GPUs and the Future of Accelerated Computing

Emerging Technology Conference 2014
University of Manchester

John Ashley
Senior Solutions Architect
jashley@nvidia.com

Disclaimers

- These are my views, not NVIDIA's (although I do hope you all will agree by the time I'm done!) , and where I've quoted others, any misrepresentation of their views is mine and mine alone.
- There are "forward looking statements" in here. You are warned that my crystal ball is no better than anyone else's, even if I'm speaking about things like NVIDIA roadmaps, etc. It's a lawyer thing.
- Trademarks of other firms are their own, etc.

Agenda

- Who is NVIDIA?
- Why should I care about accelerated computing?
- What sort of difference can CUDA make?
- Where is NVIDIA going?

Who is NVIDIA?

NVIDIA products span the power-performance spectrum

GPUs Power World's 10 Greenest Supercomputers

Green500 Rank	MFLOPS/W	Site
1	4,503.17	GSIC Center, Tokyo Tech
2	3,631.86	Cambridge University
3	3,517.84	University of Tsukuba
4	3,185.91	Swiss National Supercomputing (CSCS)
5	3,130.95	ROMEO HPC Center
6	3,068.71	GSIC Center, Tokyo Tech
7	2,702.16	University of Arizona
8	2,629.10	Max-Planck
9	2,629.10	(Financial Institution)
10	2,358.69	CSIRO
37	1959.90	Intel Endeavor (top Xeon Phi cluster)
49	1247.57	Météo France (top CPU cluster)

GPUs are becoming…

GPUs are becoming…pervasive

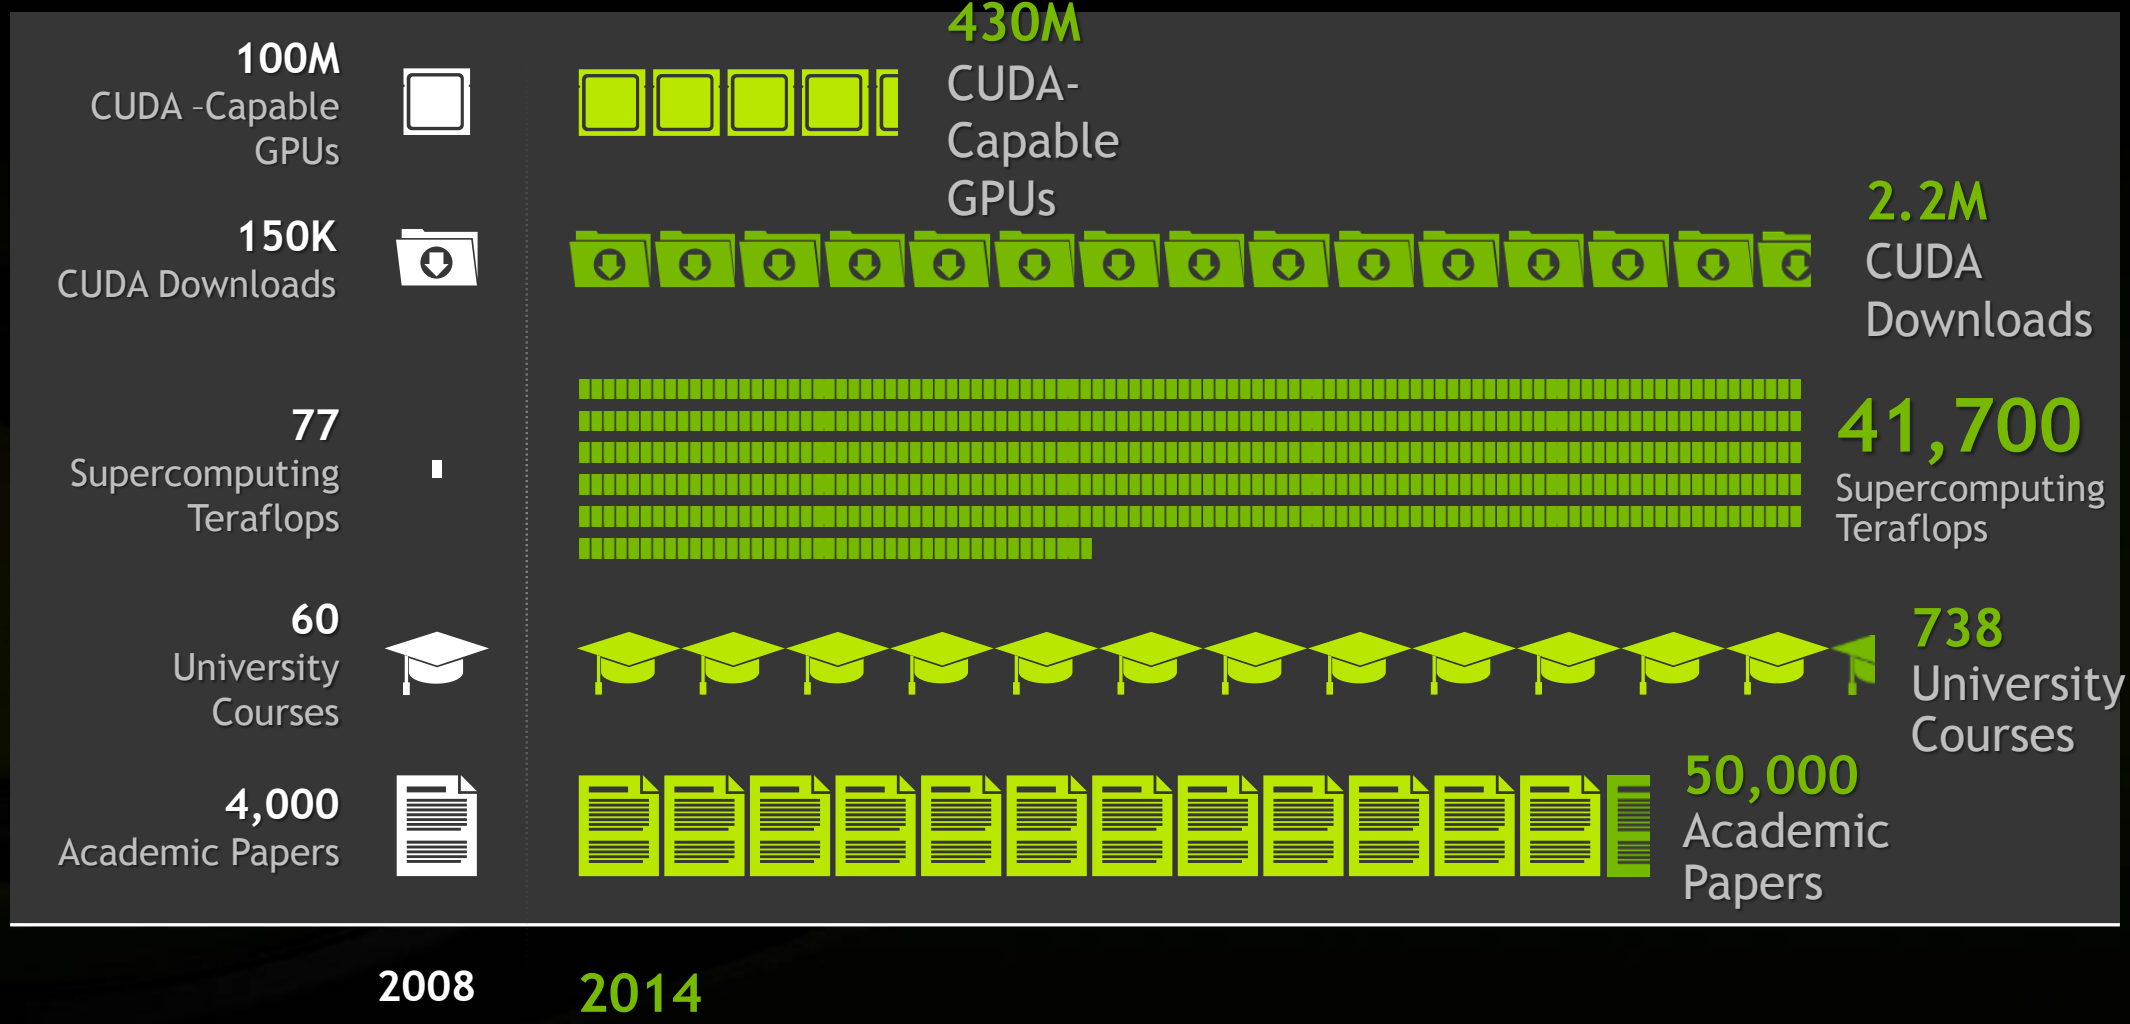

Why should I care about accelerated computing?

Moore's Law isn't what it used to be.

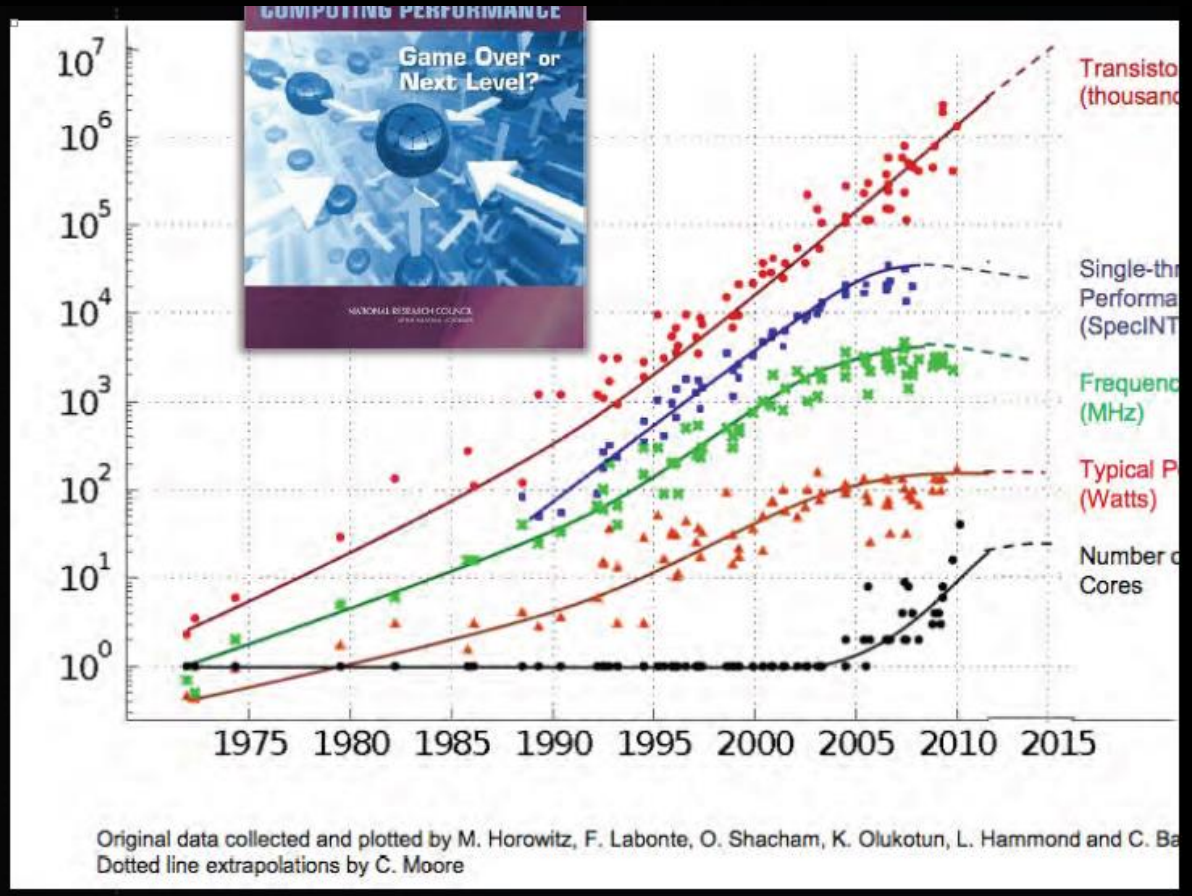

Moore's law is alive and well, but…

- Instruction-level parallelism (ILP) was mined out in 2001
- Voltage scaling (Dennard scaling) ended in 2005
- Most power is spent on communication

What does this mean to you?

The future is here today …

- Performance gains come from parallelism
- Systems are power limited (efficiency IS performance)
- Systems are communication limited (locality IS performance)

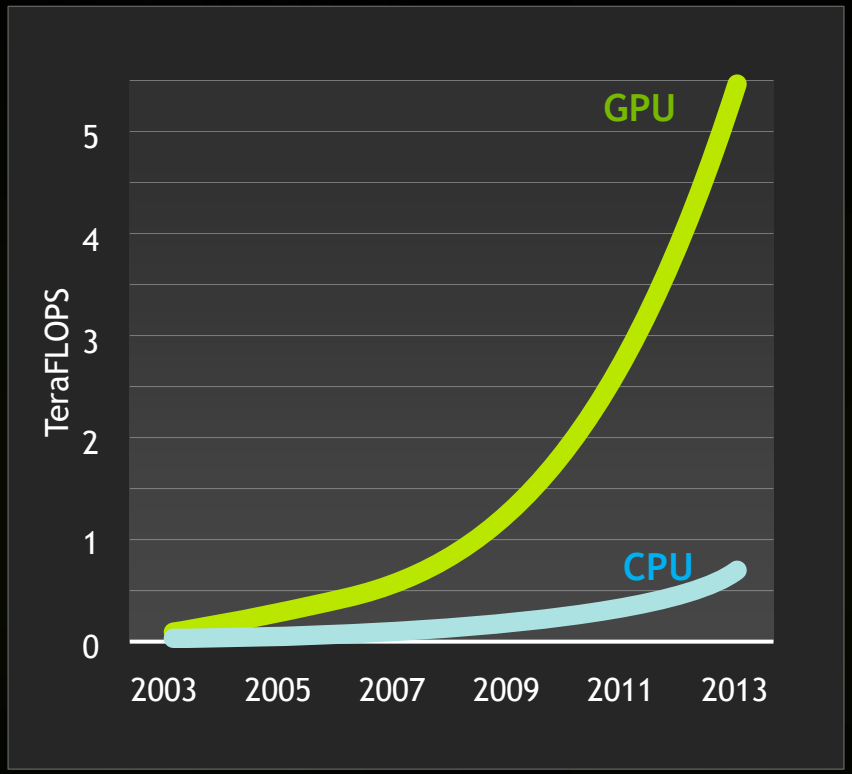

High Performance Computing is Going Hybrid

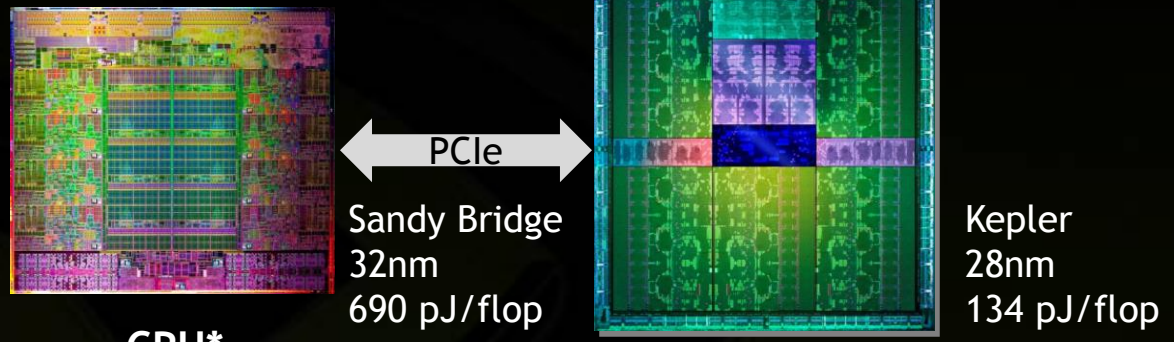

- Do most work by many cores optimized for **extreme energy efficiency**
- Still need a few cores optimized for **fast serial work**
- **Amdahl's law continues to apply to all architectures**

*x86, ARM, Power

What sort of difference can CUDA make?

CUDA Parallel Computing Platform

www.nvidia.com/getcuda

NVIDIA

Programming Approaches

- **Libraries** — "Drop-in" Acceleration
- **Directives** — Easily Accelerate Apps
- **Programming Languages** — Maximum Flexibility

Development Environment

Nsight IDE
Linux, Mac and Windows
GPU Debugging and Profiling

CUDA-GDB debugger
NVIDIA Visual Profiler

Open Compiler Tool Chain

Enables compiling new languages to CUDA platform, and CUDA languages to other architectures

Hardware Capabilities

SMX

Dynamic Parallelism

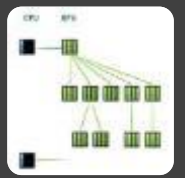

HyperQ

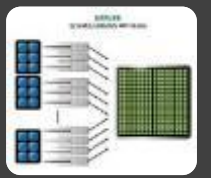

GPUDirect

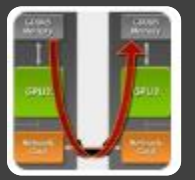

Artificial Neural Network at a Fraction of the Cost with GPUs

"*Now You Can Build Google's $1M Artificial Brain on the Cheap*"

-Wired

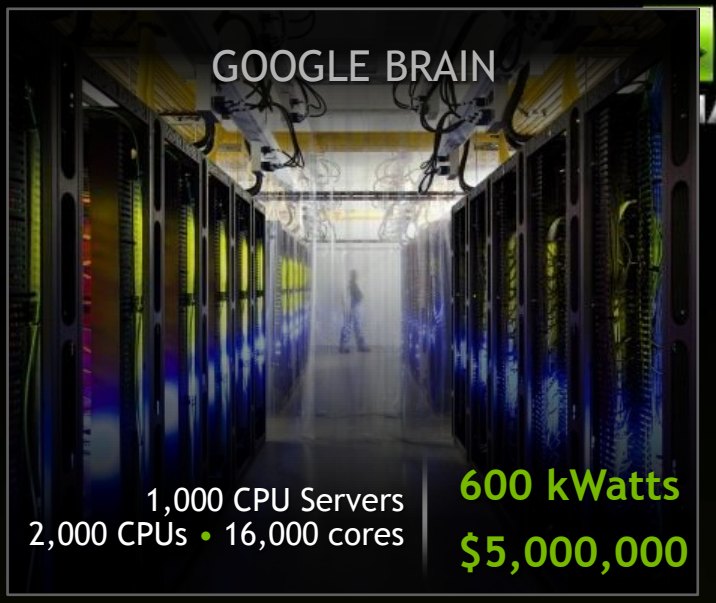

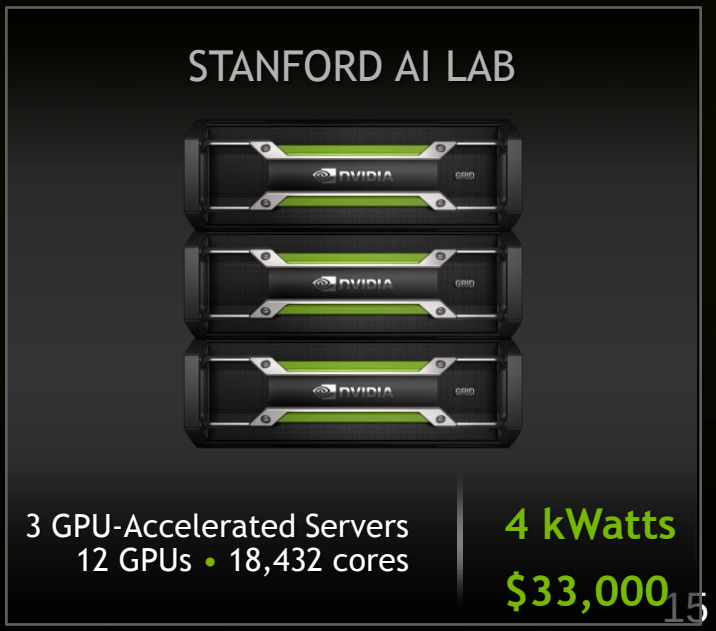

VISIONWORKS

COMPUTER VISION ON CUDA

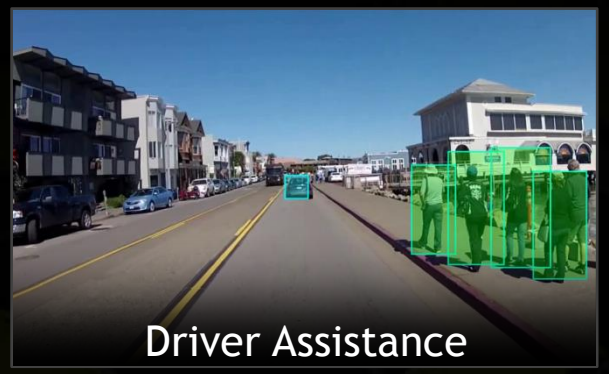

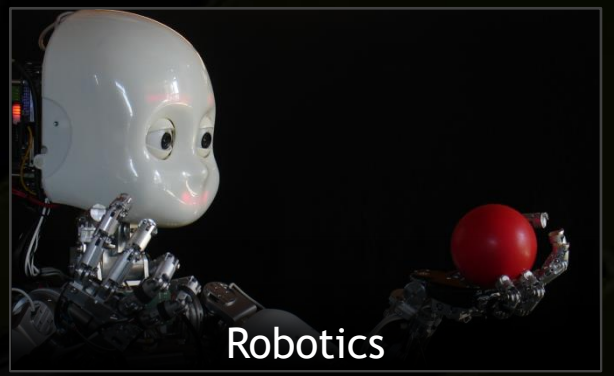

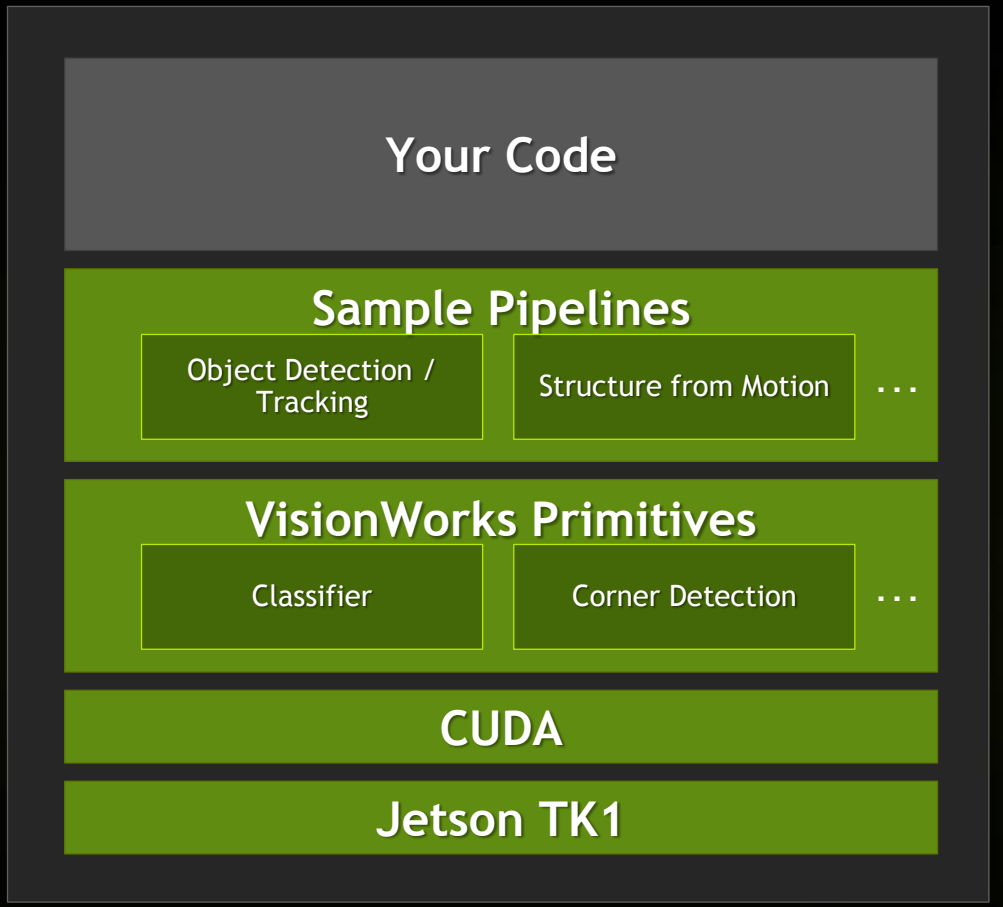

Computer Vision on CUDA

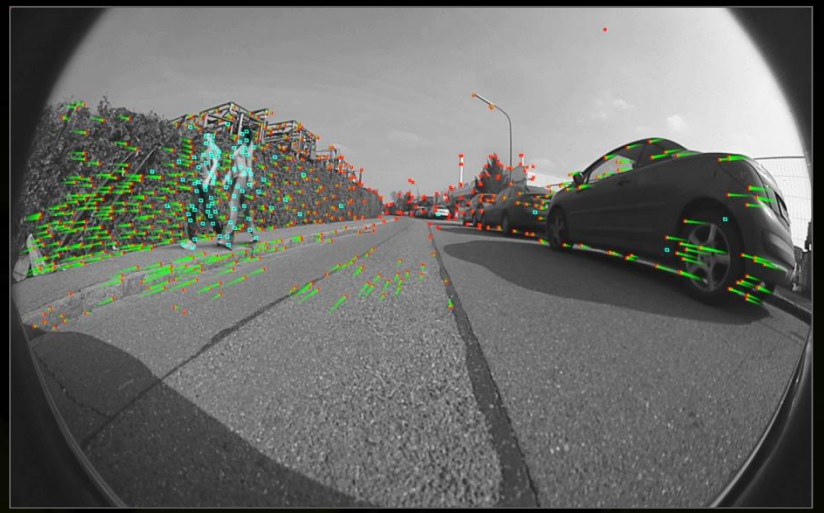

Feature Detection / Tracking

~30 GFLOPS @ 30 Hz

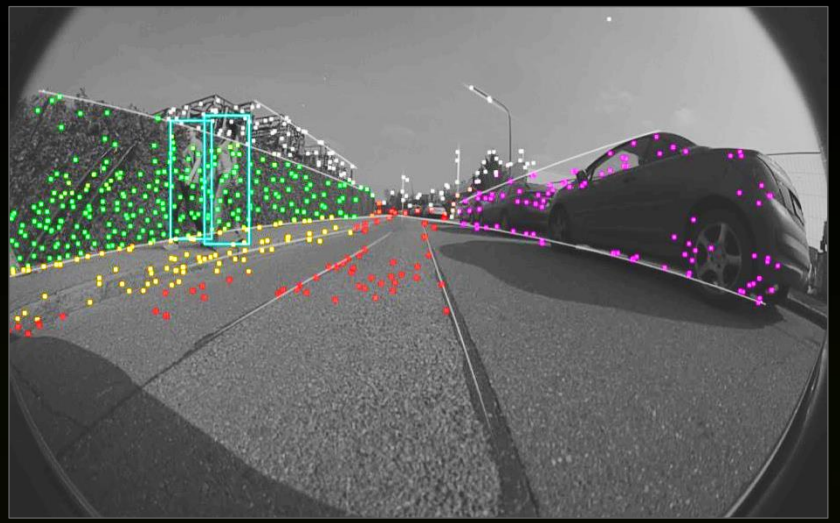

3D Scene Interpretation

~280 GFLOPS @ 30 Hz

Audi Self Driving Car Before & After CUDA

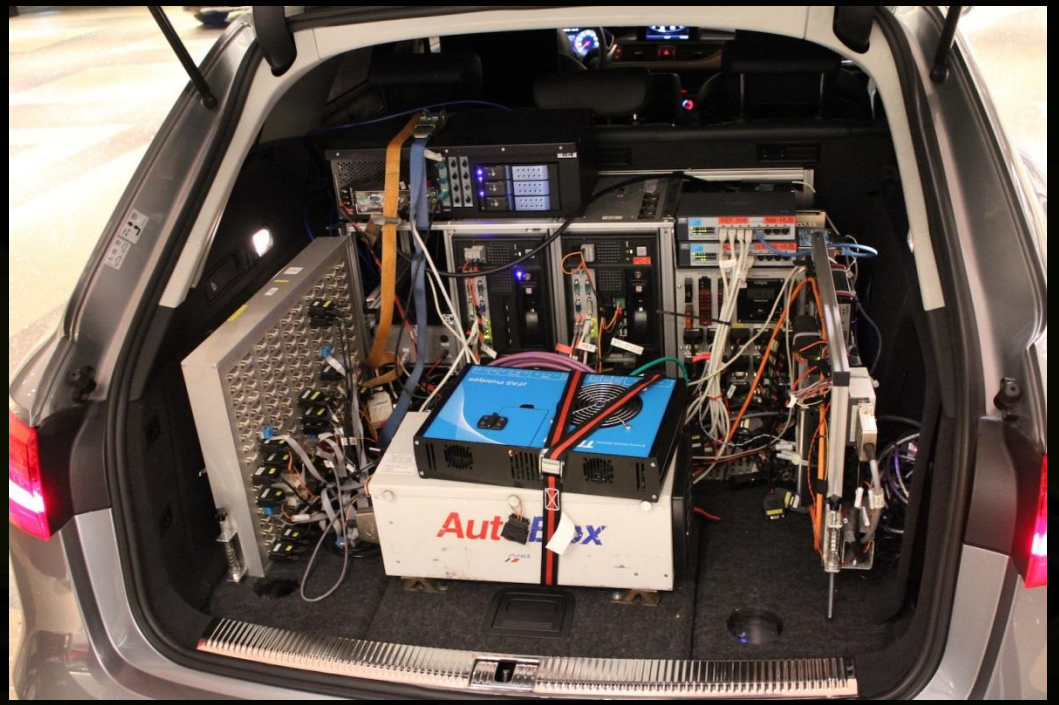

Solid Growth of GPU Accelerated Apps

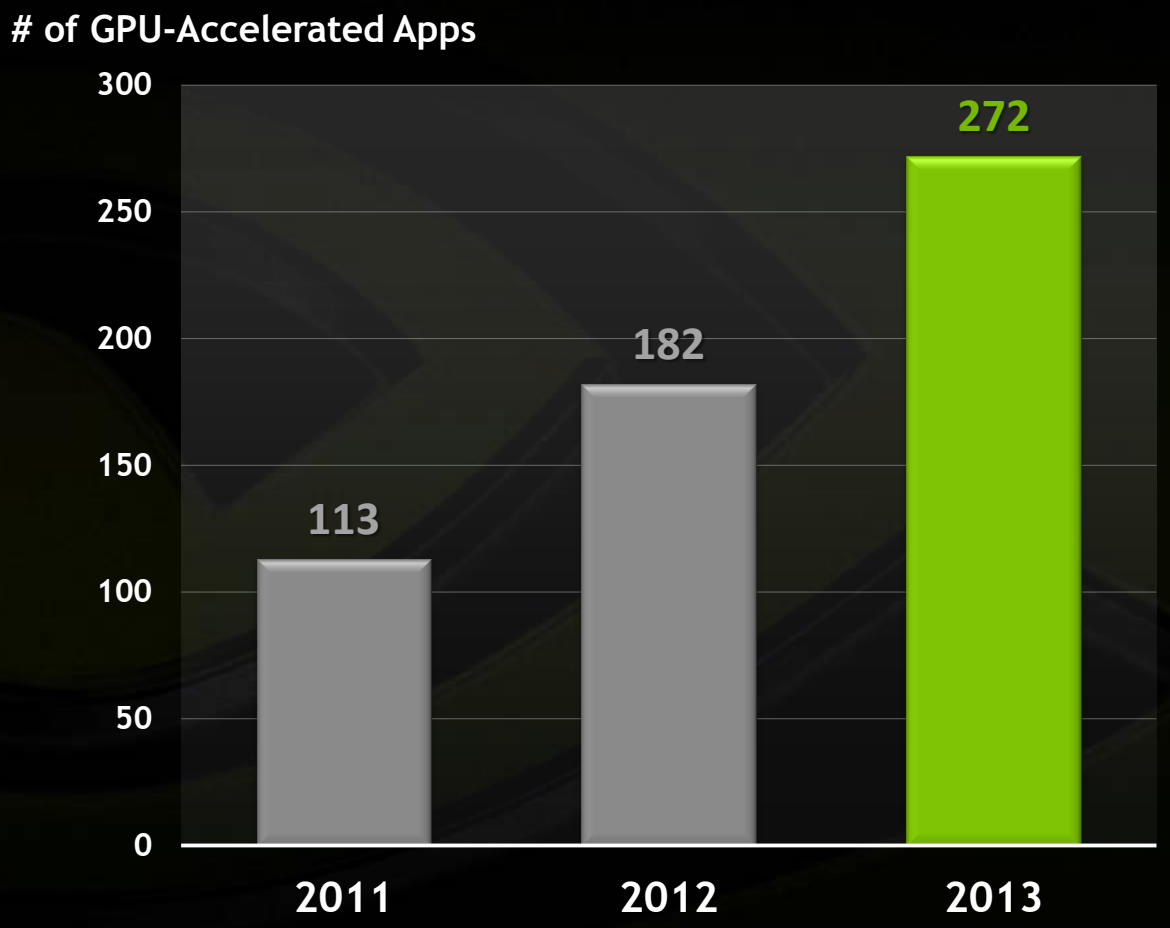

Top HPC Applications

Molecular Dynamics	AMBER CHARMM DESMOND	GROMACS LAMMPS NAMD
Quantum Chemistry	Abinit Gaussian	GAMESS NWChem
Material Science	CP2K QMCPACK	Quantum Espresso VASP
Weather & Climate	COSMO GEOS-5 HOMME	CAM-SE NEMO NIM WRF
Lattice QCD	Chroma	MILC
Plasma Physics	GTC	GTS
Structural Mechanics	ANSYS Mechanical LS-DYNA Implicit MSC Nastran	OptiStruct Abaqus/Standard
Fluid Dynamics	ANSYS Fluent	Culises (OpenFOAM)

Accelerated, In Development


NVIDIA
POPULAR GPU-ACCELERATED APPLICATIONS
02 Research: Higher Education and Supercomputing
06 Defense and Intelligence
07 Computational Finance
08 Manufacturing: CAD and CAE
10 Media and Entertainment
14 Oil and Gas
Research: Higher Education and Supercomputing
COMPUTATIONAL CHEMISTRY AND BIOLOGY
Bioinformatics
Molecular Dynamics
272 GPU-Accelerated Applications
www.nvidia.com/appscatalog

Top Applications Now with Built-in GPU Support

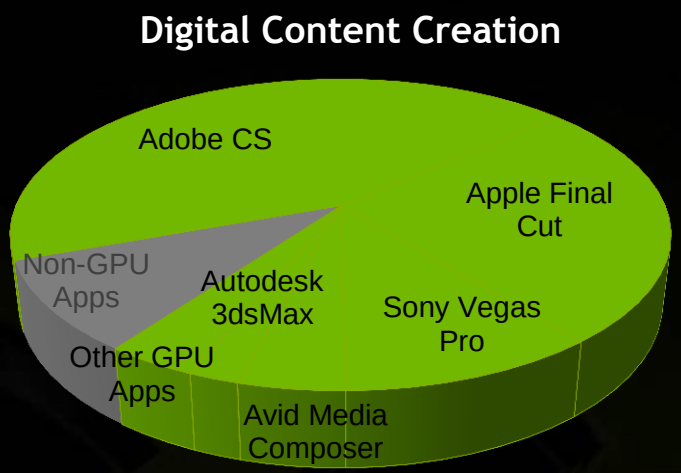

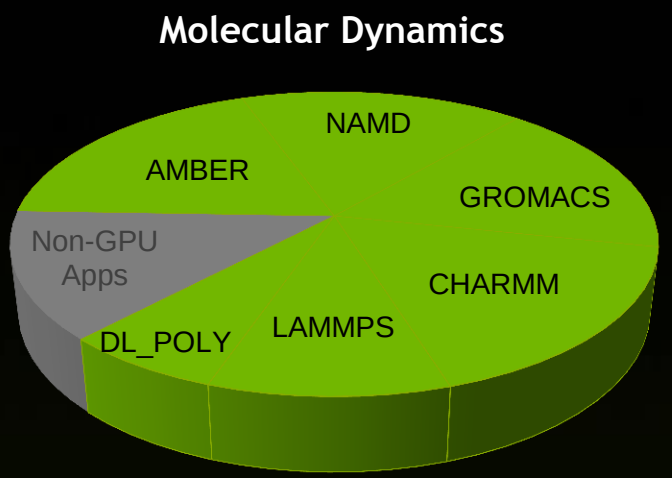

Application Market Share by Segment

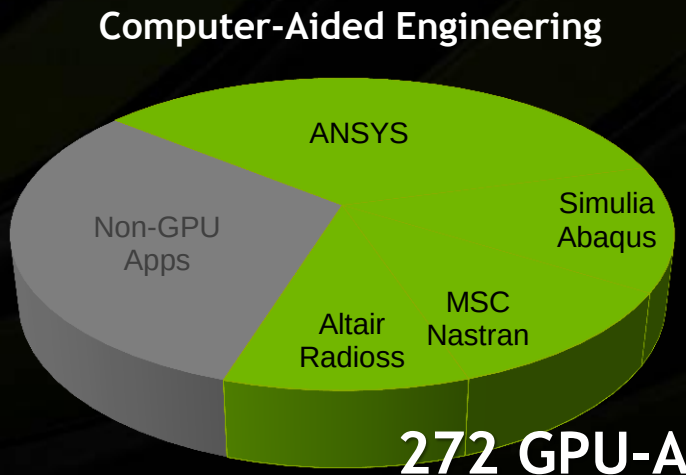

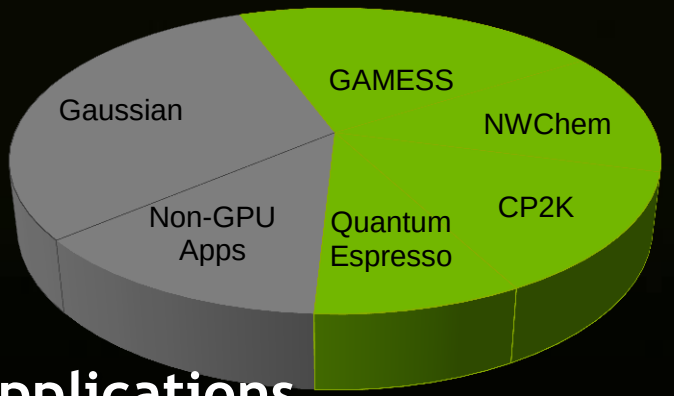

272 GPU-Accelerated Applications

www.nvidia.com/appscatalog

Performance on Leading Scientific Applications

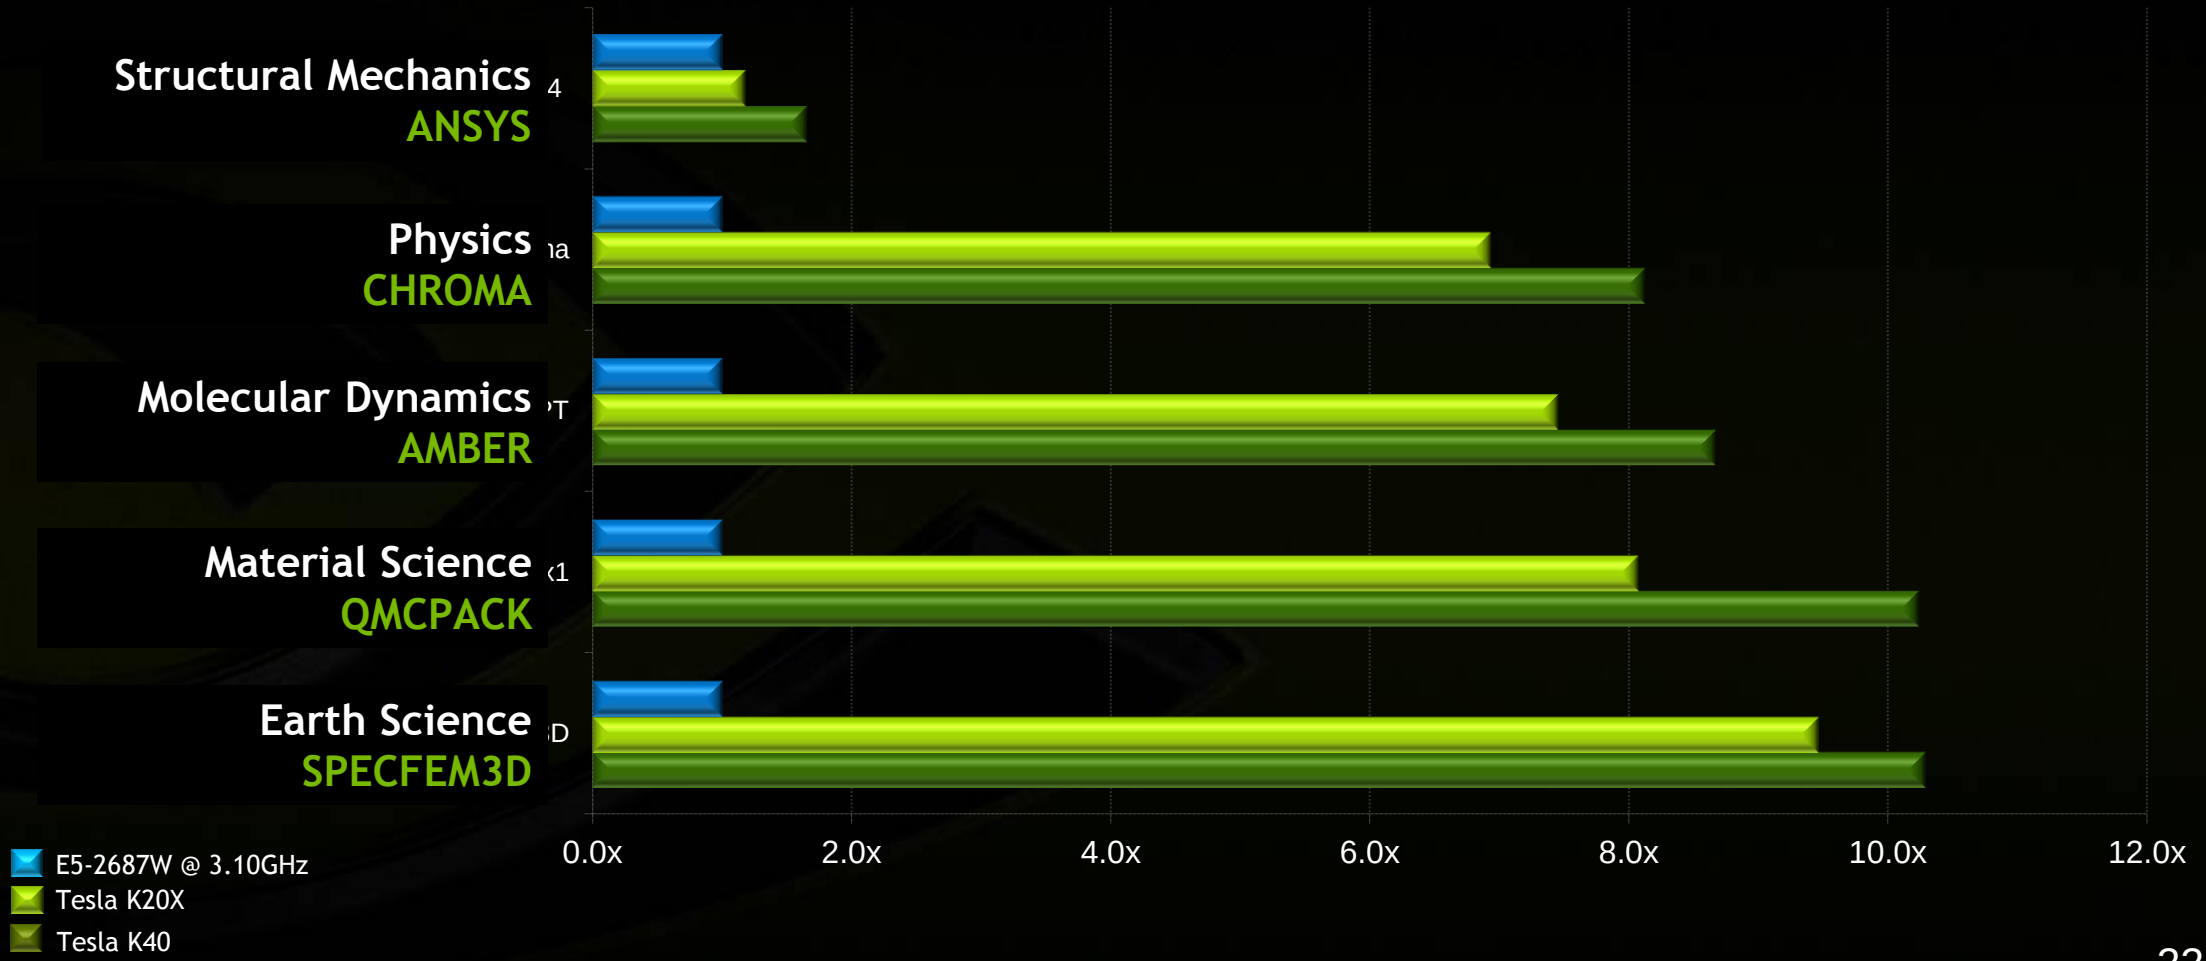

Where is NVIDIA going?

Fast Paced CUDA GPU Roadmap

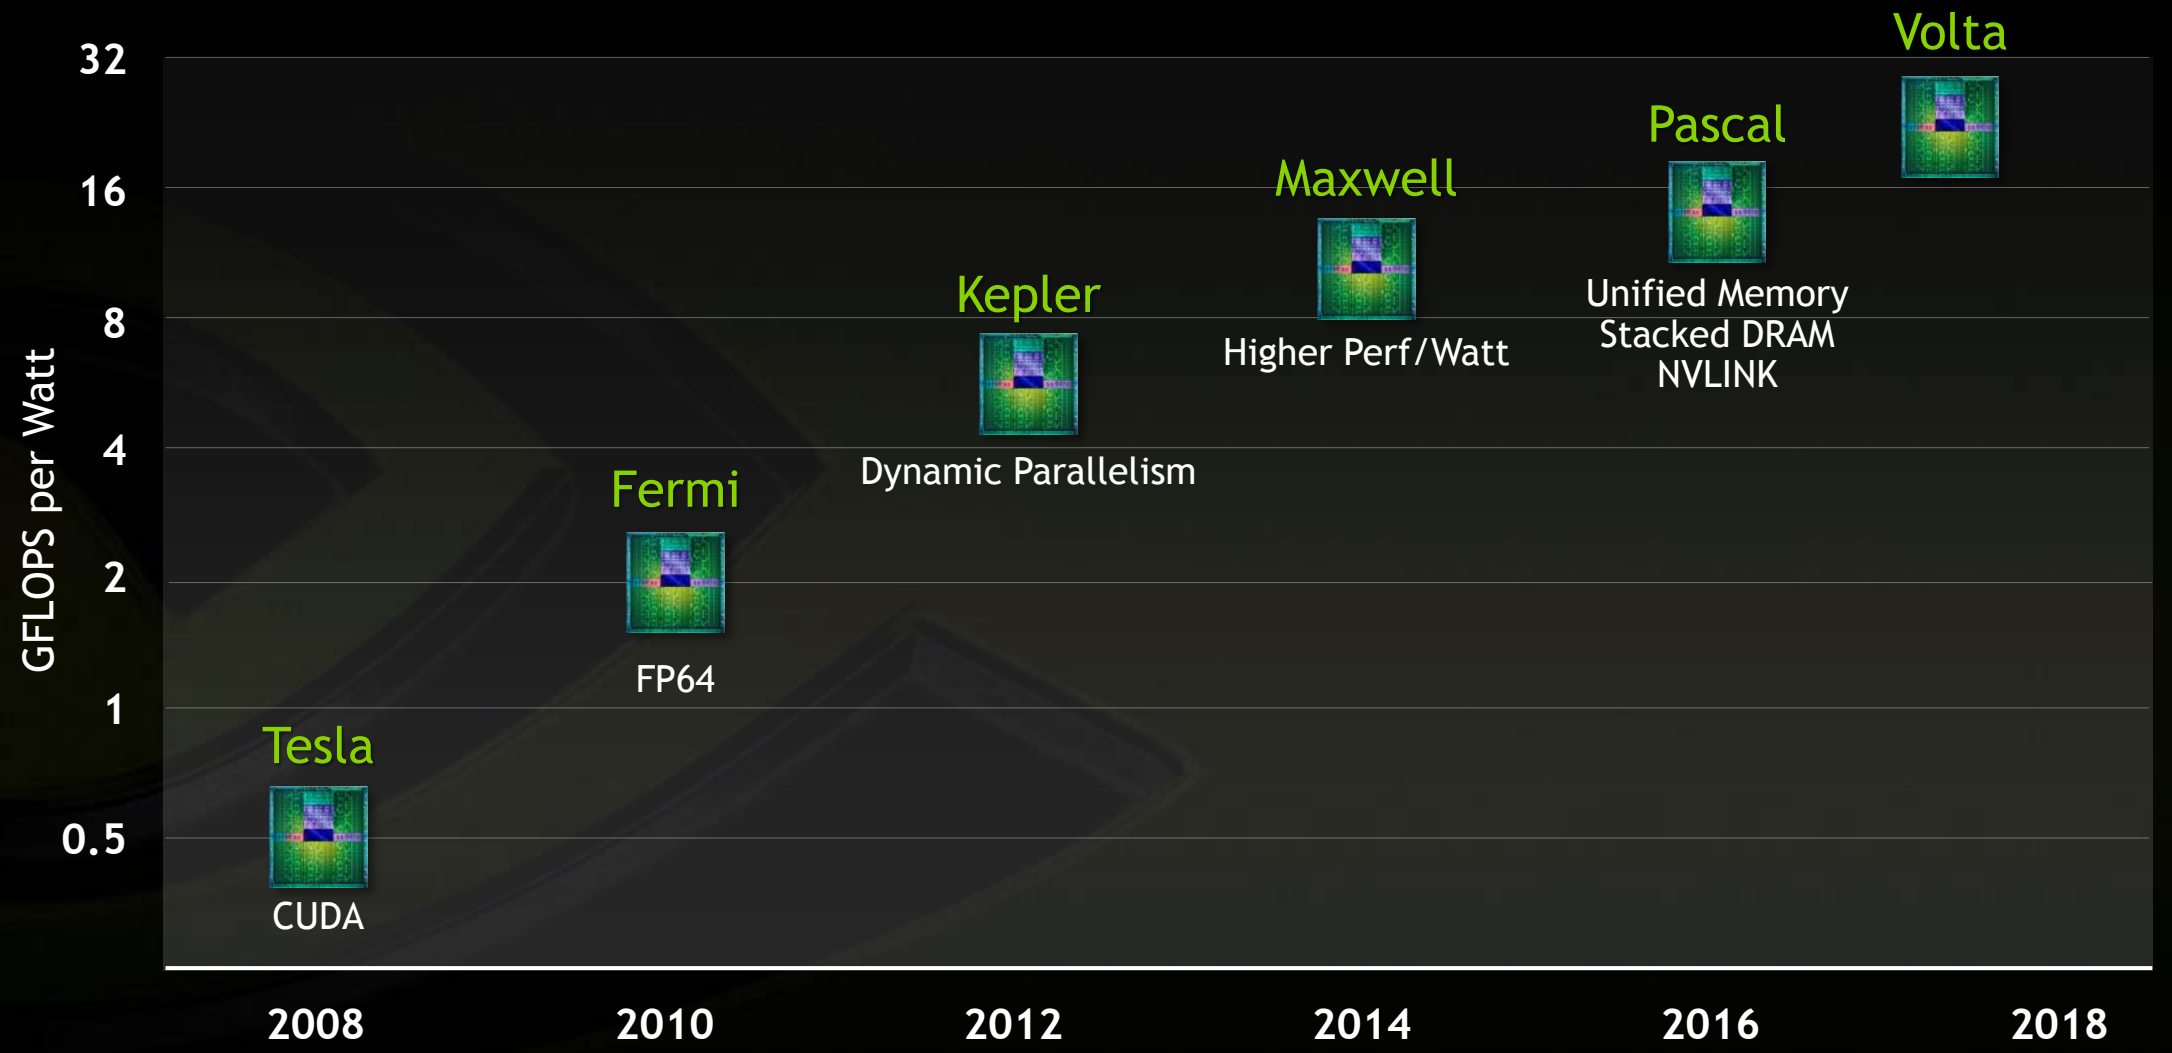

Introducing NVLINK and Stacked Memory

NVLINK

- GPU high speed interconnect
- 5-12x PCIe Gen 3 Bandwidth
- Drastically reduced energy/bit

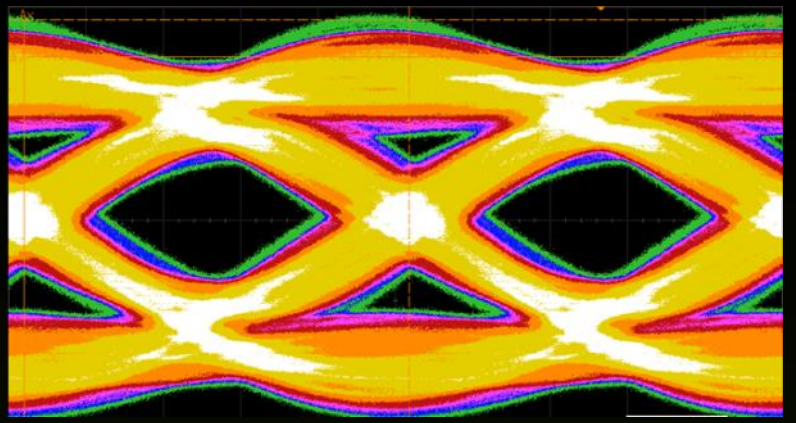

Stacked Memory

- 2-4x Capacity & Bandwidth
- 3-4x More Energy Efficient per bit
- Leaves more power for compute

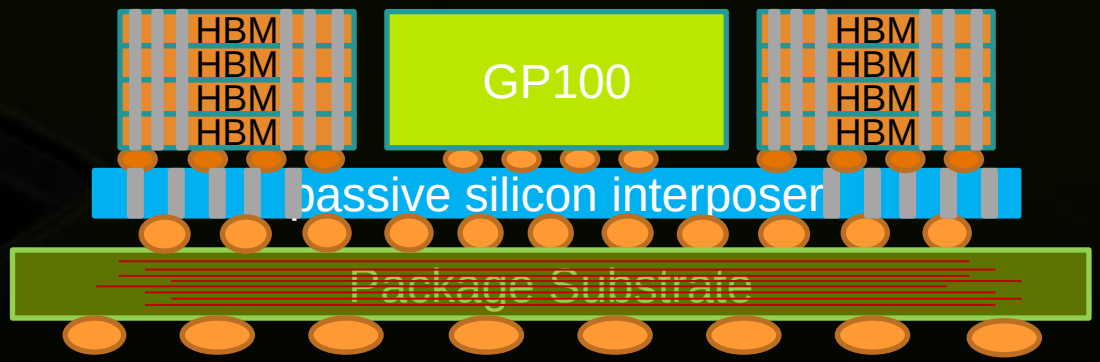

IBM Partners with NVIDIA to Build Next-Generation Supercomputers

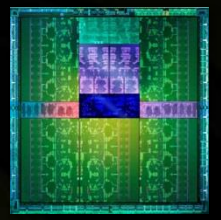

Tesla
GPU

+

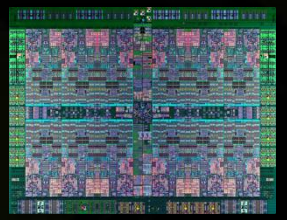

POWER 8
CPU

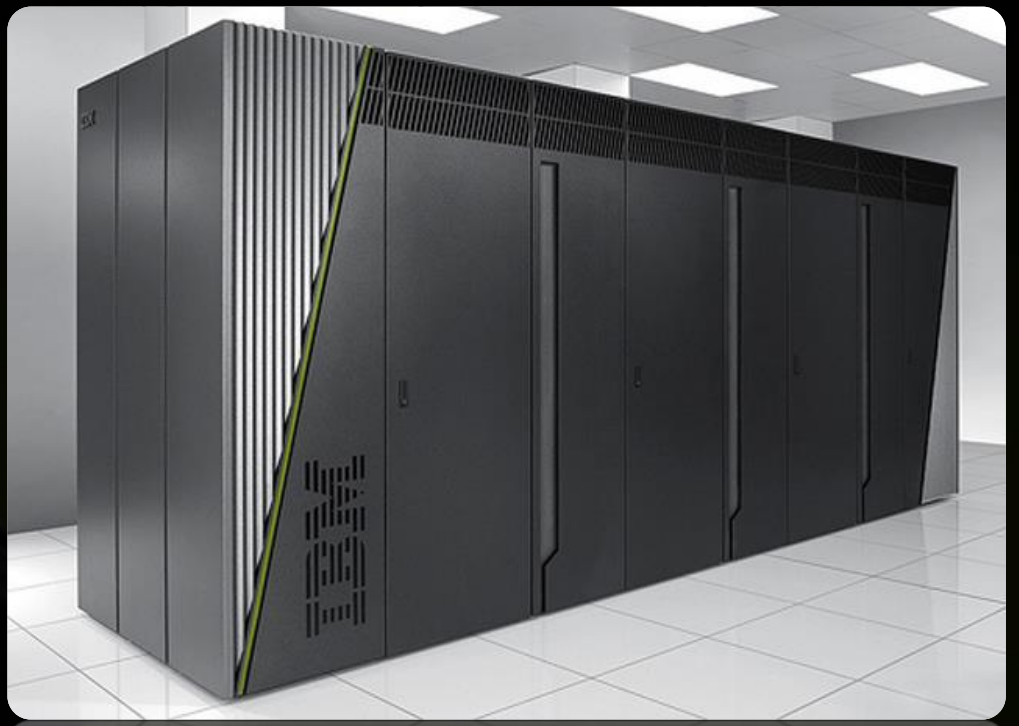

GPU-Accelerated POWER-Based Systems Available in 2014

Unify GPU and Tegra Architecture

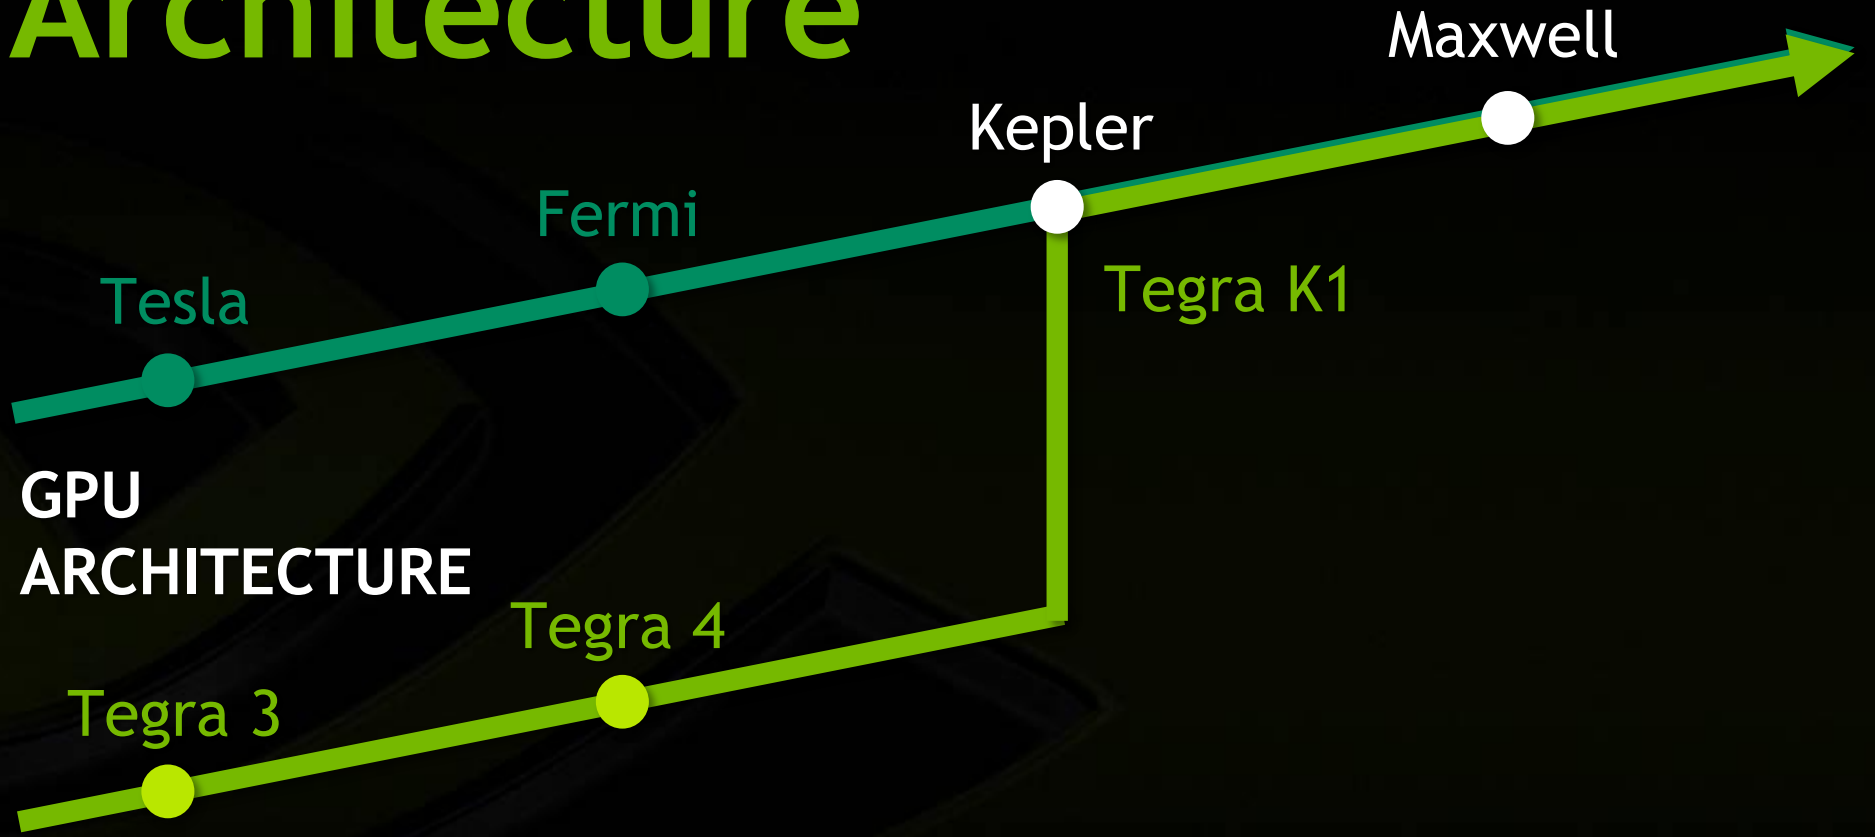

Jetson TK1 Development Kit

- 32 Bit ARM+Kepler SMX SOC
- CUDA capable
- $192 from US retailers, cost and availability will vary elsewhere
- https://developer.nvidia.com/jetson-tk1
- http://devblogs.nvidia.com/parallelforall/jetson-tk1-mobile-embedded-supercomputer-cuda-everywhere/

Summary

- **Who is NVIDIA?**
  A: The world's leading Visual Computing company, from consumer devices through to world class supercomputers

- **Why should I care about accelerated computing?**
  A: Because parallelism and heterogeneous computing is the future of big compute and big data

- **What sort of difference can CUDA make?**
  A: Order of magnitude improvements in performance and efficiency are possible with CUDA

- **Where is NVIDIA going?**
  A: Relentlessly forward to a CUDA enabled parallel & energy efficient computing future

Questions ?