GPU accelerated Cray XC systems: Where HPC meets Big Data

Peter Messmer (NVIDIA), Chris Lindahl (Cray), Sadaf Alam (CSCS)

Cray User Group Meeting 2016, London, May 8 – 12, 2016

GPU accelerated Cray XC systems: Where HPC meets Big Data

Peter Messmer

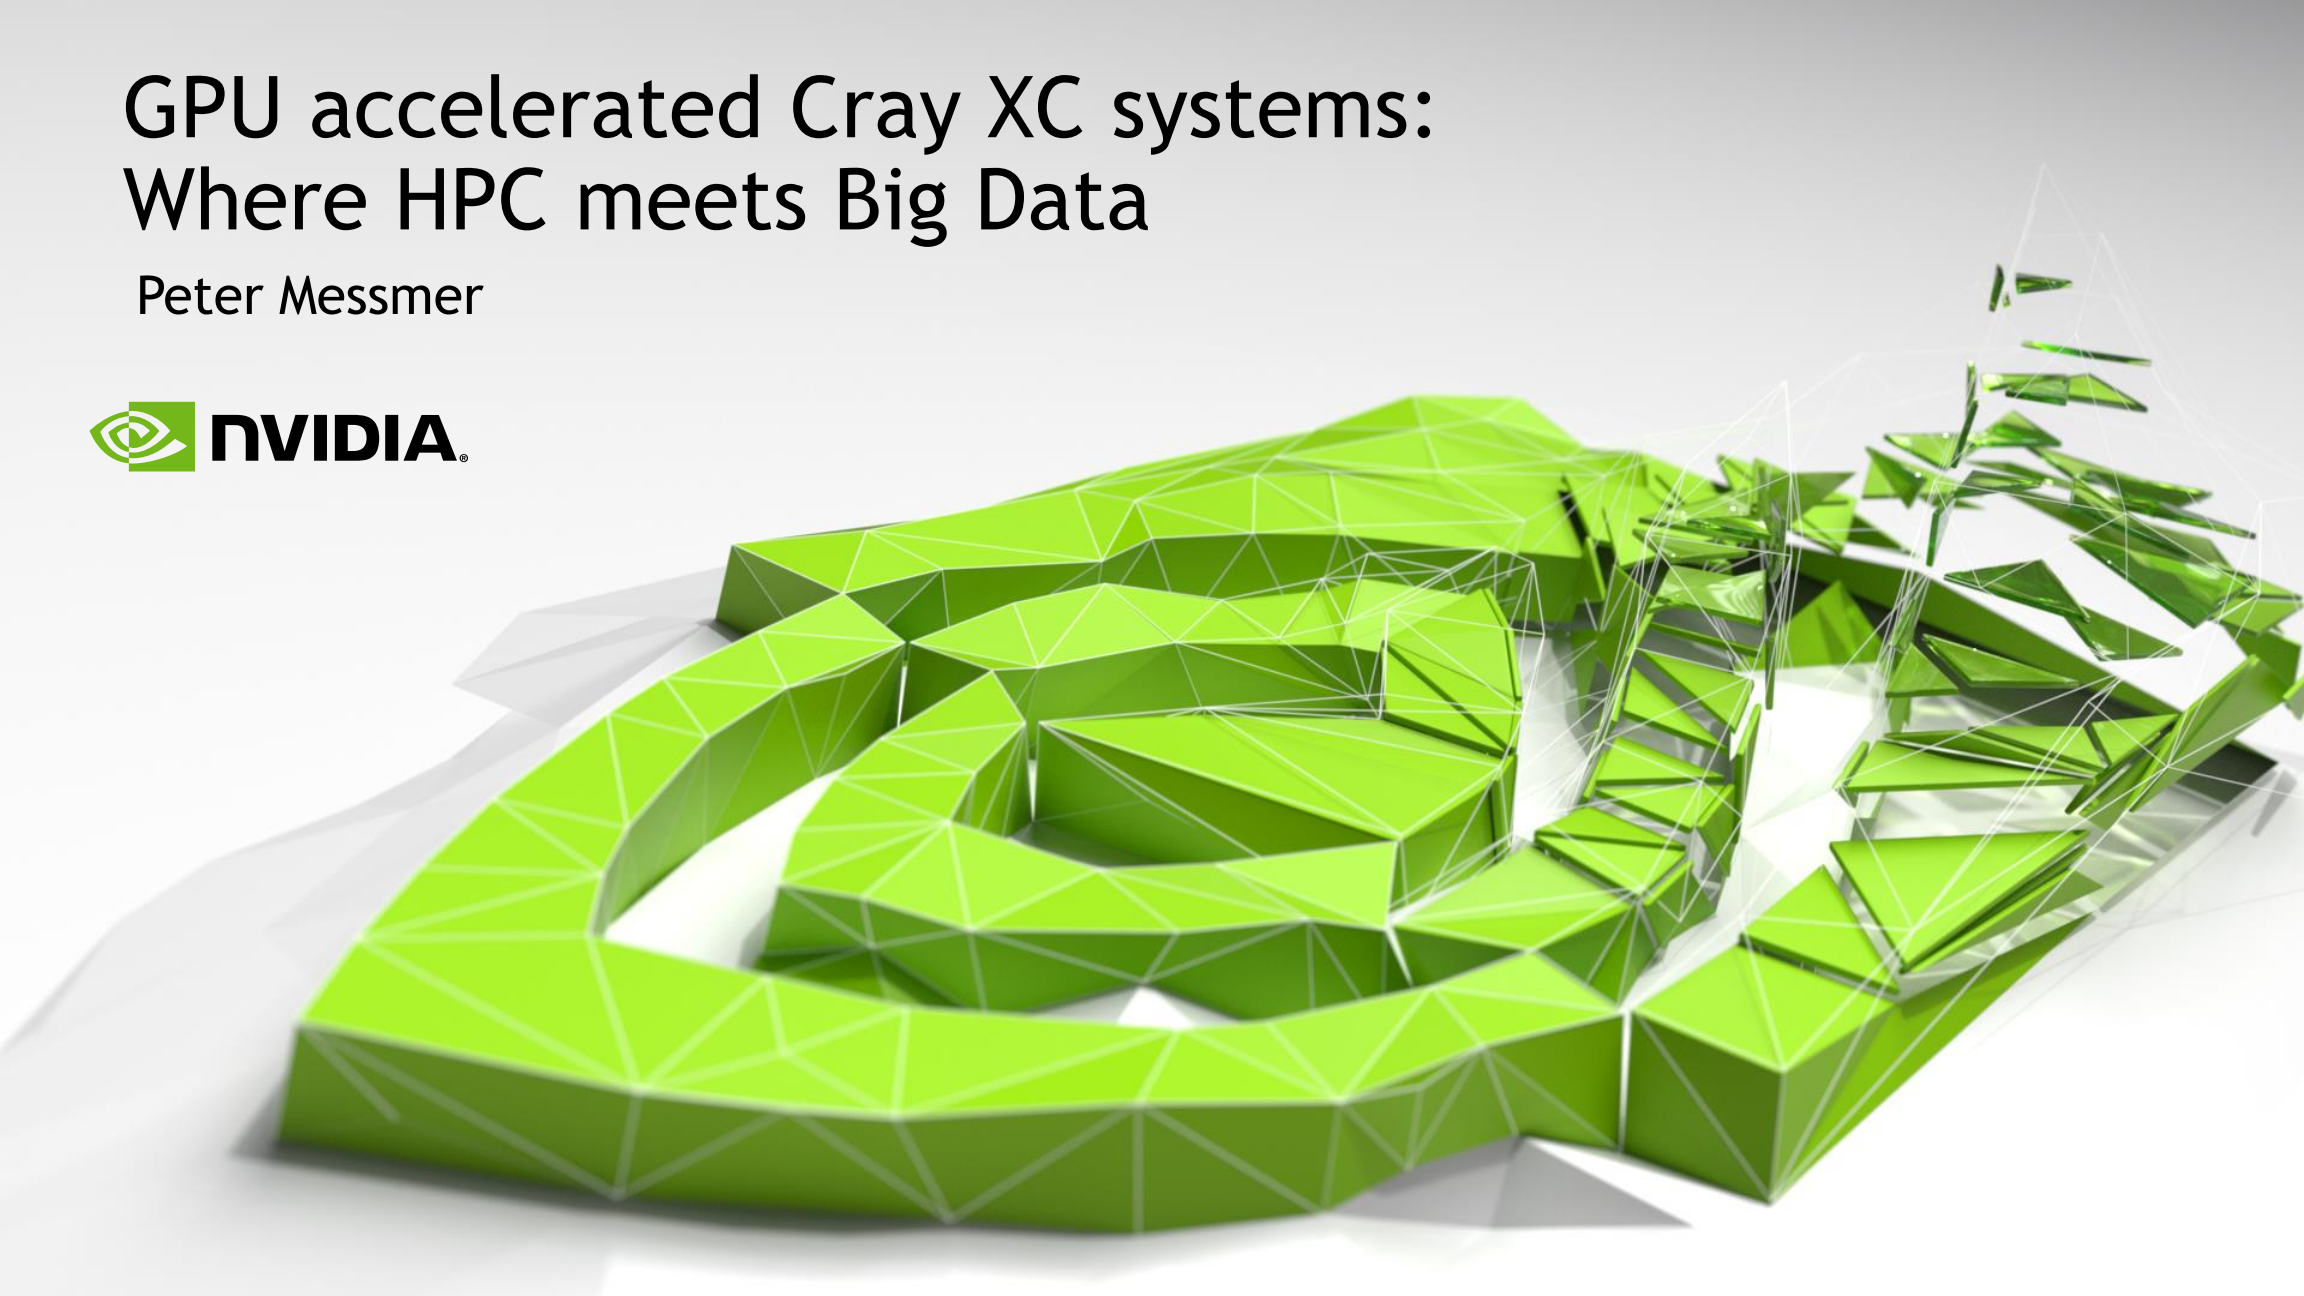

HIGH PERFORMANCE COMPUTING TODAY*

"Yes," said Deep Thought, "I can do it."

[Seven and a half million years later.... **]**

“The Answer to the Great Question... Of Life, the Universe and Everything... Is... Forty-two,' said Deep Thought, with infinite majesty and calm.”

— Douglas Adams, Hitchhiker’s Guide to the Galaxy

*mostly

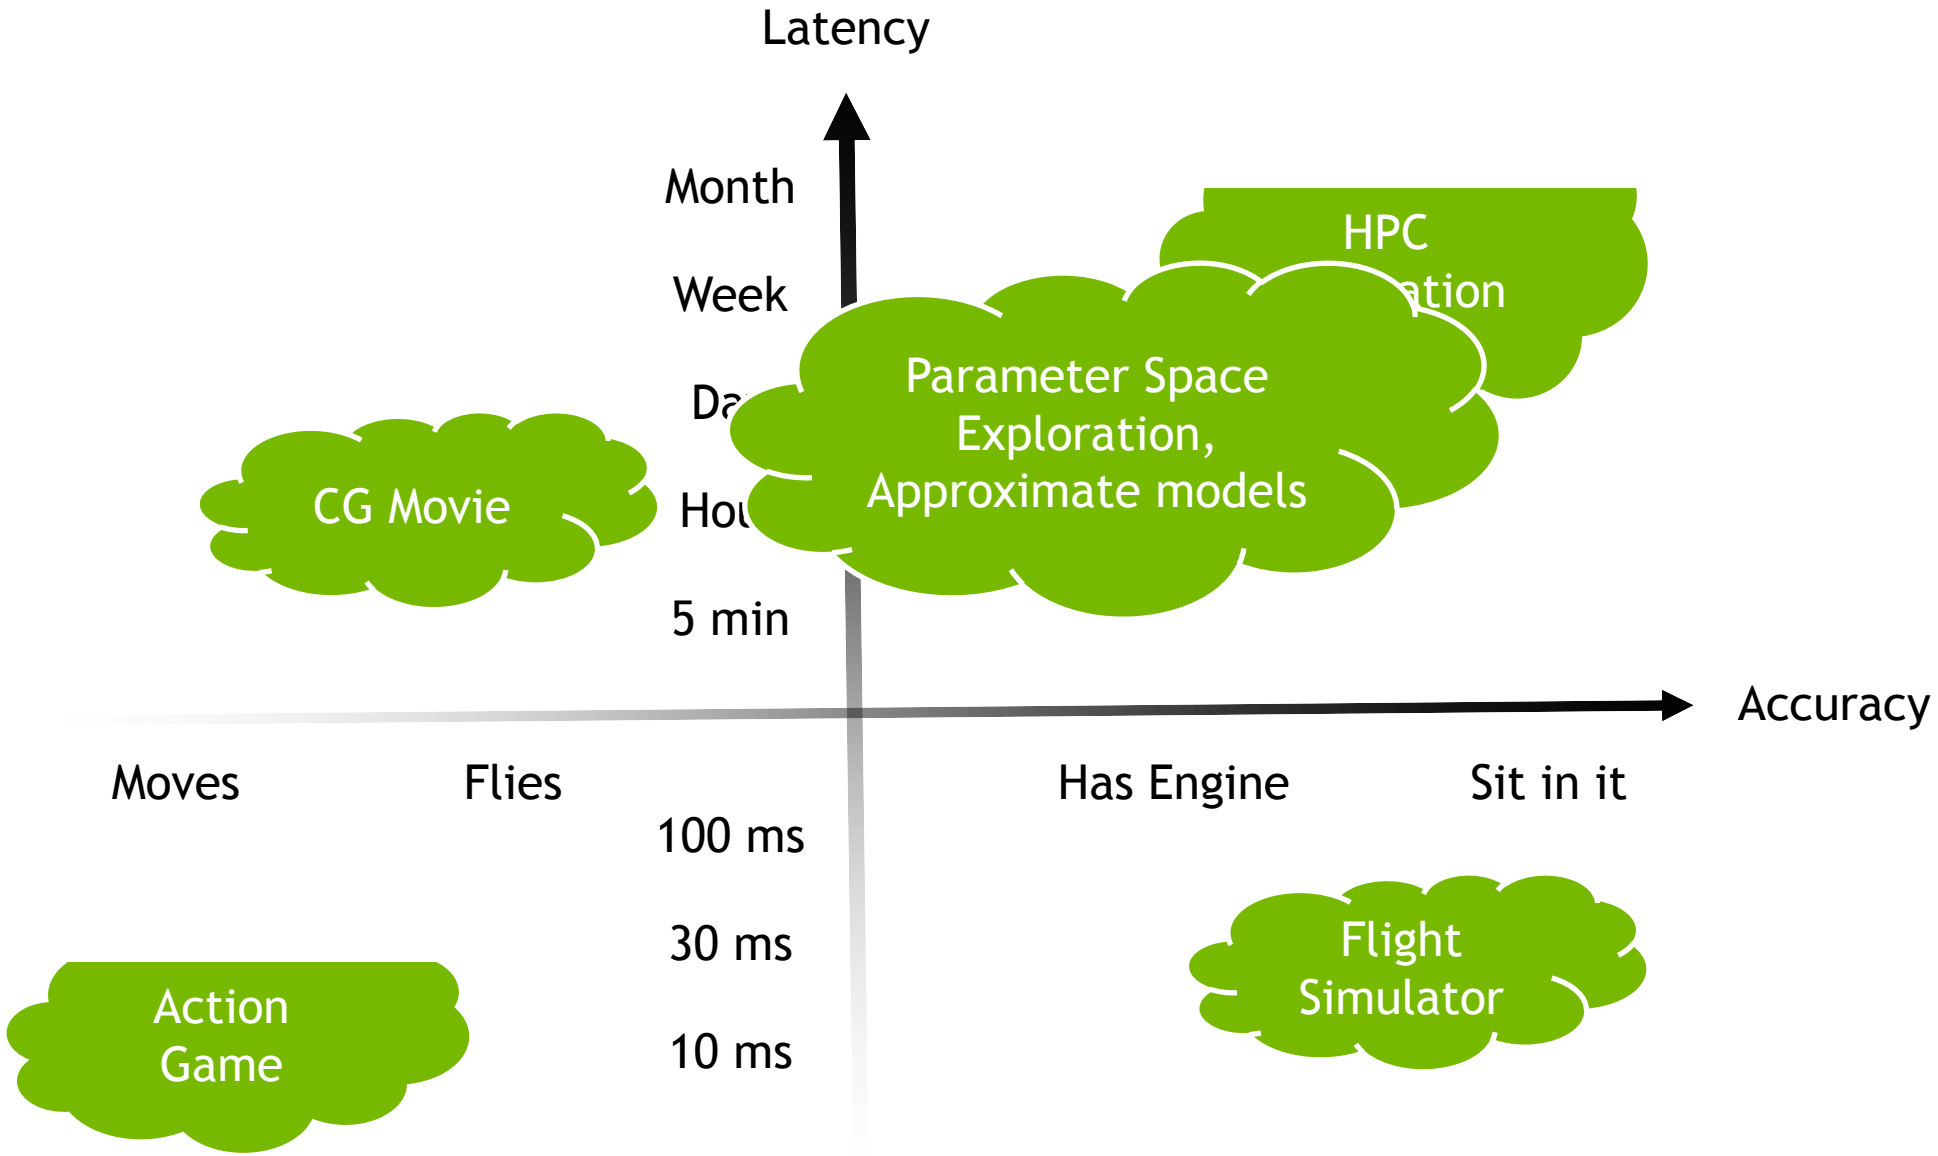
Latency
Month
Week
Da
Hou
5 min
HPC
ation
Parameter Space
Exploration,
Approximate models
CG Movie
Accuracy
Moves
Flies
Has Engine
Sit in it
100 ms
30 ms
10 ms
Action
Game
Flight
Simulator

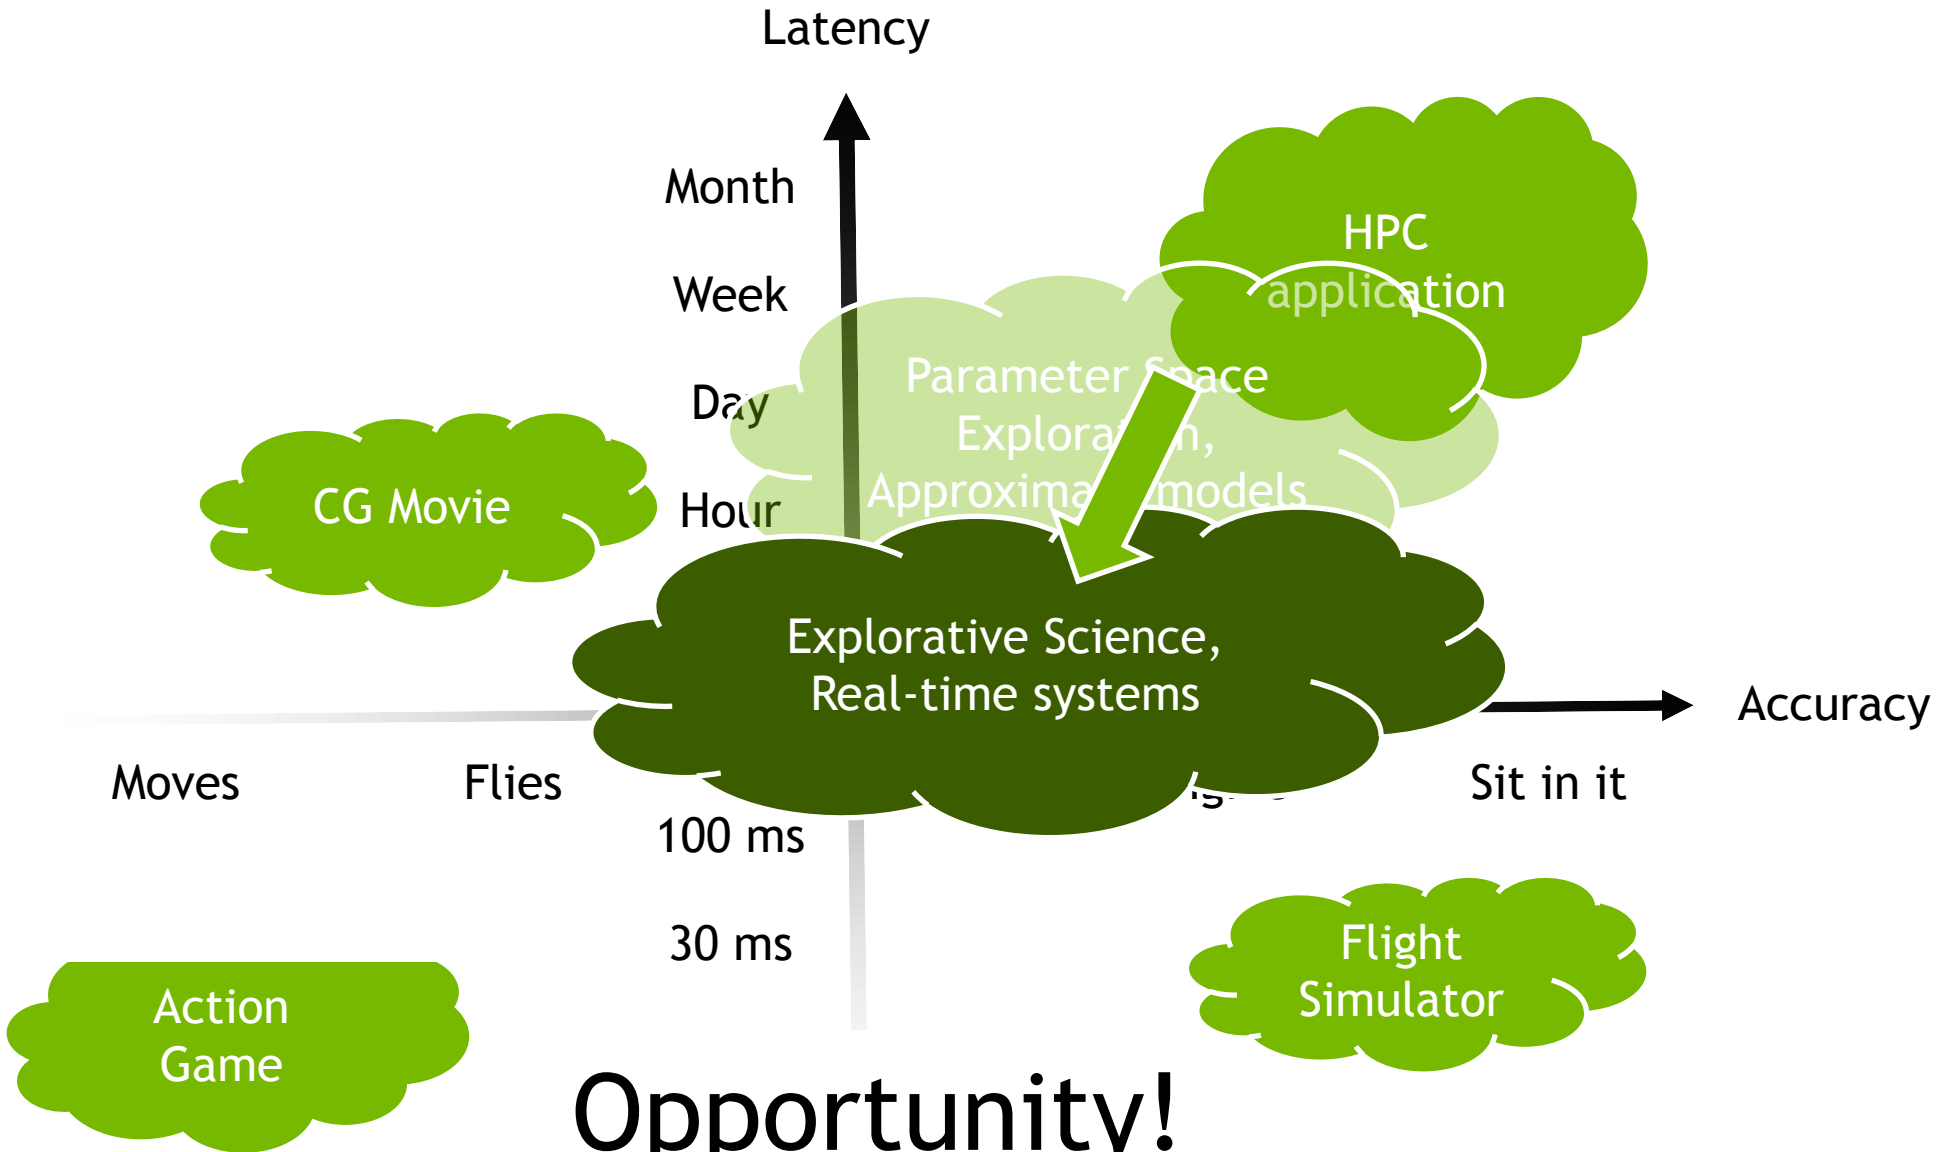

Opportunity!

Bonsai With In-Situ Viz On Piz Daint

Presented at SC14, streaming from CSCS/Switzerland to New Orleans

Bonsai With In-Situ Viz On Piz Daint

Presented at SC14, streaming from CSCS/Switzerland to New Orleans

WIKIPEDIA
The Free Encyclopedia

Main page
Contents
Featured content
Current events
Random article

Hoag's Object

From Wikipedia, the free encyclopedia

Hoag's Object is a non-typical galaxy of the type known as a ring galaxy. The appearance of this object has interested amateur astronomers as much as its uncommon structure has fascinated professionals. The galaxy is named after Arthur Allen Hoag who discovered it in 1950 and identified it as either a planetary nebula or a peculiar galaxy[3] with eight billion

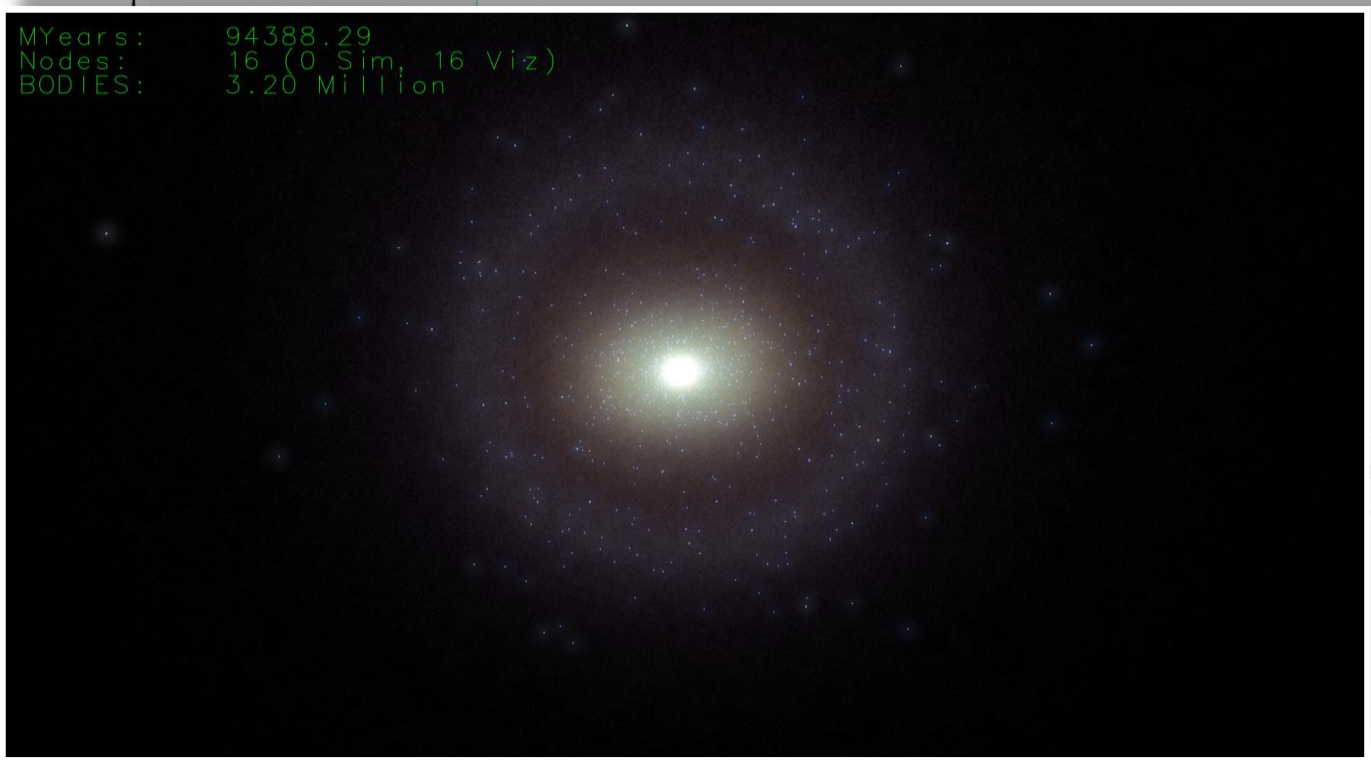

us of this ring galaxy

Hoag's Object

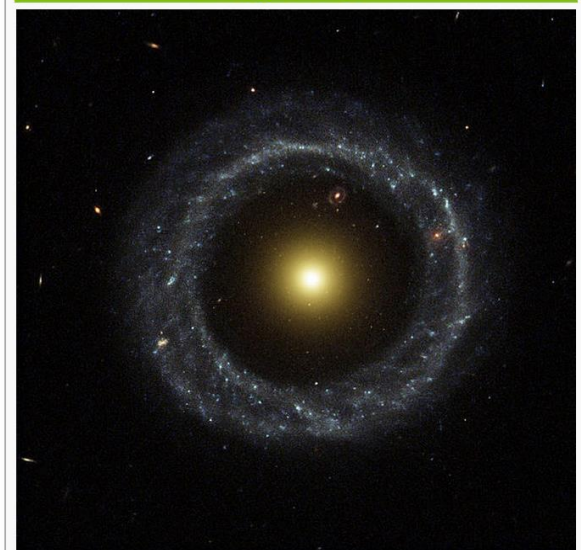

Hoag's Object taken by the Hubble Space Telescope.
Courtesy of NASA/ESA

Visualization-enabled supercomputers

CSCS Piz Daint

Galaxy formation

http://blogs.nvidia.com/blog/2014/11/19/gpu-in-situ-milky-way/

NCSA Blue Waters

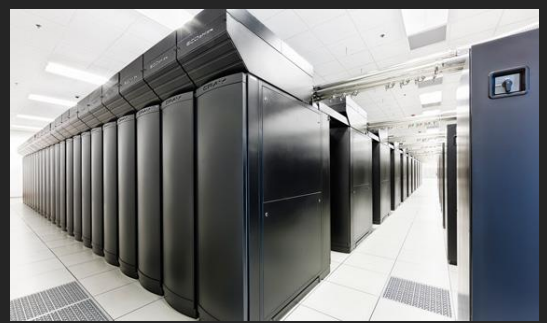

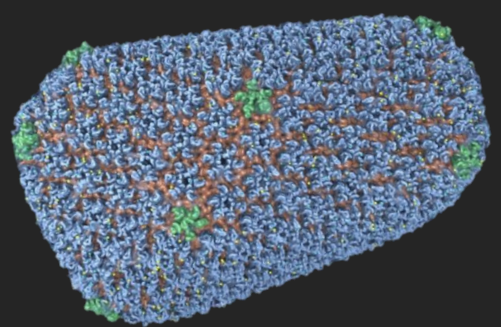

Molecular dynamics

http://devblogs.nvidia.com/parallelforall/hpc-visualization-nvidia-tesla-gpus/

ORNL Titan

Cosmology

http://www.sdav-scidac.org/29-highlights/visualization/66-accelerated-cosmology-data-anal.html

Compute+Vis supports multiple workflows

LEGACY WORKFLOW

Separate compute & vis system

Communication via file system

PARTITIONED SYSTEM

Different nodes for different roles

Communication via high-speed network

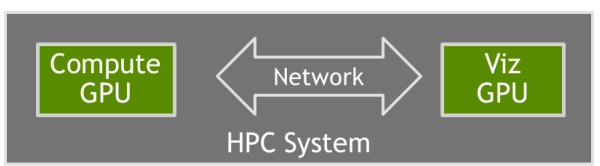

CO-PROCESSING

Compute and visualization on same GPU

Communication via host-device transfers or memcpy

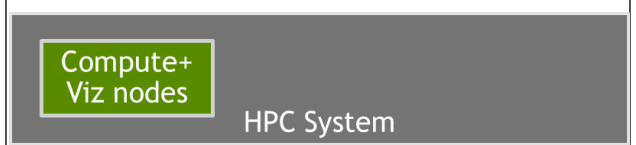

EGL Context Management

Leaving it to the driver

Top systems support OpenGL under X

EGL: Driver based context management

Support for full OpenGL*, not only GL ES

Available in e.g. VTK, ParaView, Ensight, VMD,..

New opportunities for CUDA/OpenGL** interop

*Full OpenGL in r355.11; **CUDA interop in r358.7

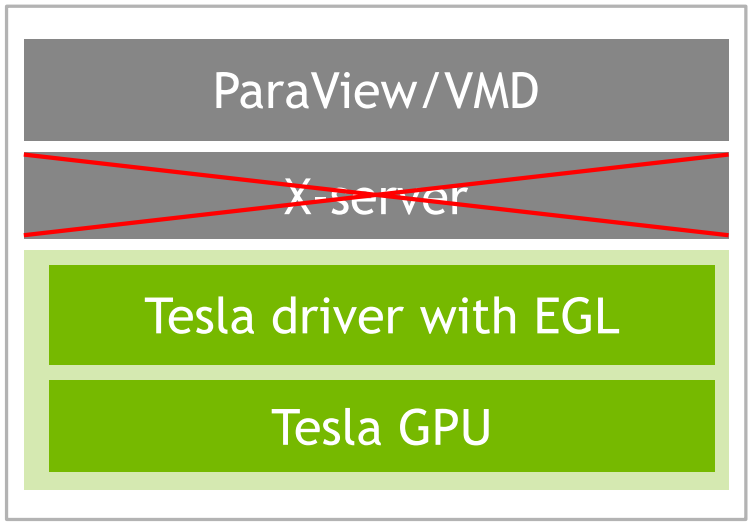

Vis Tools embrace OpenGL on EGL

Streamlined GPU accelerated off-screen rendering

- Prior to EGL: X server required for GPU accelerated OpenGL
- Full OpenGL on EGL announced at SC16
- With EGL: OpenGL without X
- Major enabler for GPU rendering in HPC, incl. Cray systems*
- Quick adoption by vis tool developers
- https://devblogs.nvidia.com/parallelforall/egl-eye-opengl-visualization-without-x-server/
- * Requires driver version 358.7 or newer required

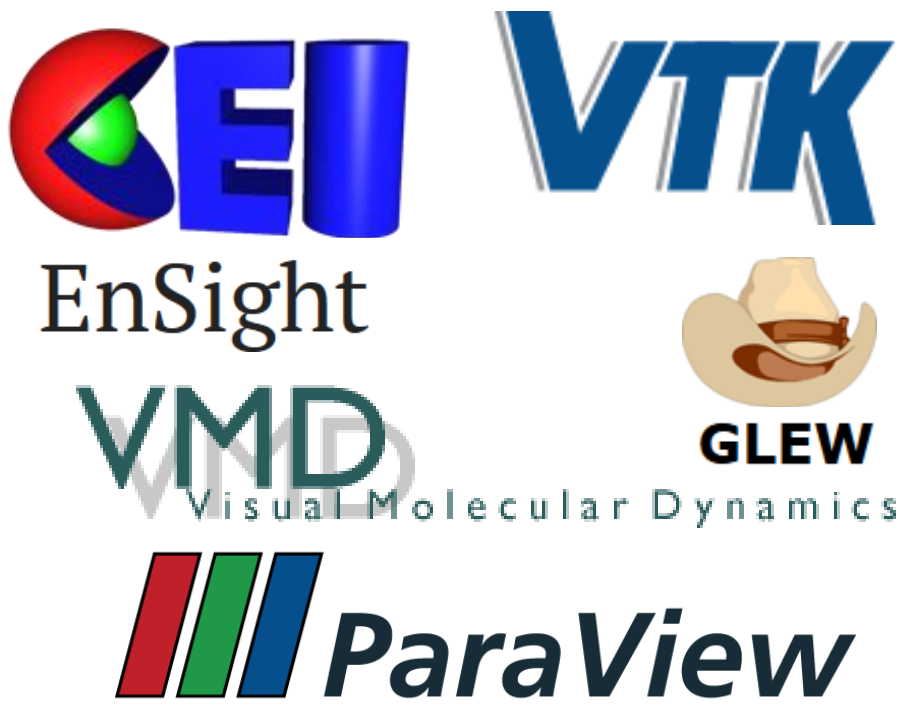

Modern OpenGL for HPC viz

Mandatory to access advanced rendering features

VTK supports now OpenGL 3.2

Enables advanced shaders (AO, VXGI, ..)

Some algorithms well suited for

distributed memory rendering

GPU hardware support

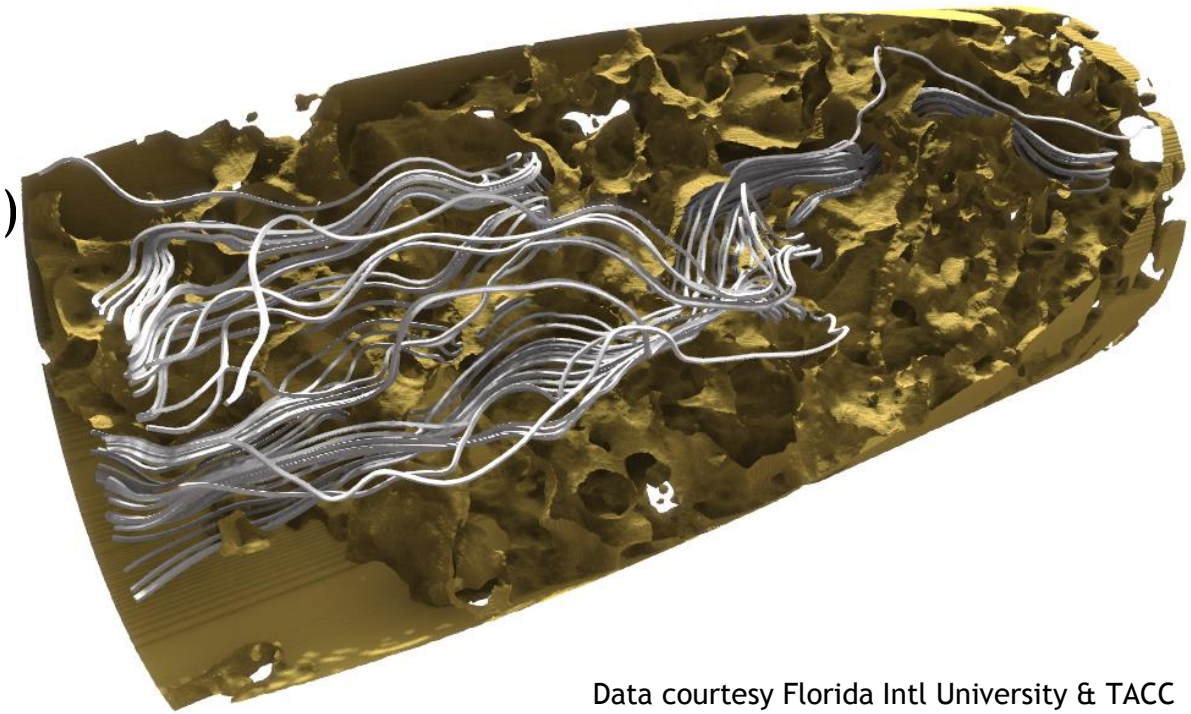

Data courtesy Florida Intl University & TACC

Advanced Rendering in scientific visualization

Better insight with visual cues

Two lights, no shadows

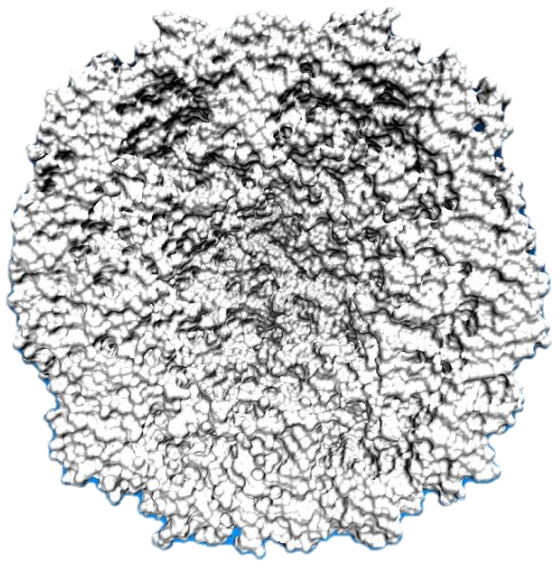

Two lights, hard shadows, 1 shadow ray per light

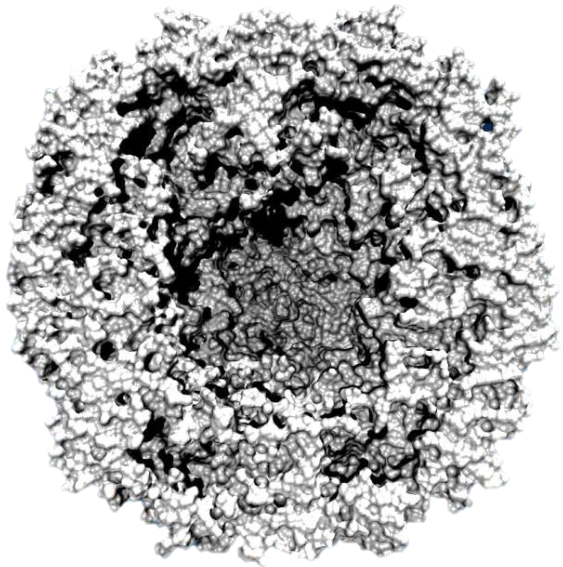

Ambient occlusion + two lights, 144 AO rays/hit

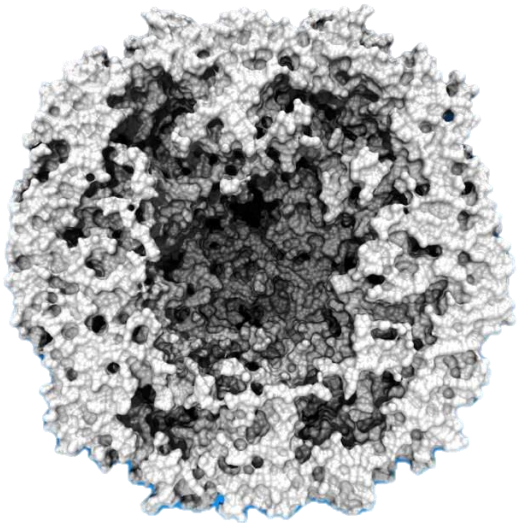

- **Ray tracing offers ambient occlusion lighting, shadows, high quality transparent surfaces**

Courtesy of John Stone, UIUC

OpenGL not limited to Rendering Tasks

Interop goes both ways, esp with EGL

CUDA->OpenGL typically one-way only

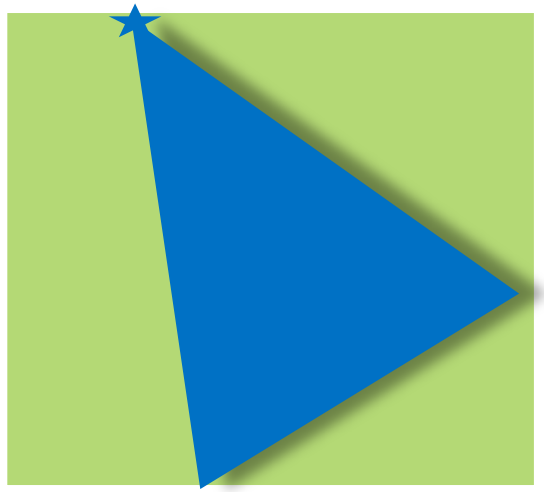

EGL enables lighter weight access to OpenGL

- No X server needed

Potential use of OpenGL for rasterization-like problems?

- Determine covered “pixels”
- 3D ordering/occlusion via Z-buffer

OpenGL Rendering Powerhouse

OpenGL vs OpenSWR

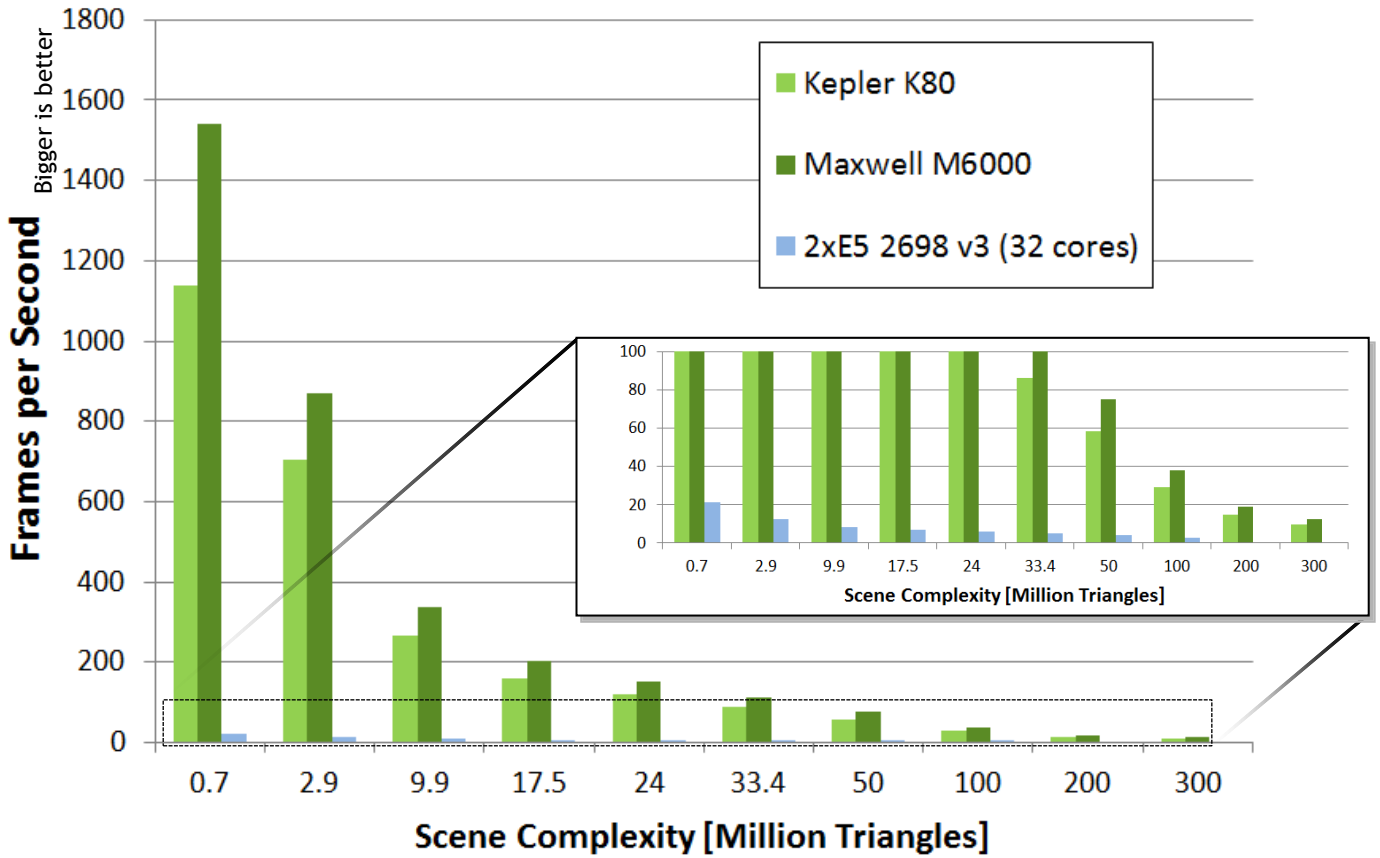

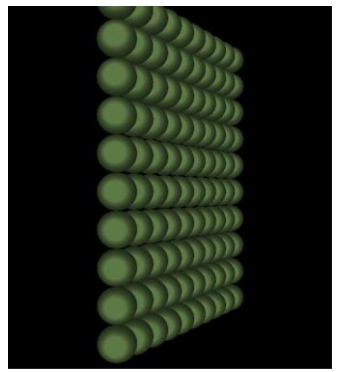

Accelerated remote rendering with Video Encoding

Interactivity over large distances

Lossy and loss-less (Maxwell +) H264 encoder

Separate unit, does not consume “GPU resources”

Leveraged by commercial, free tools

Available on e.g. Titan

Possible use for non-video data

https://developer.nvidia.com/nvidia-video-codec-sdk

NICE DCV running on Titan in user space

INTRODUCING TESLA P100

New GPU Architecture to Enable the World's Fastest Compute Node

Pascal Architecture

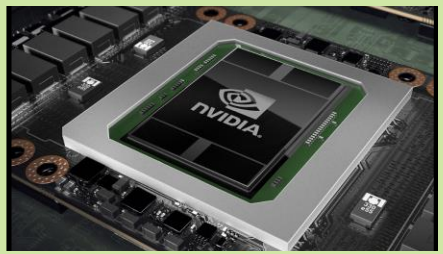

Highest Compute Performance

NVLink

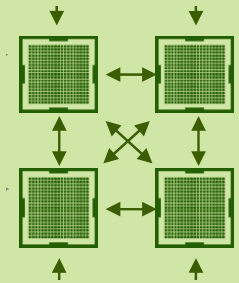

GPU Interconnect for Maximum Scalability

CoWoS HBM2

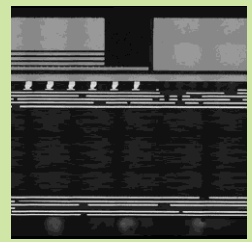

Unifying Compute & Memory in Single Package

Page Migration Engine

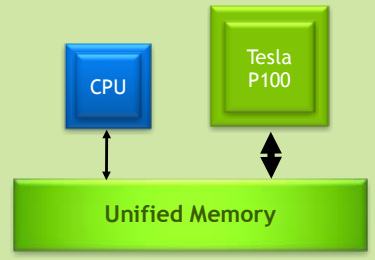

Simple Parallel Programming with Virtually Unlimited Memory

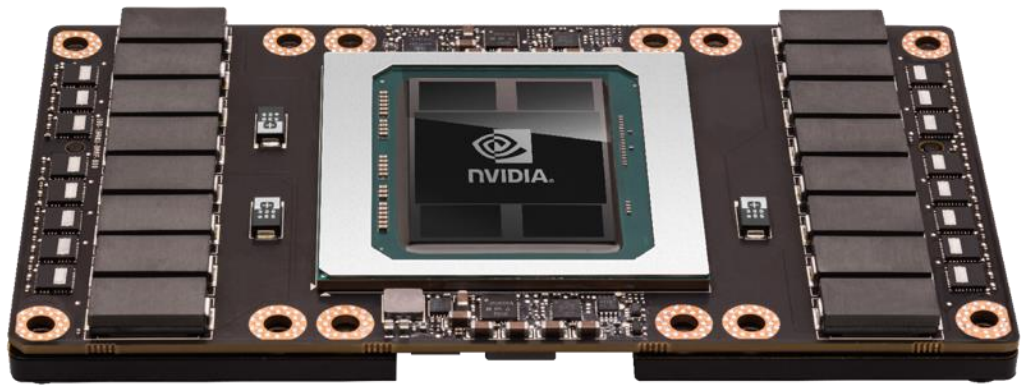

Giant leaps in everything

PASCAL ARCHITECTURE

21 Teraflops of FP16 for Deep Learning

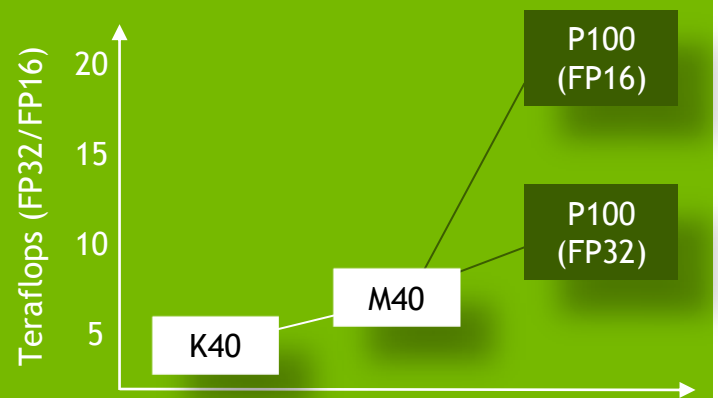

NVLINK

5x GPU-GPU Bandwidth

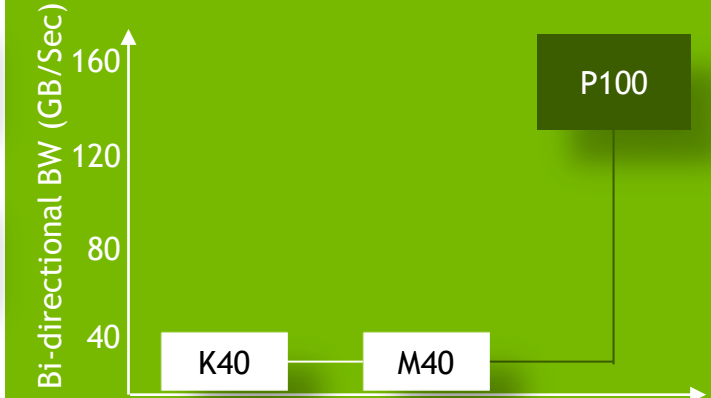

CoWoS HBM2 Stacked Mem

3x Higher for Massive Data Workloads

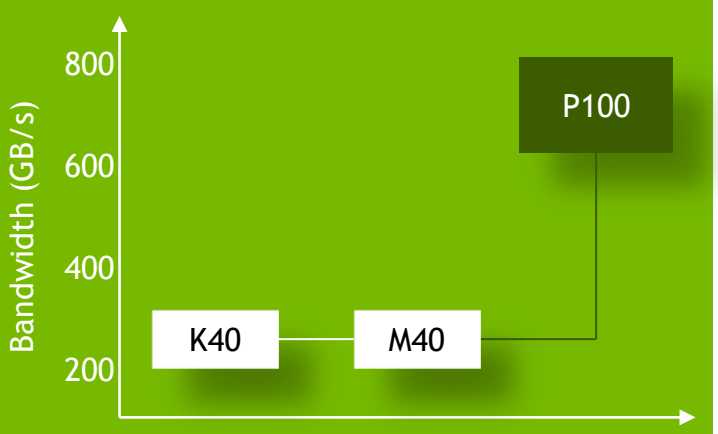

PAGE MIGRATION ENGINE

Virtually Unlimited Memory Space

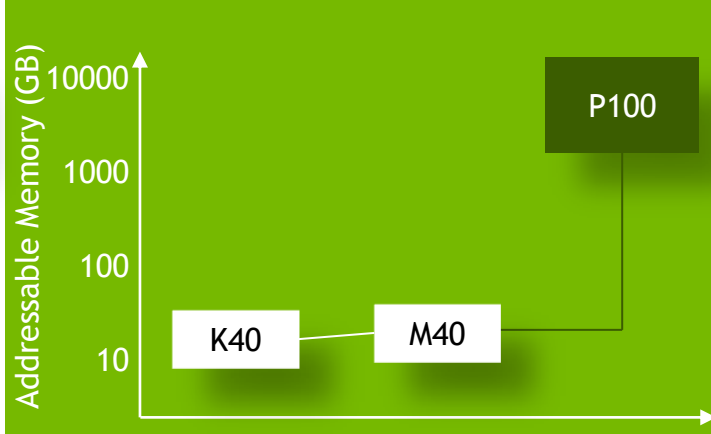

nvGRAPH

Accelerated Graph Analytics

nvGRAPH for high performance graph analytics

Deliver results up to 3x faster than CPU-only

Solve graphs with up to 2.5 Billion edges on 1x M40

Accelerates a wide range of graph analytics apps:

PageRank	Single Source Shortest Path	Single Source Widest Path
Search	Robotic Path Planning	IP Routing
Recommendation Engines	Power Network Planning	Chip Design / EDA
Social Ad Placement	Logistics & Supply Chain Planning	Traffic sensitive routing

developer.nvidia.com/nvgraph

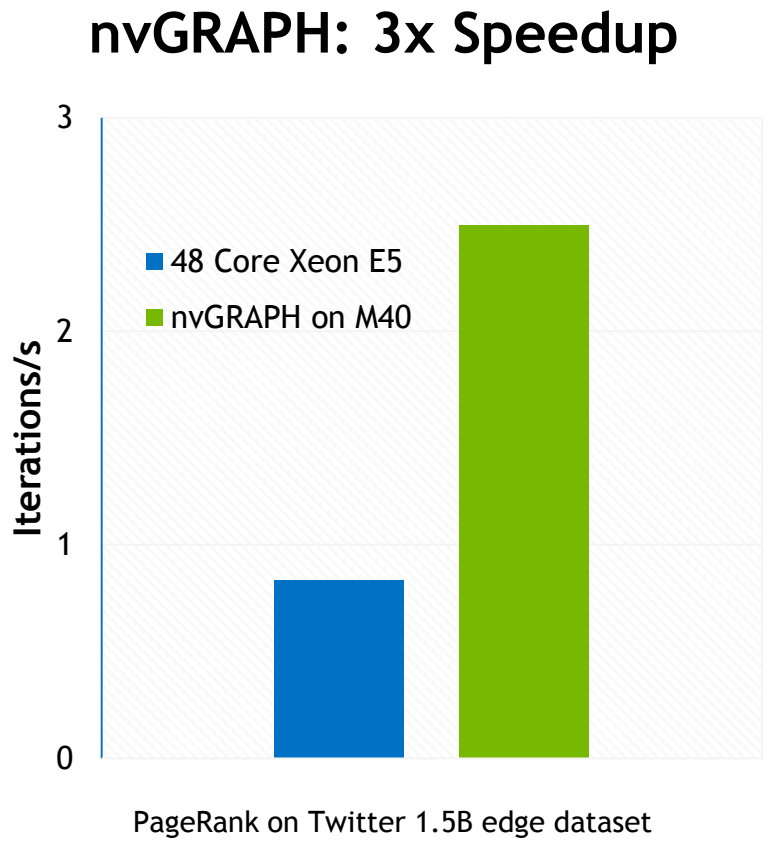

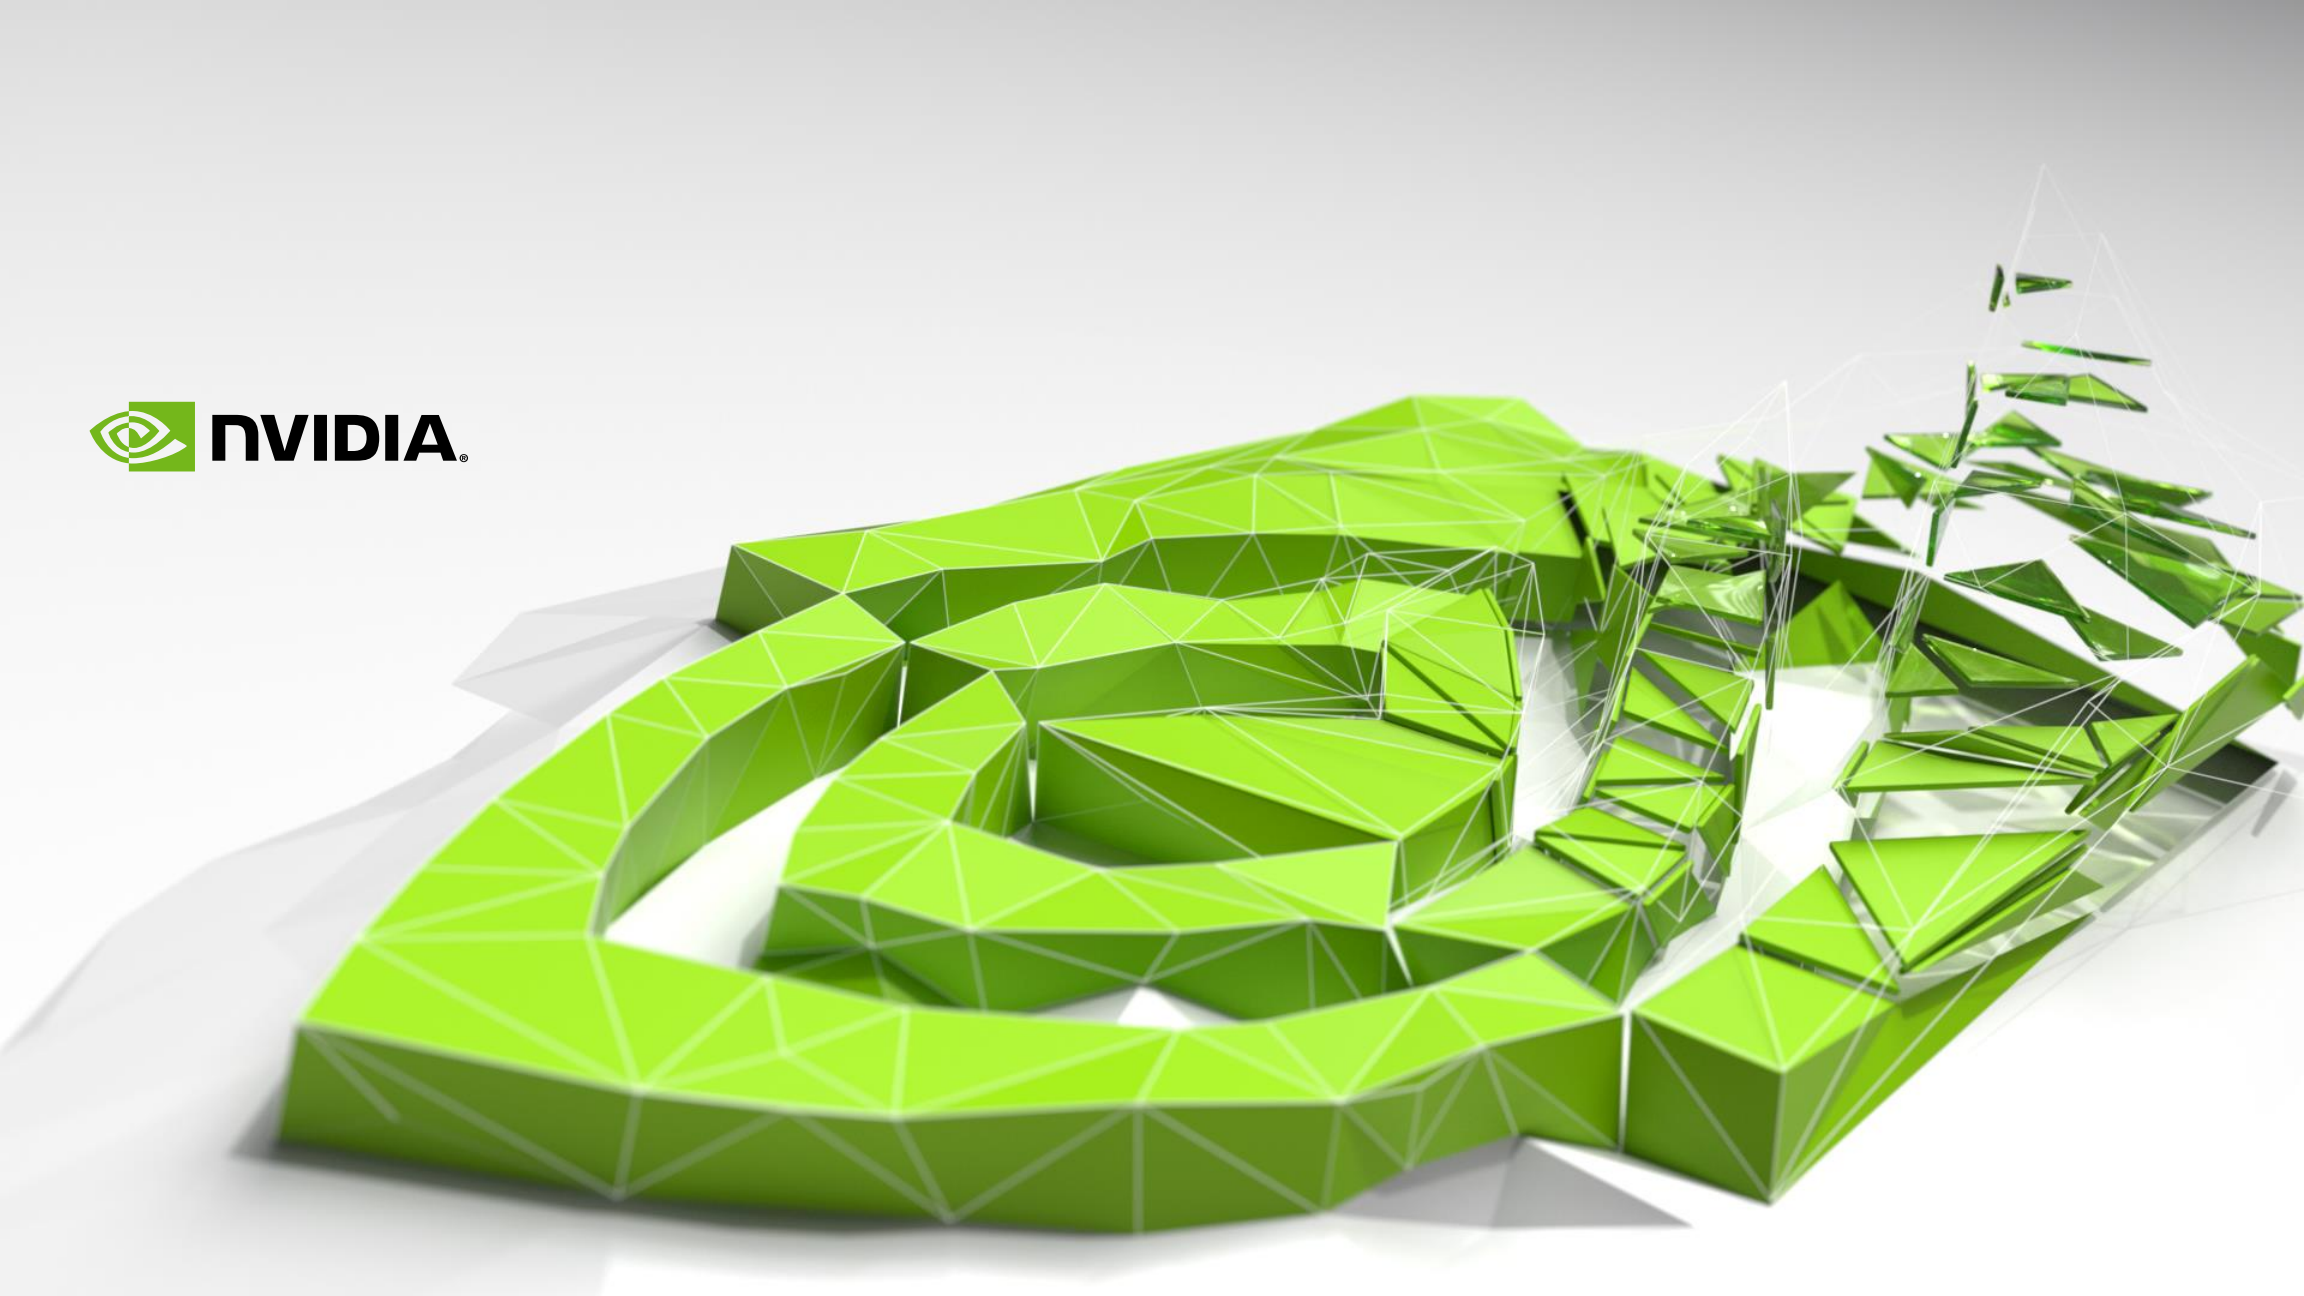
NVIDIA

CSCS
Centro Svizzero di Calcolo Scientifico
Swiss National Supercomputing Centre

ETH zürich

GPUs in XC = GPUs elsewhere (almost)

GPU Accelerated Cray XC Systems: Where HPC Meets Big Data
CUG 2016 BOF

Sadaf Alam, CSCS

May 10, 2016

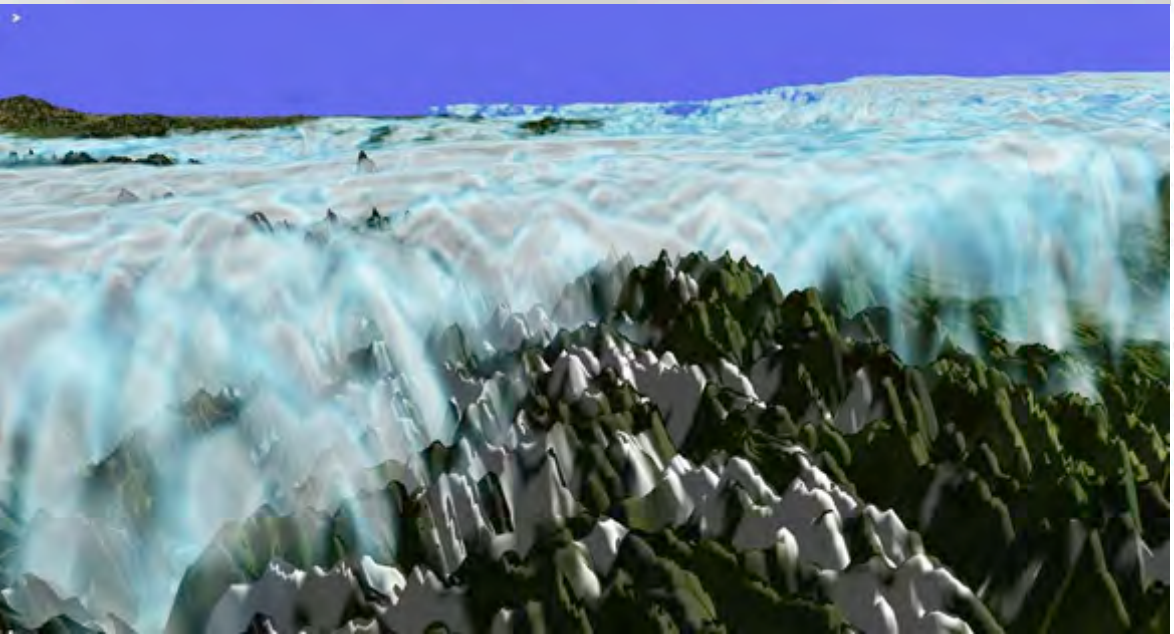

Papers with Highest Journal Impact Factor[1)]

Chemical Reviews
Impact Factor: 45.66

E. Brunk, U. Rothlisberger, Mixed quantum mechanical/molecular mechanical molecular dynamics simulations of biological systems in ground and electronically excited states, Chemical Reviews, DOI: 10.1021/cr500628b.

Nature
Impact Factor: 42.35

F. Joos, Global warming: growing feedback from ocean carbon to climate, Nature, DOI: 10.1038/522293a.

F. Miniati, A. Beresnyak, Self-similar energetics in large clusters of galaxies, Nature, DOI: 10.1038/Nature14552.

Nature Materials
Impact Factor: 36.43

G. Autès, A. Isaeva, L. Moreschini, J. C. Johannsen, A. Pisoni, R. Mori, W. Zhang, T. G. Filatova, A. N. Kuznetsov, L. Forró, W. Van den Broek, Y. Kim, K. S. Kim, A. Lanzara, J. D. Denlinger, E. Rotenberg, A. Bostwick, M. Grioni, O. V. Yazyev, A novel quasi-one-dimensional topological insulator in bismuth iodide β-Bi4I4, Nature Materials, DOI: 10.1038/nmat4488.

T. O. Thalaimar, L. Amis, F. A. Herzog, R. Schwartländer, J. Mastbergen, V. Vogel, Differential basal-to-apical accessibility of Lamin A/C epitopes in the nuclear lamina regulated by changes in cytoskeletal tension, Nature Materials, DOI: 10.1038/nmat4193.

Nature Nanotechnology
Impact Factor: 33.26

C. Becher, L. Maurel, U. Aschauer, M. Lilienblum, C. Magén, D. Meier, E. Langenberg, M. Trassin, J. Blasco, I. P. Krug, P. A. Algarabel, N. A. Spaldin, J. A. Pardo, M. Fiebig, Strain-induced coupling of electrical polarization and structural defects in SrMnO3 films, Nature Nanotechnology, DOI: 10.1038/nnano.2015.109.

Science
Impact Factor: 31.48

J. P. Gattuso, A. Magnan, R. Billé, W. W. L. Cheung, E. L. Howes, F. Joos, D. Allemand, L. Bopp, S. R. Cooley, C. M. Eakin, O. Hoegh-Guldberg, R. P. Kelly, H. O. Pörtner, A. D. Rogers, J. M. Baxter, D. Laffoley, D. Osborn, A. Rankovic, J. Rochette, U. R. Sumaila, S. Treyer, C. Turley, Contrasting futures for ocean and society from different anthropogenic CO2 emissions scenarios, Science, DOI: 10.1126/science.aac4722.

Nature Physics
Impact Factor: 20.60

F. Schulz, M. Ijäs, R. Drost, S. K. Hämäläinen, A. Harju, A. P. Seitsonen, P. Liljeroth, Many-body transitions in a single molecule visualized by scanning tunnelling microscopy, Nature Physics, DOI: 10.1038/NPHYS3212.

Energy & Environmental Science
Impact Factor: 15.49

C. Yi, J. Luo, S. Meloni, A. Boziki, N. Ashari-Astani, C. Grätzel, S. M. Zakeeruddin, U. Rothlisberger, M. Grätzel, Entropic stabilization of mixed A-cation ABX3 metal halide perovskites for high performance perovskite solar cells, Energy & Environmental Science, DOI: 10.1039/C5EE03255E.

Nature Climate Change
Impact Factor: 15.30

C. Schär, The worst heat waves to come, Nature Climate Change, DOI: 10.1038/nclimate2864.

Nano Letters
Impact Factor: 12.94

D. Alexeev, J. Chen, J. H. Walther, K. P. Giapis, P. Angelikopoulos, P. Koumoutsakos, Kapitza resistance between few-layer graphene and water: liquid layering effects, Nano Letters, DOI: 10.1021/acs.nanolett.5b03024.

Coordination Chemistry Reviews
Impact Factor: 12.10

Q. Sun, S. Mosquera-Vazquez, Y. Suffren, J. Hankache, N. Amstutz, L. M. Lawson Daku, E. Vauthey, A. Hauser, On the role of ligand-field states for the photophysical properties of Ruthenium(II)

Platform Consolidation: Enabling Workloads & Workflows Diversity

- Simulations using a range of programming paradigms
  - MPI, OpenMP, …
  - CUDA, OpenCL, OpenACC, …
  - Optimized libraries
  - Domain specific libraries
  - …
- In-situ and large-scale visualization tools
  - VisIt
  - ParaView
  - Application specific
- New domains (e.g. DL)

Two finalists in supercomputing

20.11.2015 | News
By: Simone Ulmer

Every year at the SC conference, the Gordon Bell Prize is presented to recognise outstanding work in the field of high-performance computing. This year, two research groups from ETH Zurich competed in the finals, a remarkable achievement considering only five teams have been preselected world-wide.

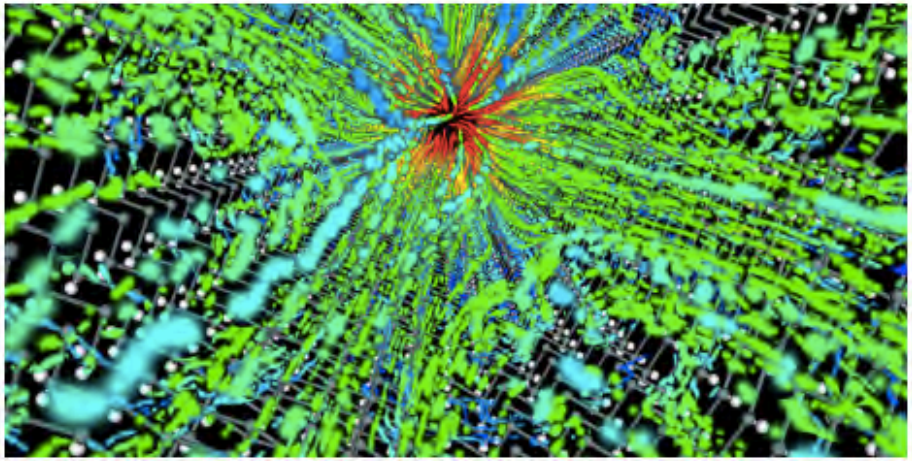

Unusual perspective: electrons flow inside a nano-transistor in direction of the viewer. (Graphic: Jean Favre / CSCS)

Continuous Co-design & Integration

- GPUs in Cray XK6/7 and XC environments have come a long way
  - GPUDirect and CUDA aware MPI
  - Quick CUDA releases and updates
  - Availability of complete toolchain & “user controlled” compute modes and features
- Potential with containers
  - Variety of workloads and workflows with other dependencies (e.g. OS, Python, etc.)
  - Data science applications & workloads
  - DL solutions and frameworks
- Emerging technologies from Nvidia readily available on Cray XC

Further Collaboration with Sites

- Strengthen partnership with Cray and Nvidia … and other CUG sites
  - Enable a richer experiences for users
- Take advantage of
  - Excellent interconnect with GPU aware MPI
  - Flexible programming and execution model
- Leverage & develop open-source solutions
- Nvidia DL SDK (accelerated frameworks)

The in silico lab-on-a-chip

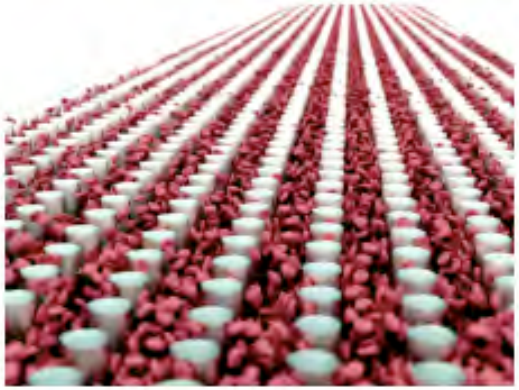

Simulated flow of red blood cells through a predefined structure. (Graphics: CSElab)

The second research team from ETH that has been nominated for the Gordon Bell Prize is led by ETH Professor Petros Koumoutsakos of the Computational Science & Engineering Laboratory → (CSElab). In collaboration with researchers from the Università della Svizzera italiana, the United States (Brown University, NVIDIA) and Italy (CNR and University of Rome), the team used the Titan supercomputer to run state-of-the-art simulations of the flow of micron-sized red blood cells and

Caffe
Chainer
DL4J
Deeplearning4j
MINERVA
mxnet
Purine
K
KERAS
Microsoft
CNTK
MatConvNet
TensorFlow
theano
torch