have set the standard for communication and collaboration for more than 25 years, bringing vital and engaging experiences to people across media and to every screen in their lives, at work and at play. The impact of Adobe® software is evident almost everywhere you look. Whether people are collaborating at work, transacting online, or socializing with friends, businesses use Adobe software and technologies to turn digital interactions into richer, high-value experiences that reach across computing platforms and devices to engage people anywhere, anytime. 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.

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.

Microsoft

Microsoft Corporation is passionate about software innovation and is making a huge investment into developing products and applications for the Technical Computing world. Modeling and simulation as well as raw computational performance is critical for scientists, engineers and analysts seeking solutions in an ever larger sea of data. Professionals seek to model an ever increasing scale of scenarios using the most advanced mathematical functions in the quickest manner possible. However, the path from math to model to results often takes too long. Microsoft has formed a Technical Computing organization to enable solutions to this challenge using the full range of compute platforms; client, cluster and cloud.

Dell

Dell's research and development (R&D) efforts span the globe, driven by some of the industry's foremost product designers and engineers. At the core of Dell's research and development (R&D) efforts span the globe, driven by some of the industry's foremost product designers and engineers. At the core of Dell's innovation approach, however, remains an unwavering commitment to deliver new and better solutions that directly address customer needs. 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. The 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. For more information on Dell's high performance solutions please visit www.dell.com/precision or www.dell.com/solutions.

Supermicro

Supermicro, the leader in server technology innovation and green computing, provides customers around the world with application-optimized server, workstation, blade, storage and GPU systems. Based on its advanced Server Building Block Solutions, Supermicro offers the most optimized selection for IT, datacenter and HPC deployments. The company's system architecture innovations include Twin server, double-sided storage and SuperBlade® product families. Offering the most comprehensive product lines in the industry, Supermicro delivers energy-efficient solutions with unmatched performance and value. Founded in 1993, Supermicro is headquartered in Silicon Valley with worldwide operations and manufacturing centers in Europe and Asia. For more information, visit www.supermicro.com.

PNY

Established in 1985, PNY Technologies®, Inc. is the authorized NVIDIA® Quadro® channel partner for North America and Europe. PNY provides unsurpassed service and commitment to its professional graphics customers offering: 3 year warranty, pre and post sales support, dedicated Quadro Field Application engineers and direct tech support hot lines. The company also offers a full line of consumer graphics cards, computer memory upgrade modules, solid state drives, flash memory cards, USB flash drives, and HDMI cables. Headquartered in Parsippany, NJ, PNY maintains facilities in North America, Europe, Asia, and Latin America. For more information, please visit http://www.pny.com.

Cooley

Cooley LLP 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 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.

Synnex Corp

SYNNEX Corporation, a Fortune 500 corporation, is a leading business process services company, servicing resellers and original equipment manufacturers in multiple regions around the world. The Company provides services in IT distribution, supply chain management, contract assembly and global business services. Founded in 1980, SYNNEX employs over 7,000 associates worldwide and operates in the United States, Canada, China, Japan, Mexico, the Philippines and the United Kingdom. Our value-added service model streamlines business processes to help customers across the globe lower their costs and create greater efficiencies. We provide a variety of professional and marketing services, including: demand generation, education and training, pre-and post-sale technical support, end-user enablement, server assessment, design and integration, recycling and trade-in, contract design and assembly, and IT resource planning.

Acer

Established in 1976, globally Acer ranks No. 2 for total PCs and notebooks. A profitable and sustainable Channel Business Model is instrumental to the company's continuing growth, while its multi-brand approach effectively integrates Acer, Gateway, Packard Bell, and eMachines brands in worldwide markets. Overcoming the barriers between people and technology: this is Acer's long-term mission, to allow anyone to use and benefit from technology. Acer is renowned for the development and manufacture of sophisticatedly, environmentally friendly and intuitively designed, easy to use products. For further information, please visit the website acer-group.com.

Gateway

The Californian company Gateway is a historical brand in the IT market, and has been a leading company in the field of computers and notebooks for nearly 25 years. Now, thanks to the creation of a new structure exclusively focused on businesses, Gateway enters Europe as a player that is fully committed to the professional market and that can develop complete and reliable business solutions for partners dedicated to medium-sized and large companies. Gateway products are available in Europe only through indirect distribution channels and, in particular, through its network of selected Partners. For further information, please visit the website gateway.com

TSMC

TSMC is the world's largest dedicated semiconductor foundry, providing the industry's leading process technology and the foundry's largest portfolio of process-proven libraries, IP, design tools and reference flows. The Company's total managed capacity in 2008 exceeded 9 million 8-inch equivalent wafers, including capacity from two advanced 12-inch - GigaFabs ™, four eight-inch fabs, one six-inch fab, as well as TSMC's wholly owned subsidiaries, WaferTech and TSMC (China), and its joint venture fab, SSMC. TSMC is the first foundry to provide 40nm production capabilities. Its corporate headquarters are in Hsinchu, Taiwan. For more information about TSMC please visit http://www.tsmc.com.

Samsung

Samsung Electronics Co., Ltd. is a global leader in semiconductor, telecommunication, and digital convergence technologies. The company consists of seven independently operated business units: Visual Display, Mobile Communications, Telecommunication Systems, Digital Appliances, IT Solutions, Semiconductor and LCD. The Samsung Semiconductor Business provides the most extensive range of Flash, DRAM, SRAM and specialty memories in the world and paces the industry in advancing memory technology for home, mobile and office environments, for consumer electronics and networked applications. For more information, please visit www.samsung.com/GreenMemory.

IBM

IBM delivers powerful, innovative solutions to customers' most challenging and complex problems, enabling businesses and researchers to innovate, achieve breakthrough results and establish sustainable competitive advantage. IBM takes a holistic approach to high performance computing that involves designing and delivering complete, robust technical solutions that are easy to acquire and access, environmentally responsible, competitively priced, and backed by IBM's world-class support.

GOLD SPONSORS

Next IO

NextIO, Inc. is the leader in next-generation network consolidation solutions for today's dynamic data center in a variety of industries including enterprise, oil and gas, High Performance Computing, digital media and financial services. Leveraging PCI Express, NextIO offers true I/O consolidation for any end-point technology

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.

SILVER SPONSORS

GE Intelligent Platforms

GE Intelligent Platforms is a leading manufacturer of rugged COTS computer boards and systems for military programs. As a partner to NVIDIA for Embedded Applications, GE-IP brings GPGPU technology into a wide range of defense related programs and can now be used in ground tanks, fighter aircraft, military helicopters, and UAV's for Radar, ISR, DSP, Sensor Processing, Imaging and many other military applications.

AMAX Information Technologies

With 31 years as a leading systems manufacturer, AMAX delivers the most cutting edge GPU and GPGPU computing solutions to solve data-intensive computing challenges in today's leading industries. Using an open-architecture approach optimizing best-of-breed components for solutions designed to fit precise needs with maximize functionality, superior performance and power efficiency—this is the AMAX advantage.

SGI

SGI is focused on helping customers solve their most demanding technology challenges by delivering: high performance computing, servers, storage, data center and cloud computing solutions and professional services. We develop, market and sell a broad line of low-cost, mid-range and high-end scale-out and scale-up servers and data storage solutions as well as differentiating software.

Appro

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.

Sutter Hill Ventures

Sutter Hill Ventures has financed technology-based start-ups and assisted entrepreneurs in building market-leading companies since 1964. Through our decades of experience, we 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.

Hynix

Hynix is a leading producer of DRAM, NAND Flash memory and Image Sensors. Besides computing memory solutions, Hynix is a leader in Graphics memory with a portfolio of high performance products.The newly introduced, eco-friendly 1.35V, 44nm 2Gb GDDR5 offers 7Gbps speed targeting high end graphics and high performance computing.

Dassault Systèmes

Dassault Systèmes is a world leader in 3D and Product Lifecycle Management (PLM) solutions, bringing value to over 100,000 customers in 80 countries. 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."

Silicon Valley Bank

Silicon Valley Bank is the premier commercial bank for companies in the technology, life science, venture capital, private equity and premium wine industries. SVB provides a comprehensive suite of financing solutions, treasury management, corporate investment and international banking services to its clients worldwide. Through its focus on specialized markets and extensive knowledge of the people and business issues driving them, Silicon Valley Bank provides a level of service and partnership that measurably impacts its clients' success. Founded in 1983 and headquartered in Santa Clara, Calif., the company serves clients around the world through 26 U.S. offices and international operations in China, India, Israel and the United Kingdom. Silicon Valley Bank is a member of global financial services firm SVB Financial Group (Nasdaq: SIVB), with SVB Analytics, SVB Capital, SVB Global and SVB Private Client Services. More information on the company can be found at www.svb.com. Silicon Valley Bank is the California bank subsidiary and the commercial banking operation of SVB Financial Group. Banking services are provided by Silicon Valley Bank, a member of the FDIC and the Federal Reserve System. SVB Private Client Services is a division of Silicon Valley Bank. SVB Financial Group is also a member of the Federal Reserve System.

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.

Barco

With more than 25 years of experience in 3D, Barco designs and develops display visualization solutions for a variety of markets, including simulation, virtual reality, medical, defense, media and entertainment, and traffic, surveillance and monitoring. A global technology company, Barco (NYSE, Euronext Brussels: BAR) is active in more than 90 countries and employs 3,300 staff worldwide.

Mandel Communications Inc.

Mandel is a provider of communication skills coaching and training to leading companies worldwide. Since 1993, Mandel has been dedicated to helping executives, sales, and technical professionals achieve improved business results through effective business conversations and presentations, particularly when the stakes are high. Mandel's proprietary approach to strengthening presentation, conversation, and facilitation skills turns spoken communications into a competitive advantage. With a global presence, Mandel delivers its solutions both face-to-face and virtually in more than 14 languages.

Emerging Companies Exhibitors

3DreamTeam

3DreamTeam LLC have created a unique library of photo realistic 3D worlds and images as well as the Vizerra software platform and ecosystem which allows their in-house team of thirty designers, software engineers, technical artists or any independent studio to assemble photo realistically rendered, 3D environments and marketing content. The Vizerra Platform and Ecosystem will drive the content boom creating demand for enterprise and consumer hardware. The unique benefits of the Vizerra Platform and Ecosystem is speed to create, cost of production and size of file Vs quality of rendering.

3DTV Solutions

3DTVSolutions™ conceives and develops cameras and software suites to create and manipulate real and virtual images in FullDepth™3D to be viewed without glasses in real time. Our unique expertise in the complete chain of 3D imaging: Production-Transmission and Display opens up huge market potential for industry, healthcare, audio-visual production, and gaming. 3DTVSolutions™ will demonstrate its capabilities in terms of capture, editing and display of real and virtual 3D images to be viewed on auto-stereoscopic screens without glasses. These images come from our unique camera system that films directly in 3D and/or computer-generated images to be mixed and edited to create films or augmented reality in FullDepth™ 3D.

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.

Acceleware

Acceleware provides High Performance Computing software solutions to the Oil & Gas and Computer Engineering markets for multi-core CPUs and compute GPUs. Additionally, our HPC consulting services are utilized by enterprises needing to harness the power of parallel computing. At Acceleware the goal is always the same – Go Faster!

Allegorithmic

Allegorithmic is the first company to propose a professional toolset for compression, authoring and on-the-fly rendering of smart textures. Allegorithmic's new product range Substance is poised to redefine the development and distribution of rich content for the next generation of online and mobile games: Substances are dynamic and are typically 500-1000 times smaller than regular textures.

Binatix

Binatix develops pattern recognition software based on machine learning algorithms that mimic the human brain.

Biodigital

BioDigital is dedicated to using cutting edge biomedical visualization systems to improve training, communication and the interpretation of medical information. From 3D animation, to virtual training environments, to systems that intuitively store and visualize scientific data, BioDigital's products and services promise to revolutionize the way we understand medical subjects.

Bunkspeed

Bunkspeed is a leading global provider of 3D rendering and animation software for the design and creative industry. Bunkspeed is a private company founded in November of 2002 with the philosophy that 3D rendering software should be easy to learn, simple to use and produce stunning photographic results. Headquartered in Carlsbad, California, Bunkspeed's products include Bunkspeed Shot™, a "virtual digital camera" with interactive ray-tracing empowering anyone to quickly create photographic-quality images, Bunkspeed Move™, a "virtual movie camera" bringing products to life by quickly creating various types of animations, and Bunkspeed Drive™, a fully featured visualization application tuned for the automotive industry. Bunkspeed's customers include Frog Design, Pininfarina, Unilever, Rubbermaid, Nike, Ford Motor Company, Honda, and Tiffany's. For more information on Bunkspeed's products and services, visit www.bunkspeed.com.

Cinnafilm, Inc.

Cinnafilm™ was founded to address the absence of quality, software-based tools to visually optimize, convert and repurpose video images in the post production market. Recognizing the power of modern Graphics Processing Units (GPUs) and the accelerating migration to file-based workflows, Cinnafilm disregarded prevailing industry methods in 2005 and started from scratch to create what has become the world's fastest, most accurate GPU-based image processing engine, Pixel Strings™. Cinnafilm is a growing company that has successfully partnered with some of the strongest names in the motion picture and television industry: ARRI, Quantel, Rhozet, and NVIDIA. Our partners have recognized the power of Pixel Strings and the superlative image quality which can result when this power is properly harnessed. Cinnafilm is a privately held company, headquartered in Albuquerque, NM, amongst powerful resources such as the nation's defense laboratories and New Mexico's highly competitive film tax incentives.

Code Sourcery

CodeSourcery builds software tools that enable its customers to get the most out of hardware platforms ranging from embedded devices to supercomputers: Sourcery G++, tools for professional embedded C and C++ developers, and Sourcery VSIPL++, a C++ library for developing high performance signal- and image-processing applications.

Cooliris, Inc

Cooliris 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. Each of us is passionate about serving our users without compromise and seeing that our products deliver the best experience possible. Headquartered in Palo Alto, CA, Cooliris is backed by Kleiner Perkins Caufield & Byers, DAG Ventures, the Westly Group, and T-Venture. For more information, please visit www.cooliris.com/company.

Cyberlink

CyberLink Corp. is the leader in enabling digital multimedia on PCs and CEs. Our software solutions include 3D stereoscopic applications, Blu-ray Disc playback and creation, digital home entertainment, and touch-enabled media solutions. CyberLink's partners with worldwide business leaders in the PC industry include top-5 desktop and notebook brands, drive manufacturers, and graphics card markers. CyberLink offers the most extensive array of 3D software for the consumer market, from Blu-ray™ 3D to 3D conversion of 2D video and even 3D slideshow applications. Hardware support from NVIDIA® 3D Vision™ and optimization for NVIDIA® CUDA™ decoding and encoding technology results in a superb viewing experience with reduced loading on PC system resources.

Discretix

Security and content protection lie at the heart of the mobile and home entertainment markets. Discretix' suite of embedded security solutions includes security co-processors, security sub-systems, cryptographic cores and content protection applications. The Discretix Downloadable Secure Player allows content service providers (CSP) to target the large amount of devices already in the market, overcoming the dependency on pre-installed or embedded device applications and accelerating the deployment of new services. This unique secure player supports both industry standard and proprietary content protection schemes and is compliant with the requirements of the content owners. Suitable to a wide variety of connected devices including smartphones, e-book readers, tablet computers, netbooks and internet-enabled TVs, the secure player enables a broad range of premium content services and applications. The secure player is available on various open operating systems including iPhone OS, Android, Windows Mobile / Windows Phone, Linux and Symbian. Discretix' content protection solutions are field proven and trusted by some of the world's best-known semiconductor and device manufacturers including Intel, HTC, SonyEricsson, Acer and Motorola. For more information with visit www.discretix.com.

embodee

embodee™ gives you an Online Try-onSM service with digital versions of real-world jeans, enabling you to quickly find a great looking pair. Garment sizes are automatically matched with your body, making it easy to select a size. Embodee's GPU equipped servers run highly accurate simulations of look and fit - of any jeans style, in any size, on your body – and we deliver them on-demand to any device. www.embodee.com

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. Our most popular commercially available software product is CULA, a GPU-accelerated linear algebra library built in collaboration with NVIDIA.

Empulse GmbH

Empulse is an IT service company focussing on the development of CUDA applications for High Performance Computing. Based in Cologne/Germany, a team of highly professional developers and architects designs and implements high quality solutions. The main product of empulse is Parstream, the first database engine running on GPUs while using optimized indexing algorithms. Parstream is the optimal solution for filtering and analyzing very large structured datasets with billions of records, like web-logs, financial transaction histories, call data records etc. It needs a fraction of time compared to conventional solutions.

Filter Foundry

Filter Foundry is the first social network marketplace for the worldwide audience of designers and visual artists, enabling "members" to make money and enhance career, daily. Individuals and creative companies build and manage a profile and creative work, market talents, get work, hire and manage projects, search and answer vital questions, teach and learn, and buy digital products.

Geomerics

Geomerics is an innovation-led company specialising in graphics software. We are built on a combination of advanced in-house IP, a world-class research team, and strong management experience. Geomerics' first product is Enlighten. Enlighten redefines the way lighting is handled in computer games. Instead of pre-baking the effects of global illumination into the scene they are computed at run time, allowing for fully dynamic lighting that dramatically enhances quality. Enlighten gives artists total control over lighting, driving a new generation of games that rival film for their manipulation of mood and atmosphere. Licensees of Enlighten include AAA titles in production at EA DICE, CCP and FunCom.

HPC Projects/ Wild Systems

HPC Project is a high-tech company whose mission is to supply combined hardware and software solutions to users requesting high performance computing for their applications. Under the brand name Wild Systems, HPC Project provides dedicated appliances. Using the open source software Par4All for translating standard C code in a CUDA-capable code, the appliances take full advantage of the NVidia GPU technology. One of Wild Systems turnkey appliances is the WildLab for the simulation community. With the WildLab, users creating models in the Scilab script language seamlessly produce a GPU accelerated autonomous executable code.

Israel Economic Mission (IEM)

"The Government of Israel Economic Mission (IEM) is responsible for enhancing bi-national trade relations between the West Coast and Israeli business communities. By leveraging its networking capacity and industry knowledge in Israel and the Western United States, IEM is able to seamlessly engage prospective business partners half a world apart. Such actions manifest an array of high-level connections ranging from brokering introductions to organizing international trade missions. Our operations span a variety of industry sectors, with a focus on high-tech, security, new media, cleantech, and biotechnology."

Jedox Business Intelligence

Jedox develops a centralized, in-memory, GPU-based calculation engine that controls and stores the spreadsheet-based business Intelligence data contained in every Excel, Open Office and Google spreadsheet in an organization. This technology stops "Spreadsheet Spreadmart Chaos" (hundreds of spreadsheets with uncoupled, non-verifiable data "running amok" in an organization).

Mersive Technologies Inc

Mersive Technologies' display management software and solutions bring unprecedented simplicity and affordability to large-scale, beyond-HD displays allowing visual collaboration to go mainstream. Mersive's patented SOL software automatically aligns multiple projectors into one seamless image of extraordinary quality and resolution without the expense of specialized hardware, building infrastructure or services.

miGenius Limited

Based on the powerful combination of NVIDIA CUDA based GPU's and mental images 'iray' rendering technology, miGenius has developed an easy to use toolset, EasyRS, allowing non-technical users to create fully customized User Interfaces with a wide range of viewing and management controls that can be quickly uploaded onto either a dedicated GPU server or utilizing the rapidly emerging 'cloud computing' networks. With these toolsets, the vast market of both businesses and consumers alike will be able to rapidly transform these powerful web based rendering platform

technologies into a customisable and truly revolutionary 3D visual communication and collaboration medium.

Milabra

Milabra uses proprietary, patent pending machine vision software to create visual data about the web. We bridge the visual relationship between the Ad, the Page, and the Audience for increased advertiser performance. We analyze the visual attributes and content of webpages in order to target online display advertising, optimizing creative choice based on the visual environment within which it will appear. We analyze page elements as well as the photo and video content, providing real time data for targeting and optimization as well as defending advertisers against negative content associations.

Mingleverse

Revolutionizing the way people meet and interact, Mingleverse brings to the market an entirely new "Virtual Telepresence" communications service allowing you to talk and share media live within a browser based virtual "Mingle Room". It is simply the new, easy, and fun way to talk and meet online!

NaturalMotion Ltd

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.

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.

Perceptive Pixel

Perceptive Pixel is dedicated to the research, development and deployment of multi-touch interfaces for the knowledge worker. The company's hardware and software solutions enable both novice and expert users to manipulate complex datasets through a new class of intuitive yet powerful and visually rich interface techniques.

PhaseSpace

PhaseSpace uses Nvidia GPUs with custom 4 megapixel cameras for pose tracking, eye tracking, 3D scanning and computer vision tasks in real time. PhaseSpace's list of clients includes the U.S. Navy, Air Force, Army, NASA, Boeing, Disney, Honda, Google, Stanford, Berkeley, MIT, Cambridge, Sandia Labs, Sony, MTV, and others.

Phototour

Phototour is the "social Google streetview", created by crowd sourcing geotagged photographs and panoramas. With 40% of the revenue generated online relating to travel, exploring and discovering new places still remains a challenge. Phototour conquers this challenge by providing a convenient, intuitive, and reliable way to browse geospecific photos. Users also have the freedom to create their own virtual tour using photos or panoramas created by our patented web based photostitcher without having to download any software.

Playcast Media Systems

Playcast Media Systems brings video games to the world's largest media distribution platform – Pay TV networks. The company's solution includes a head-end based system, which streams a game's audiovisual content as a standard MPEG stream, as well as the provisioning of the content and programming itself. Playcast's media streaming systems, located in operators' headends, host the games and stream them over the existing video network to an already distributed base of set-top boxes. Playcast is a privately owned, venture capital backed company, based in Israel and the UK.

Prometech

Prometech Software, Inc. provides a particle-based CAE software "Particleworks" for Japanese manufacturing industries. Prometech, as an university-launched technology venture, has a strong technical capability of complex physics simulation, visualization and many-core acceleration and offers physics based simulation in the fields of manufacturing, VFX and scientific researches.

RealityFrontier

RealityFrontier, an affiliate of CoroWare Inc. and its partners, is dedicated to delivering services and sales enablement solutions in three converging technical fields of expertise: hardware acceleration, connected systems and telepresence devices. RealityFrontier's market is with service providers and designers. RealityFrontier long-term goal is to become an ISV in augmented reality, defined as a virtual interaction layer between the physical world and the user. To pursue that goal, RealityFrontier capitalizes on its partners' know-how in video streaming, embedded systems, user interaction, data manipulation and high-performance computation. RealityFrontier is actively reaching out to sponsors and partners.

Reservoir Labs

Reservoir Labs has developed compiler algorithms for automatically performing the resource tradeoffs (parallelism vs. locality vs. contiguity) needed to automatically translate sequential loop nests in C into CUDA. This can multiply the productivity of users bringing new algorithms to NVIDIA GPUs reaching high percentage of peak performance.

RTT

Realtime Technology (RTT) AG stands for creative and fascinating 3D visualization solutions, which bring products to life in realtime and portray them in a natural and realistic environment. The company provides its clients with assistance during each stage of the life cycle of their products – from the initial product design stage through to development and subsequent marketing and sales. The 3D data model from the product development stage serves as the basis for all the following steps in the product lifecycle. It can be used, for example, to rapidly create computer generated, photorealistic product illustrations for the marketing department or to develop a 3D online product configurator on a website. In this way, RTT doesn't just speed up decision making and development processes for its clients, but it also opens up new opportunities with regard to marketing and sales. The company was founded in 1999 and its head office is in Munich, Germany. RTT AG has over 400 employees and is represented in 14 locations worldwide. Many leading businesses have put their trust in RTT and its portfolio of clients includes names such as Adidas, Audi, BASF, BMW, Bosch, Daimler, EADS, Harley-Davidson, Miele, Porsche, Samsung, Thyssen-Krupp, Toyota and Volkswagen. RTT AG is a stock market listed company (Xetra:R1T; WKN: 701220; ISIN: DE0007012205). For more information visit www.rtt.ag.

Scalable Display Technologies

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, multiprojector 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.

ScaleForm Corporation

Scaleform is the leading provider of user interface software for the videogame and consumer electronic industries. Scaleform GFx leverages the power of the Adobe® Flash® tool set and enables developers to quickly create powerful and immersive user interface environments while stream lining workflow and improving time to market.

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.

Stonetrip

Stonetrip a leading 3D engine company for games and 3D applications. Headquartered in France, the company designs and supports ShiVa, powerful easy-to-use tools for creating amazing 3D real-time applications and games. Stonetrip continues to add additional platforms as it extends its reach to new markets, maintaining its position as the most cross-platform-compatible solution in the market today.

Softkinetic

Softkinetic is the leader in natural interfaces that transform the way people interact with the digital world. We provide the most advanced software platform for building immersive, transparent and intuitive user experiences within the fields of Interactive Digital Entertainment, Serious Games, Interactive Marketing and Consumer Electronics. More information: www.softkinetic.net

Tide Powerd Ltd.

Founded in 2009 at the University of Alabama, TidePowerd Ltd. was the first-ever participant accepted into Red Gate Software's (http://www.red-gate.com) "Springboard" incubator program. Our core goal is to provide developers with powerful, yet easy-to-use tools for numerical and high-performance computing (HPC) - and to provide those tools with the best value and technical support possible. To this end, we've created GPU.NET, a system that allows developers to write their GPU code in any .NET-supported language (e.g., C#, F#, IronPython). GPU.NET opens up the exciting world of GPU computing to millions of new developers worldwide, and we hope it will help to make GPU computing more popular than ever before.

Trinity Racing

Trinity Racing Concepts produces high quality motorsports simulation products and services. Trinity is the leader in the integration of Stereoscopic 3D into hyper-realistic control systems to develop the most accurate and immersive motorsports simulation experience available. In addition to creating custom and retail simulators, Trinity also develops and manages race-themed promotional tours, trade shows, company parties, team-building activities, and other events.

Universal Robotics

Universal Robotics creates software that enables machines to learn from their experiences, react and adapt to their surroundings, and perform tasks that are costly, dangerous or difficult for humans to undertake. The company's signature technology, Neocortex, which was developed over seven years at NASA and Vanderbilt University, will increase efficiency and worker safety across industries in applications including warehousing, mining, handling hazardous waste and automating vehicles such as forklifts.

Useful Progress

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 highperformance visualisation.

VisiSonics Corporation

VisiSonics RealSpace© tools combine proprietary true real-time acoustic signal processing algorithms with computer vision on CUDA platforms to provide enabling technologies addressing real world solutions. VisiSonics RealSpace offerings demonstrate advanced telepresence combining vision and sound in a unified solution space. Applications include industrial acoustic analytics, virtual reality, surveillance and HRTF.

GTC Exhibitors

Ace Computers

Ace computers is a 27 year old system integrator with a focus on high performance computing, workstations, servers and storage. We hold WSCA contract B27157, GSA schedule GS-35F-0400T and other BPAs including major universities and federal agencies.

Aspen Systems

Aspen Systems, founded in 1982, is an established, privately-held, two time Inc. 500 corporation that designs, manufactures, and services computing products including high-performance compute clusters, systems software, storage/file systems, and visualization. Aspen Systems places its highest priority on first class technical support and the creation of fully customized products that always incorporate the latest technologies. This allows our customers to enjoy the highest performing solutions at very competitive prices.

Boxx Technologies

BOXX Technologies offers high-performance solutions that empower your animation, VFX, architectural or industrial design software and renders your creations faster than any other. From our XXtreme, 4 GPU workstations to our state-of-the-art rendering systems, BOXX solutions are built by the one company that understands and supports your creative business.

Bright Computing

Bright Computing is a specialist in cluster management software and services for high-performance computing (HPC) in general and GPU clusters in particular. Its flag-ship product — Bright Cluster Manager — makes clusters of any size easy to install, use and manage. Bright Cluster Manager has earned a reputation for being the best cluster management software on the market for GPU clusters. Bright Computing works closely with NVIDIA and has customers in industry, government and academia across the world.

CIARA Technologies

CIARA designs, develops, markets, services, and supports a variety of High Performance systems including TITAN Systems based on NVIDIA® Tesla™ 20-series, NEXXUS-4000® and CX1 Personal Cluster, the acclaimed VXRACK® high density blade server, VXPRO® rack-mount/tower servers and GRAPHIXX® high-end workstations.

Cirrascale Corporation

Cirrascale Corporation is a premier developer of blade-based cloud servers, GPGPU systems and scale-out storage solutions. Cirrascale's patented Vertical Cooling Technology allows it to provide the industry's densest, most energy-efficient, reliable standards-based solutions. To learn more about Cirrascale and its unique blade-based solutions, please visit http://www.cirrascale.com or call (888) 942-3800.

Colfax International

Buy it from a trusted expert. Colfax provides the most comprehensive range of innovative, cutting-edge and highly customized GPU solutions. Colfax has been first-to-market with a 4GPU PSC and the revolutionary CXT8000 - the world's first 8GPU server unveiled at GTC 09. Leading universities, labs and companies are accelerating their research and business outcomes with optimally configured, ready-to-go Colfax GPU solutions. Join us for an in-person conversation at Booth #xxxx. Or visit www.colfax-intl.com

Creative Consultants

Creative Consultants LLC is a New Mexico based business producing efficient high performance computing for Scientists, and Engineers. We design, and support computing systems advancing leading edge, but sensible technologies. We specialize in the integration of outside source and in-house developed software with our custom GPU accelerated hardware to produce complete solutions. Creative Consultants has serviced America's National Laboratories for more than twenty years. Our company offers an ecosystem of products enabling the creation of GPU clustered

appliances. This includes the Stelletto dual node workstation, our eStella cloud service, and Stella, a high throughput conSTELLAtion cluster for local GPU supercomputing.

Cubix Corporation

Cubix Visual & GPU Compute Solutions is a new division of Cubix Corporation, a vertically-integrated manufacturer with 35 years of manufacturing experience, focused on delivering deskside and rackmount modular, scalable, GPU Compute hardware solutions for demanding applications such as physically-based rendering, animation, visualization, and cloud computing applications.

Exxact

Founded in 1992, Exxact develops and manufactures high performance computing workstations, servers, and clusters that serve compute-intensive applications such as space exploration, medical imaging, financial modeling, broadcasting, and industrial design. Exxact is also a value-added distributor of professional workstation graphics by ATI FirePro™ and Nvidia® Quadro® by PNY.

GIADA

Giada is a leading innovator of small form-factor PCs. Its state-of-the-art technology and stylish design is exemplified in its flagship product, the $299 Slim-N20. At only 1.2lbs, it is the world's smallest PC that uses NVIDIA's next generation ION processor, and is the perfect digital home entertainment solution, providing a robust platform for Internet entertainment. Giada, already a leading brand in the entertainment category for mini-PCs in China, is establishing a presence in the U.S. to promote its innovative products, including all-in-one PCs, OEM motherboards, graphics cards and Netbooks.

GraphStream

GraphStream is a leading provider of custom-integrated scalable computing+storage systems with GPU acceleration. GraphStream partnered with NVIDIA, Mellanox, and Supermicro in 2003 to deliver the world's first commercially integrated scalable GPU-computing systems with InfiniBand cluster interconnect. Since then, GraphStream has delivered advanced GPU-computing systems to more than 200 customers worldwide, including four Top10 supercomputing sites.

James River Technology Inc.

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.

JMR Electronics Inc

JMR is a leading value provider of scalable storage systems for performance and capacity driven applications in the government, DCC, VOD, video surveillance and Web 2.0 markets. Headquartered in Chatsworth, California, JMR has been developing reliable, high performance RAID storage technologies since 1982. JMR's complete line of BlueStor PeSAN™ DAS, NAS and SAN solutions, manufactured entirely in the U.S.A., are ideal for nearly every IT and video production need. For further information please visit, www.jmr.com.

Koi Computers

Koi Computers, Inc. has over fifteen years of experience in the IT hardware and systems integration industry. We offer a wide range of custom configured computer systems and an extensive catalog of technology products. Our core competencies include computer high performance clusters; server, storage, and blade solutions; desktops, laptops, and workstations; mounting and cabling solutions; and strategic sourcing for technology products.

Los Alamos National Labs

Los Alamos National Laboratory is a premier national security research institution, delivering scientific and engineering solutions for the nation's most crucial and complex problems. Our primary responsibility is ensuring the safety, security, and reliability of the nation's nuclear deterrent.

Mathworks

Over 1,000,000 engineers and scientists in more than 100 countries, on all seven continents, use MATLAB® and Simulink®. These products have become fundamental tools for work at the world's most innovative technology companies, government research labs, financial institutions, and at more than 5000 universities. For more information, visit www.mathworks.com

Mazda Technologies

Mazda Technologies creates significant business advantages for its customers by providing differentiated, value-added hardware and software IT solutions and professional services based on GNU/Linux and other open source and current OS solutions. At Mazda Technologies our main goal and mission is to always ensure that our team is knowledgeable, competent and customer focused in dealing with cutting edge technology thus providing our customers the absolute best value for their IT dollars and the necessary peace of mind to run their IT Operations efficiently. We strive to deliver custom high performance solutions that will fit the needs of the customer.

Mellanox Technologies

Mellanox Technologies is a leading supplier of end-to-end connectivity solutions for servers and storage that optimize data center performance.

Micoy

Micoy is the only true provider of full omni-directional 3D Stereo Immersive systems. This revolutionary 3D format is a breakthrough patented method in 360 degree panospheric & cylindrical 3D optics. With a secure patent portfolio for the technology Micoy is looking for development partners for applications of its technology and co-development opportunities in real time solutions.

Microway Inc.

Since 1982, Microway has earned an international reputation for building robustly designed and cooled HPC clusters and WhisperStations. Since 2007, these included Tesla GPUs. Utilizing multi-core hosts with high-efficiency power, excellent cooling, and QDR InfiniBand interconnects, Microway's GPU clusters deliver more TFLOPs with fewer watts. Microway is known for the quality of its design, integration, and support teams and high level of customer satisfaction.

Morgan Kauffman

Morgan Kaufmann has been bringing the knowledge of experts to the computing community since 1984. Our goal is to provide timely yet timeless content to research and development professionals, business leaders and IT managers, everyday practitioners, and academia. We publish textbooks and references in Artificial Intelligence, Computer Networking, Computer Architecture, Computer Graphics & Game Development, Data Management & Business Intelligence, Software Engineering, and User Experience & Human Computer Interaction. For more information, visit mkp.com.

Next Computing

NextComputing manufactures high-performance portable and small form-factor workstations and servers, designed to run high-end applications across a range of industries. Leveraging NVIDIA's 'Fermi' GPU architecture, NextComputing offers unique portable and small footprint platforms for professional power-users who require more performance and flexibility than ordinary notebook and desktop computers can provide.

PEER 1 hosting

One of the world's leading IT hosting providers built on two obsessions: Ping & People. A 10Gb SuperNetwork™ connects 17 datacenters and, paired with 24x7 FirstCall Support, helps power over 10,000 customers worldwide. PEER 1 recently launched the first ever large-scale GPU Cloud hosted in North America and UK.

Penguin Computing

Penguin Computing is a global leader in high-performance computing (HPC), delivering complete, integrated HPC solutions, from the workstation to the cloud. With a focus on cutting-edge technology, ease-of-use and exceptional customer service, Penguin cost-effectively meets the needs of the world's most demanding HPC users, including Caterpillar, Lockheed Martin, the U.S. Air Force, and the U.S. Navy. Today, Penguin delivers a range of solutions, from massive Linux clusters to "Penguin on Demand" (POD), a new service that provides a complete HPC solution in the cloud. Penguin has been an innovator in HPC solutions for over a decade, and the company's founder Donald Becker is recognized as the "Father of Linux Clustering." For more information about Penguin Computing and Penguin products please go to www.penguincomputing.com.

Platform Computing

Platform Computing is the leader in cluster, grid and cloud management software – serving more than 2,000 of the world's most demanding organizations. For 18 years, our workload and resource management solutions have delivered IT responsiveness and lower costs for enterprise and HPC applications. Platform has strategic relationships with Cray, Dell, HP, IBM, Intel, Microsoft, Red Hat, and SAS. GPU-accelerated clusters are rapidly growing in popularity as powerful, cost-effective High Performance Computing (HPC) solutions. NVIDIA's GPU hardware, along with the CUDA computing environment, is delivering impressive results for commercial HPC workloads. Platform HPC suite, with CUDA kit, enables analysts, engineers, and scientists to unlock the power of NVIDIA GPU Clusters, making them easier to deploy, run and manage.

Polywell

Polywell is a manufacturer of high quality computer products ranging from PCs to high-performance workstations, storage solutions and high-end servers, and has been serving the needs of consumers, businesses and government entities since 1987. Its direct sales channel and Just-in-time manufacturing processes allow competitive pricing. There are system specialists to support the various vertical markets, including data centers, IPTV, entertainment, gaming, biotech, oil/gas, CAD/CAM, animation, and content creation. Polywell also provides OEM service for PC/Linux appliances in digital signage, set-top box, POS, kiosks, medical equipment and network appliances. Its full-range service includes product design, prototyping, production, fulfillment, and warranty repair.

Portland Group

The Portland Group® (PGI®) offers high performance parallel compilers and tools for workstations, servers and clusters based on 64-bit x86 processors with NVIDIA CUDA-enabled GPUs running under Linux, MacOS or Windows operating systems. PGI GPU accelerator products include directive-based PGI Accelerator™ Fortran and C compilers and CUDA Fortran.

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.

Tech- X Corporation

Tech-X offers products and services for high-performance computing. GPULib enables users of MATLAB and IDL to take advantage of GPUs from within these high-productivity languages. We offer consulting, training, and custom software development to migrate customers' scientific computing problems onto hardware accelerated architectures, using technologies like CUDA, OpenCL, or MPI.

T-Platform

T-Platforms is the major Russian supercomputing group. T-Platforms installations occupy 38% of the current Top50 list of Russia's most powerful computer systems. T-Platforms designed and manufactured 5 supercomputers featured in the global Top500 list, 36th being the highest rank. T-Platforms has accumulated substantial expertise with over 200 HPC installations.

Wolfram Research

Wolfram Research, Inc. is the technical innovation powerhouse behind the world's most powerful global computation system Mathematica, and the world's first-ever computational knowledge engine, Wolfram|Alpha. Wolfram Research continues its strong commitment to technology and education with resources like MathWorld, load-on-demand curated data, and the Wolfram Demonstrations Project. www.wolfram.com

Marketing Partners

VISUALIZE **GREEN**

VISUALIZE A GREEN EVENT

Place compostables 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 convention center's greenhouse gas is offset
- Extensive composting and recycling
- Producers and vendors agree to green guidelines
- Minimizing printed materials
- Using recycled and biodegradable paper/non-toxic inks
- Monitoring lighting and A/C usage
- Local and organic food options
- Non-toxic cleaning materials

LET'S DO AMAZING

WITH WORKSTATION AND SERVER PRODUCTS KNOWN FOR

INNOVATION

AWARD WINNING PRODUCTS

PERFORMANCE

BROAD GPU ENABLED PORTFOLIO

RELIABILITY

APPLICATION FOCUSED DESIGN

Find your AMAZING with HP

In the **Exhibit Hall** or **Session #2233**

Solving Your GPU Computing Needs

Join Us Tues. 2pm Room 13

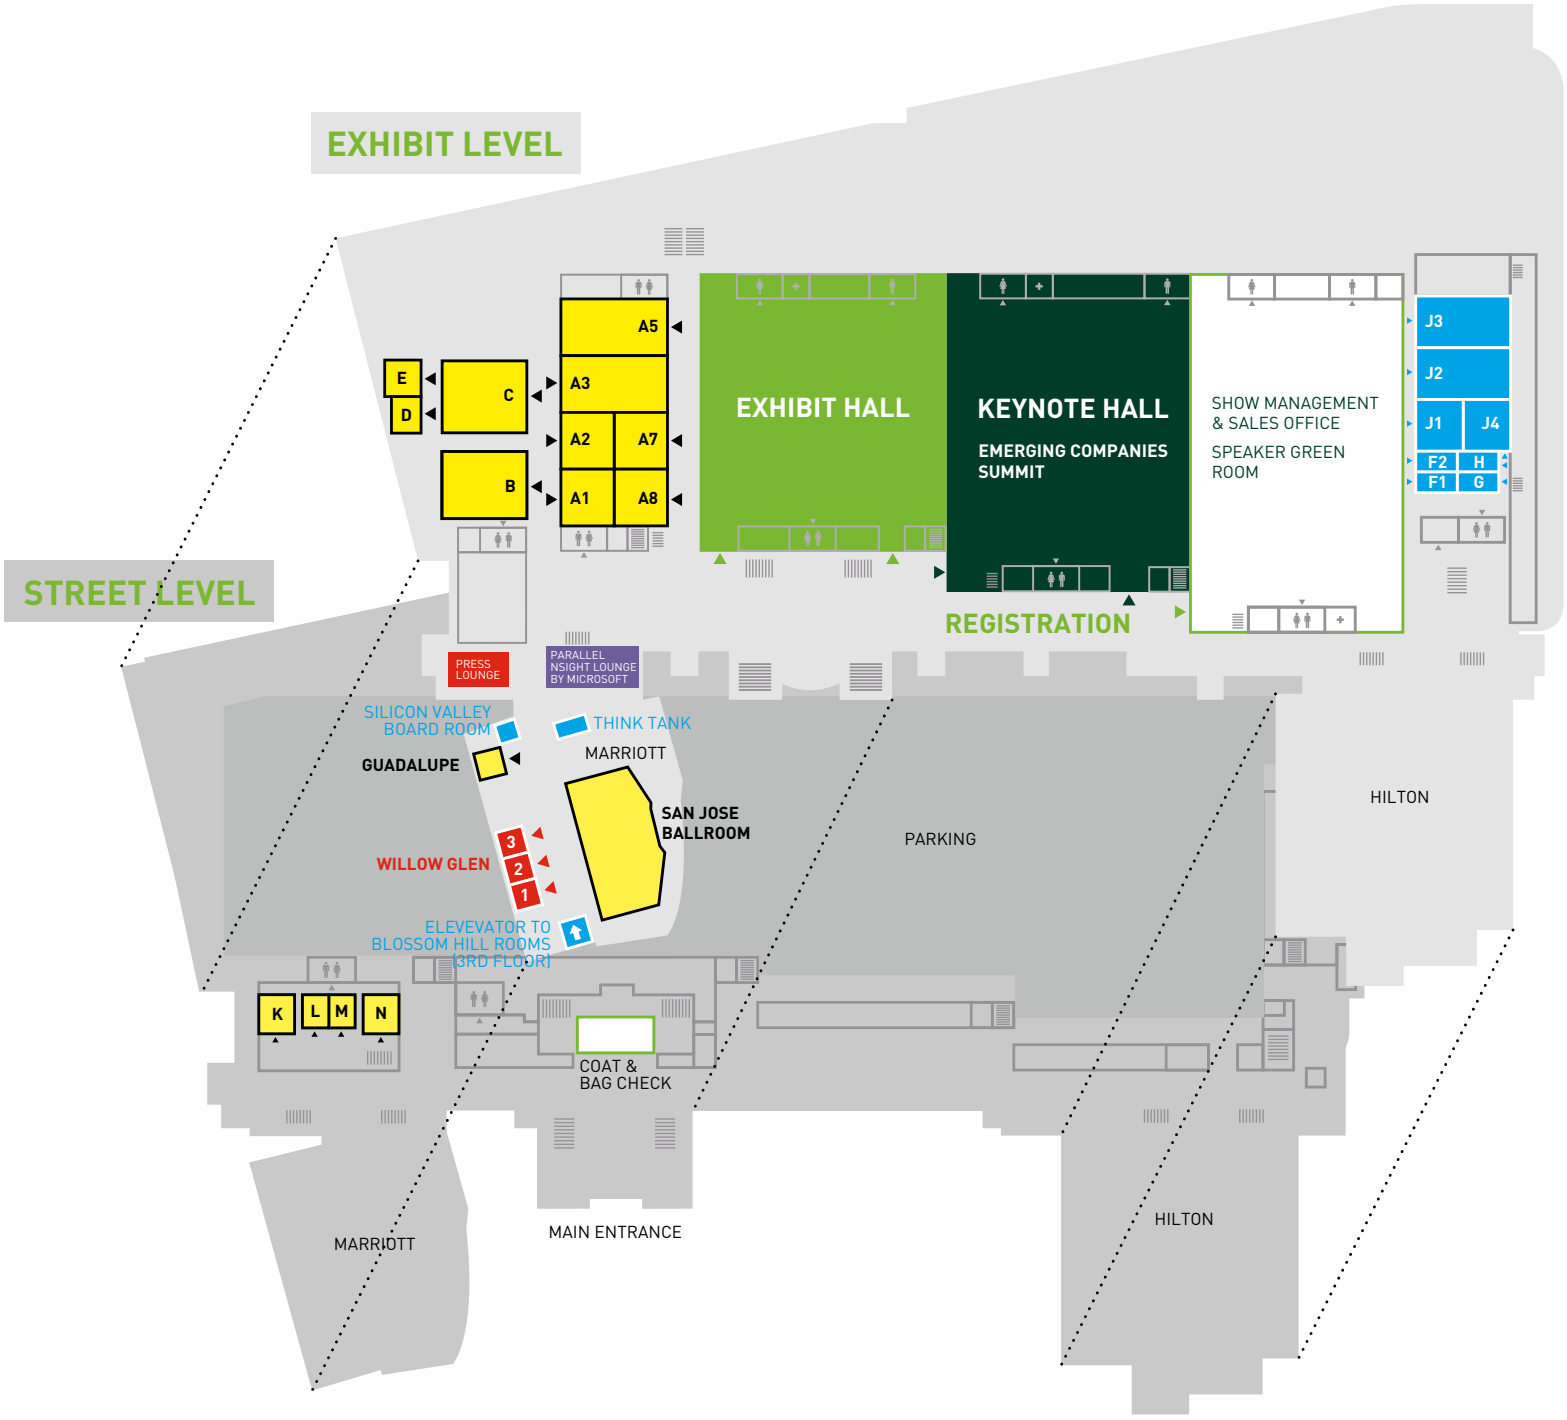

GPU TECHNOLOGY CONFERENCE