MedeA instrument integrated into HPC Clouds and how this benefits LAMMPS users worldwide

T.L.Thomas, UNM

Instruments like the
MedeA instrument integrated into HPC Clouds
and how this benefits LAMMPS users worldwide

A Vision:

Instruments like the
MedeA instrument integrated into HPC Clouds
and how this benefits users worldwide

MedeA instrument

what's an Instrument?

integrated into HPC Clouds

what's a Cloud?

and how this benefits LAMMPS users worldwide

Intel Paragon

risc 1 of 1

From Wikipedia, the free encyclopedia

This article includes a list of references, related reading or external links, but **its sources remain unclear because it lacks inline citations**. Please improve this article by introducing more precise citations. *(March 2013)*

This article **relies on references to primary sources**. Please add references to secondary or tertiary sources. *(March 2013)*

The **Intel Paragon** was a series of massively parallel supercomputers produced by Intel. The **Paragon XP/S** was a productized version of the experimental *Touchstone Delta* system built at Caltech, launched in 1992. The Paragon superseded Intel's earlier iPSC/860 system, to which it was closely related.

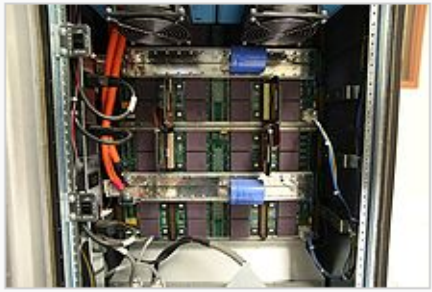

Mesh interconnect on XP-E cabinet

The Paragon series was based around the Intel i860 RISC microprocessor. Up to 2048 (later, up to 4000) i860s were connected in a 2D grid. In 1993, an entry-level **Paragon XP/E** variant was announced with up to 32 compute nodes. The system architecture was a partitioned system, with the majority of the system comprising diskless compute nodes and a small number of I/O nodes interactive service nodes. Since the bulk of the nodes had no permanent storage, it was possible to "Red/Black switch" the compute partition from classified to unclassified by disconnecting one set of I/O nodes with classified disks and then connecting an unclassified I/O partition.

Intel Paragon XP-E single cabinet system
Cats

Intel intended the Paragon to run the OSF/1 AD distributed operating system on all processors. However, this was found to be inefficient in practice, and a light-weight kernel called SUNMOS was developed at Sandia National Laboratories to replace OSF/1 AD on the Paragon's compute processors.

The prototype for the Intel Paragon was the Intel Delta, built by Intel with funding from DARPA and installed operationally at the

ASCI Red

From Wikipedia, the free encyclopedia

ASCI Red (also known as **ASCI Option Red** or **TFLOPS**) was the first computer built under the Accelerated Strategic Computing Initiative (ASCI),[5][6] the supercomputing initiative of the United States government created to help the maintenance of the United States nuclear arsenal after the 1992 moratorium on nuclear testing.

ASCI Red was built by Intel and installed at Sandia National Laboratories in late 1996. The design was based on the Intel Paragon computer. The original goals to deliver a true teraflop machine by the end of 1996 that would be capable of running an ASCI application using all memory and nodes by September 1997 were met.[7] It was used by the US government from the years of 1997 to 2005 and was the world's fastest supercomputer until late 2000.[4][6] It was the first ASCI machine that the Department of Energy acquired,[6] and also the first supercomputer to score above one teraflops on the LINPACK benchmark, a test that measures a computer's calculation speed. Later upgrades to ASCI Red allowed it to perform above two teraflops.

ASCI Red earned a reputation for reliability that some veterans say has never been beat. Sandia director Bill Camp said that ASCI Red had the best reliability of any supercomputer ever built, and "was supercomputing's high-water mark in longevity, price, and performance." [8]

ASCI Red was decommissioned in 2006.[2]

ASCI Red

Active	Two-Thirds Operational March 1997, Fully Operational June 1997,[1] decommissioned 2006[2]
Sponsors	Intel Corporation[1]
Operators	Sandia National Laboratories, US Department of Energy
Location	Sandia National Laboratories, United States
Power	850 kW
Space	1,600 sq ft (150 m^2)[3]
Memory	1212 gigabytes
Speed	1.3 teraflops (peak)[1]
Ranking	TOP500: 1, June 2000[4]
Purpose	nuclear materials testing, other
Legacy	First Supercomputer to achieve over 1.0 teraflops on LINPACK test

Contents [hide]

1 System Structure
2 Technical Specifications
3 First to TeraFlop
4 Operating System
5 References
6 External links

Intel, the ASCI Red Computer is also the first large scale supercomputer to components.[10]

Cougar 1 of 3

All of ASCI Red's partitions are interconnected to form one supercomputer, however at the same time none of the nodes support global shared memory. Each of the nodes works in its own memory, and each shares data with the others through "explicit message-passing".[9]

Technical Specifications [edit source | edit beta]

The Computer itself took up almost 1600 square feet of space,[3] and is made up of 104 "cabinets". Of those cabinets, 76 are computers (processors), 8 are switches, and 20 are disks. It has a total of 1212 GB of RAM, and 9298 separate processors. The original machine used Intel Pentium Pro processors each clocked at 200 MHz. These were later upgraded to specially packaged Pentium II Xeon processors, each clocked at 333 MHz. Overall, it required 850 kW of power (not including air conditioning). What sets ASCI Option Red aside from all of its predecessors in supercomputing is its high I/O bandwidth. Previous supercomputers had multi-gigaflops performance, yet their slow I/O speeds would slow down, or bottleneck the systems. Intel's TFLOPS PFS is an extremely efficient "Parallel File System" that can sustain transfer speeds of up to 1 GB/s, eliminating bottlenecks.[11]

First to TeraFlop [edit source | edit beta]

In December, 1996, three quarters of ASCI Red was measured at a world record 1.06 TFLOPS on MP LINPACK and held the record for fastest supercomputer in the world for several consecutive years, maxing out at 2.38 TFLOPS after a processor and memory upgrade in 1999.[7][12] The system utilized Pentium Pro processors when initially constructed and when it recorded performance above one TeraFlop. In that configuration, when fully built it recorded 1.6 TeraFlops of performance. Upgrades later in 1999, to specially packaged Pentium II Xeon processors, pushed performance to 3.1 TeraFlops.[8]

Operating System [edit source | edit beta]

The different partitions of ASCI Red run on different Operating Systems. For example, users of the computer work in an environment called "Teraflops OS", an operating system (once called Paragon OS) that was originally developed for the Intel Paragon XP/S Supercomputer.[5] ASCI Red's Compute partition runs on an operating system named Cougar.[9] Cougar is a Sandia Labs and University of New Mexico collaboration; it is a lightweight OS based on PUMA and SUNMOS, two systems that were also designed for use on the Paragon supercomputer.[9] It consists of a light weight kernel, the Process Control Thread, and other utilities and libraries. The Linux 2.4 kernel was ported to the system and a custom CNIC driver was written, but the heavy weight OS did not perform as well as the Cougar lightweight kernel on many benchmarks.[9]

References [edit source | edit beta]

List	Rank	System	Vendor	Total Cores	Rmax (GFlops)	Rpeak (GFlops)	Power (kW)
11/2005	276	ASCI Red	Intel	9,632	2,379.00	3,207.00	
06/2005	139	ASCI Red	Intel	9,632	2,379.00	3,207.00	
11/2004	80	ASCI Red	Intel	9,632	2,379.00	3,207.00	
06/2004	61	ASCI Red	Intel	9,632	2,379.00	3,207.00	
11/2003	27	ASCI Red	Intel	9,632	2,379.00	3,207.00	
06/2003	17	ASCI Red	Intel	9,632	2,379.00	3,207.00	
11/2002	15	ASCI Red	Intel	9,632	2,379.00	3,207.00	
06/2002	7	ASCI Red	Intel	9,632	2,379.00	3,207.00	
11/2001	4	ASCI Red	Intel	9,632	2,379.00	3,207.00	
06/2001	3	ASCI Red	Intel	9,632	2,379.00	3,207.00	
11/2000	2	ASCI Red	Intel	9,632	2,379.00	3,207.00	
06/2000	1	ASCI Red	Intel	9,632	2,379.00	3,207.00	
11/1999	1	ASCI	Intel	9,632	2,379.00	3,207.00	

en.wikipedia.org/wiki/File:Supercomputers.png

Other bookmarks

File Talk

Search

File:Supercomputers.png

From Wikipedia, the free encyclopedia

File File history File usage Global file usage

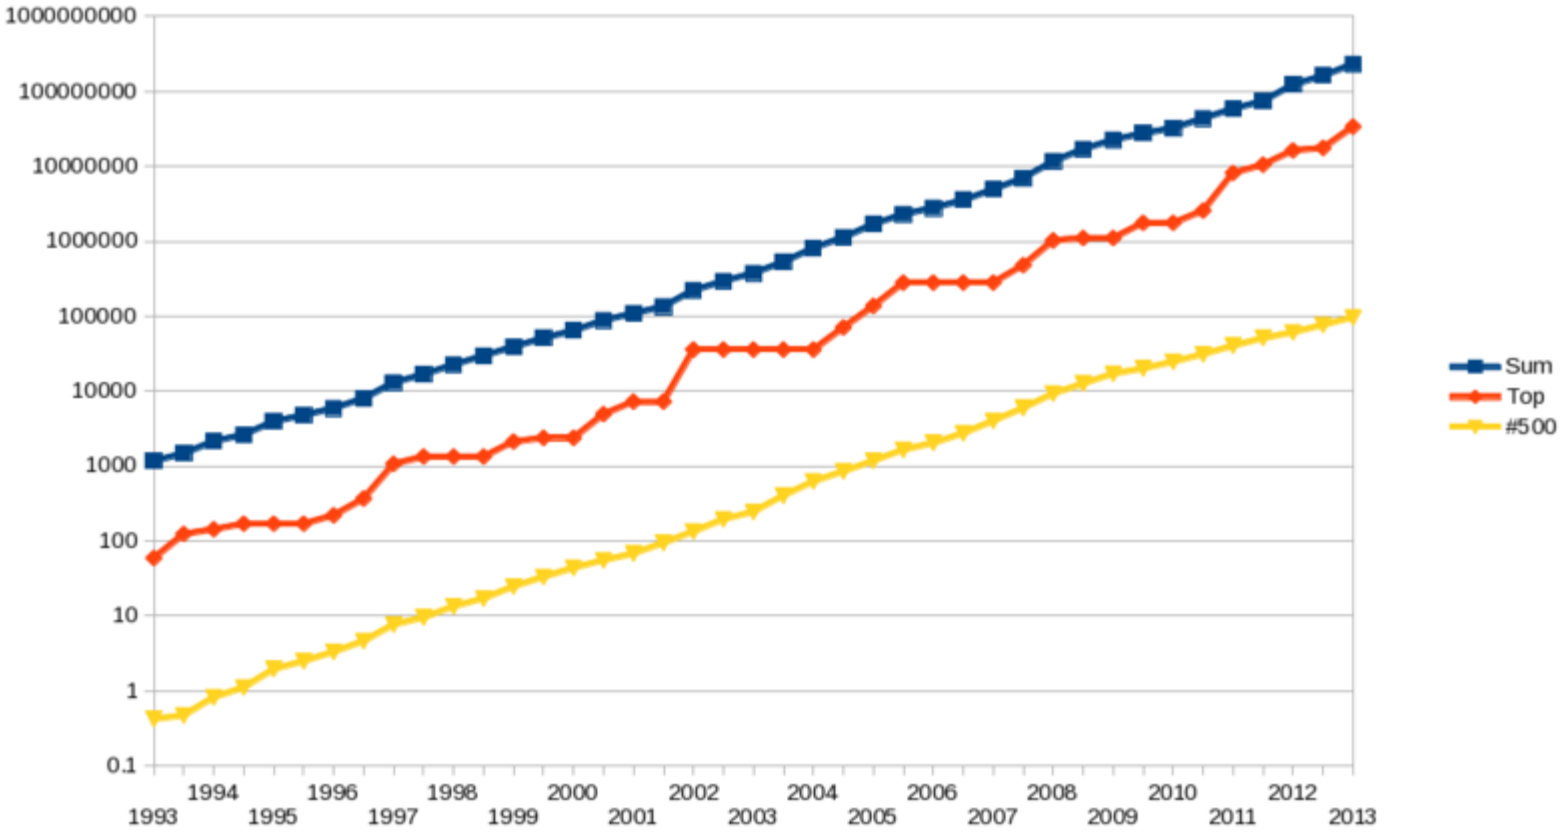


IBM
Blue Gene supercomputer

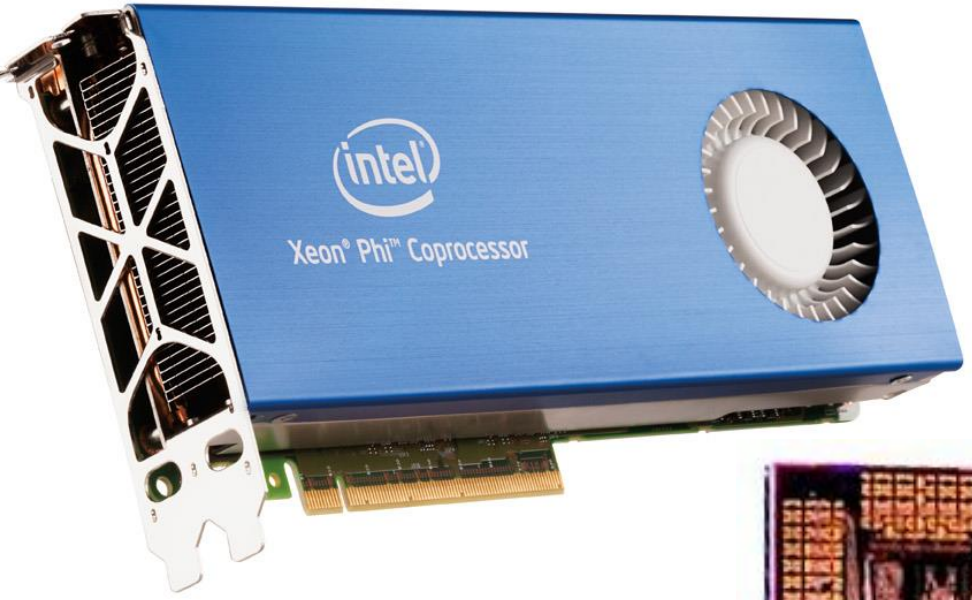
intel
Xeon® Phi™ Coprocessor

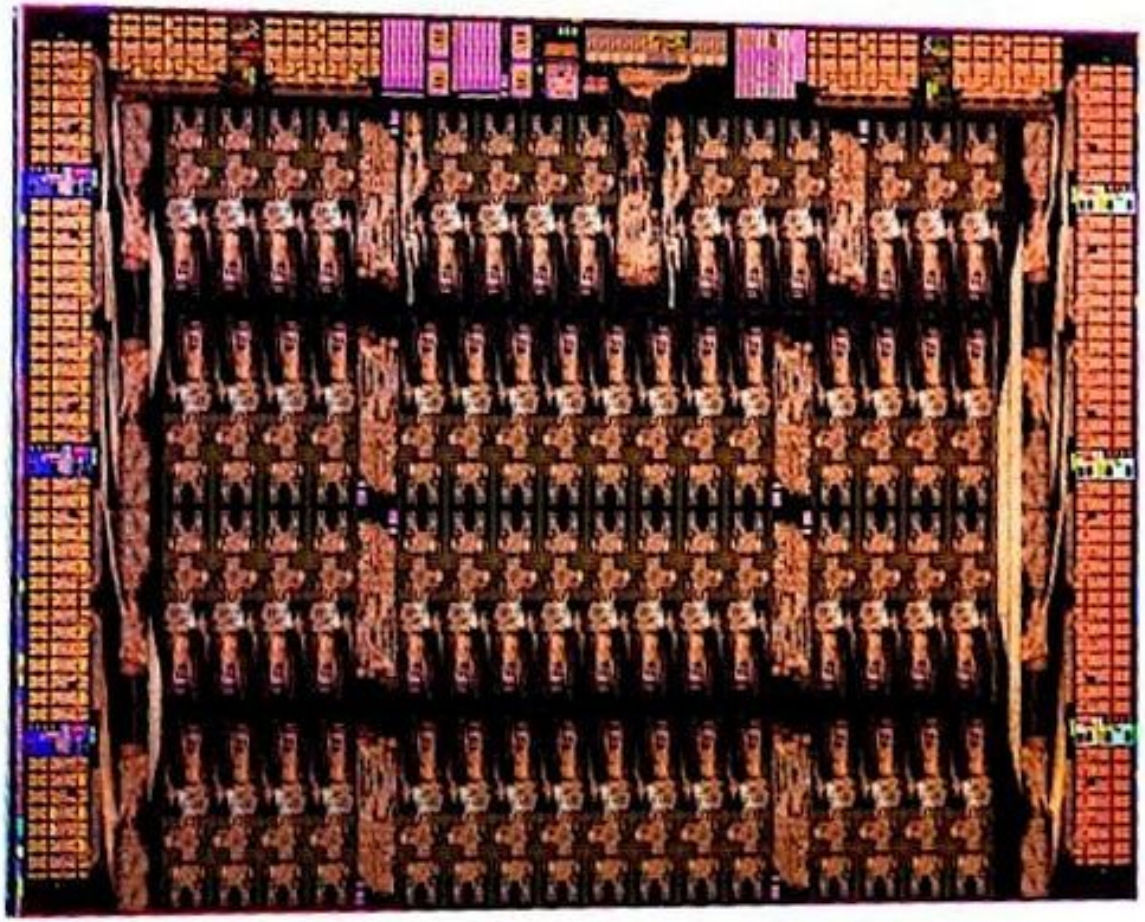

Titan: ~~World's Fastest~~ Supercomputer

18,688 Tesla K20X GPUs

27 Petaflops Peak: 90% of Performance from GPUs

17.59 Petaflops Sustained Performance on Linpack

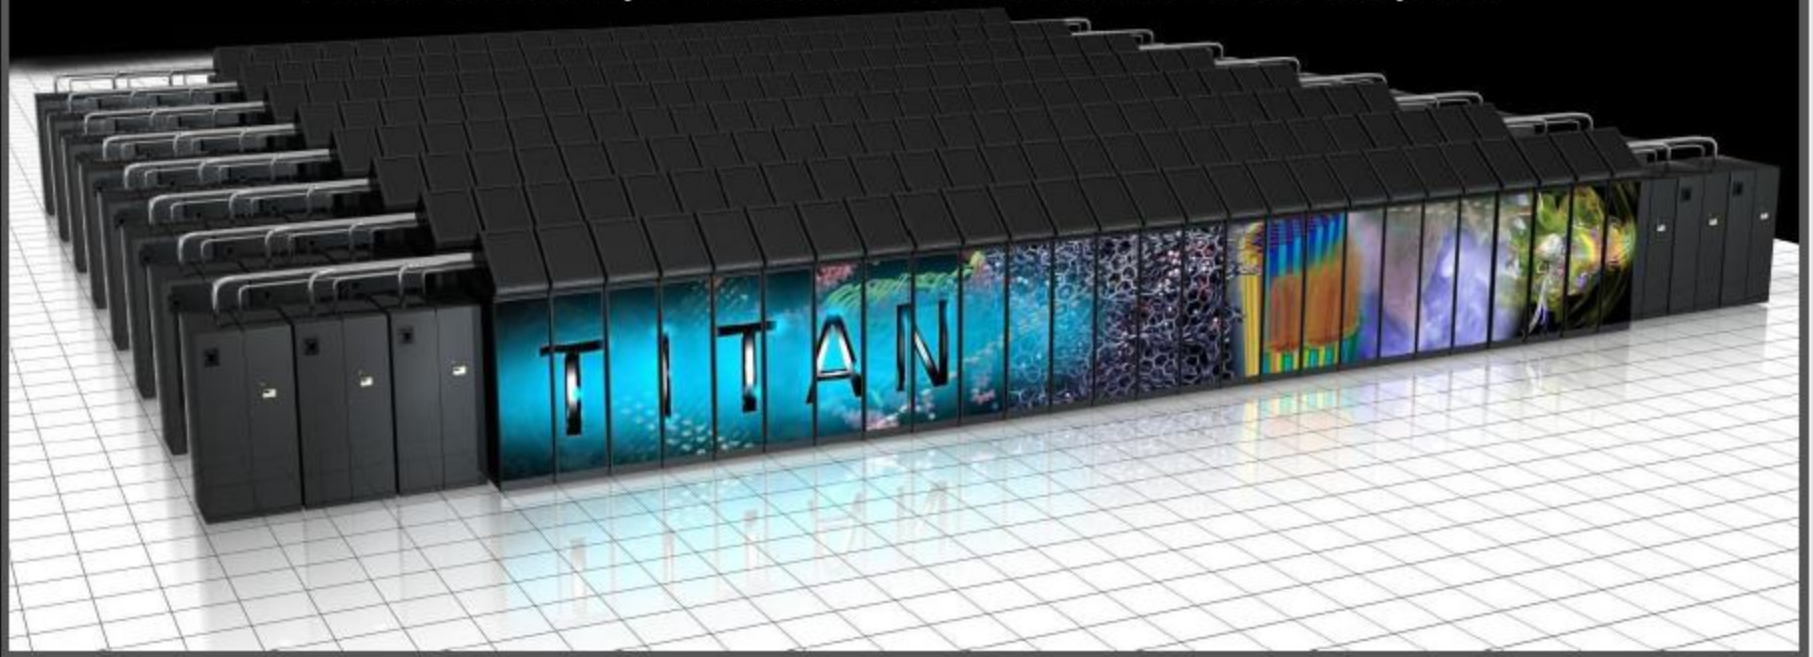

Tesla K20 Family: 3x Faster Than Fermi

Tesla K20X

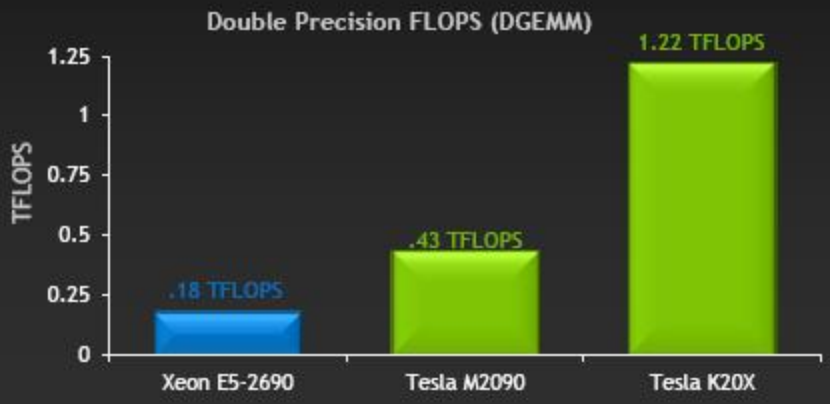

	Tesla K20X	Tesla K20
# CUDA Cores	2688	2496
Peak Double Precision	1.32 TF	1.17 TF
Peak DGEMM	1.22 TF	1.10 TF
Peak Single Precision	3.95 TF	3.52 TF
Peak SGEMM	2.90 TF	2.61 TF
Memory Bandwidth	250 GB/s	208 GB/s
Memory size	6 GB	5 GB
Total Board Power	235W	225W

Kepler GK110 Block Diagram

Architecture

- 7.1B Transistors
- 15 SMX units
- > 1 TFLOP FP64
- 1.5 MB L2 Cache
- 384-bit GDDR5

Kepler GK110 SMX vs Fermi SM

32

12*16= 192

inside-tesla-kepler-k20-family.pdf - Adobe Reader

File Edit View Window Help

12 / 46 125% Tools Sign Comment

Kepler GK110 SMX vs Ferm

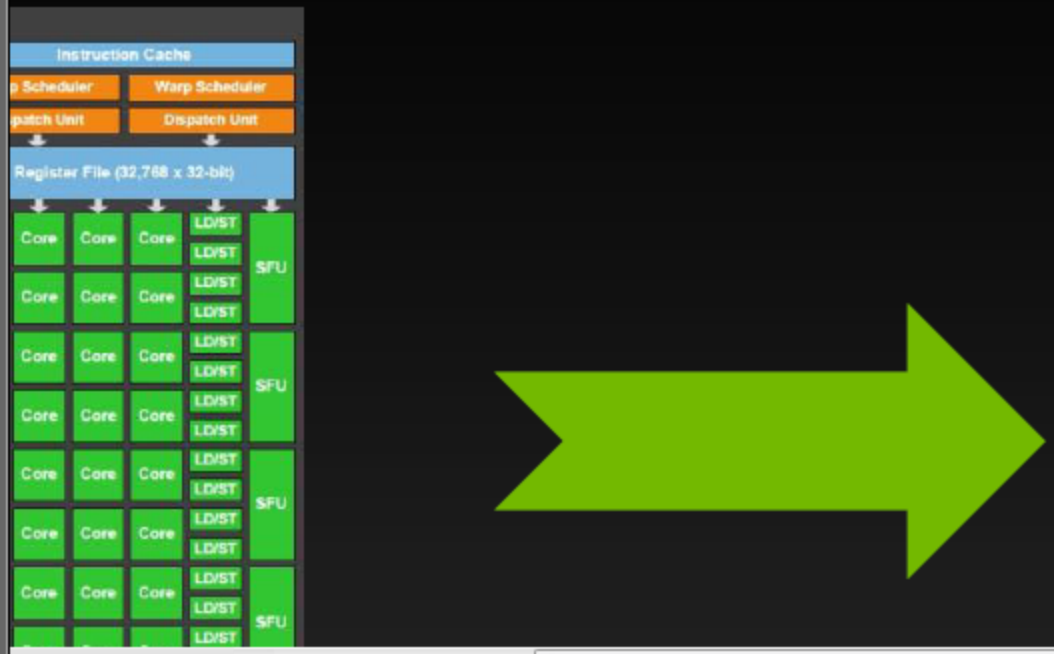

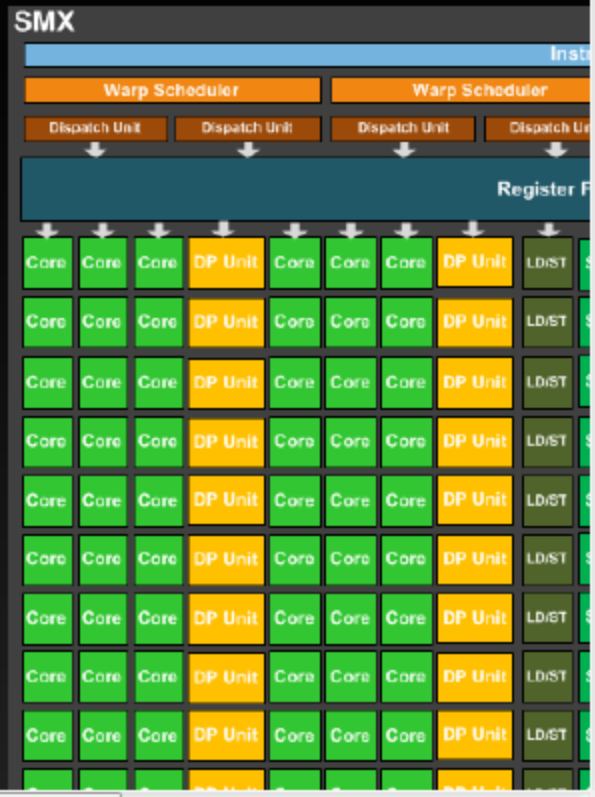

12.00 x 6.75 in

NVIDIA Updates GPU Roadmap; Announces Volta Family For Beyond 2014

17 Comments
+ Add A Comment

by Ryan Smith on March 19, 2013 7:15 PM EST

Posted in GPUs NVIDIA GTC 2013 Maxwell Volta

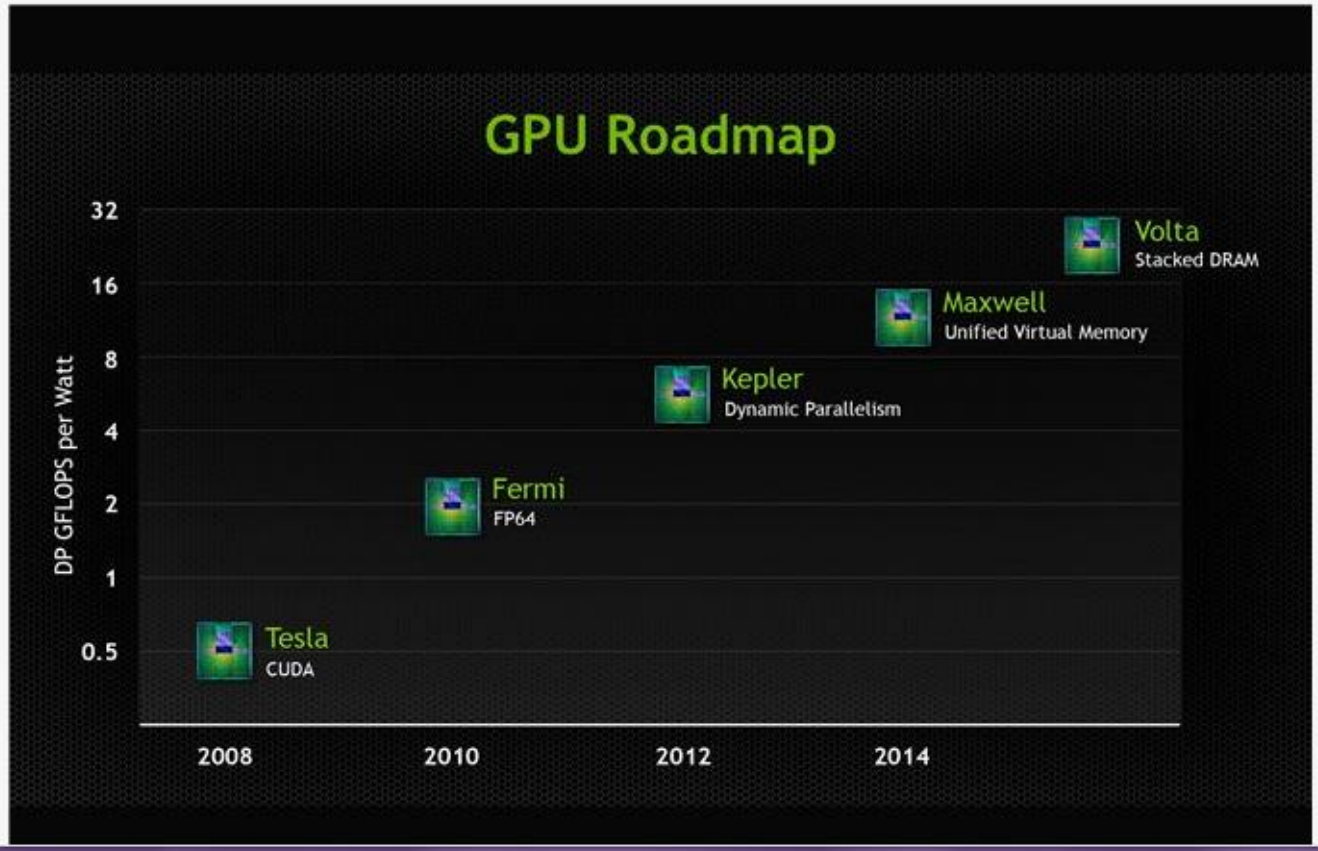

Clouds

A typical data center of a well-known, very-large-scale Cloud Service Provider

(retrieved via Google Images)

A typical data center of a well-known, very-large-scale Cloud Service Provider

(retrieved via Google Images)

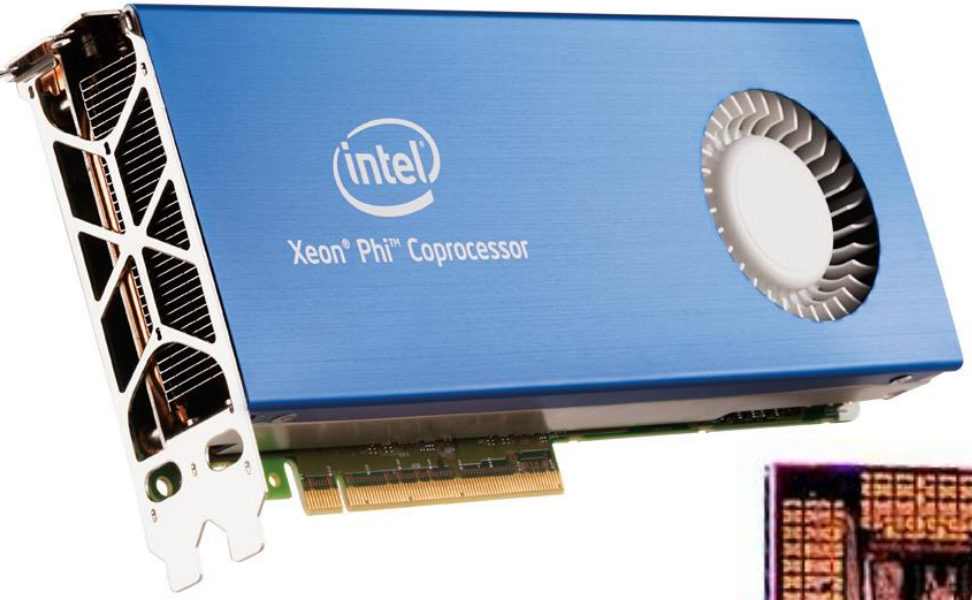
intel
Xeon® Phi™ Coprocessor

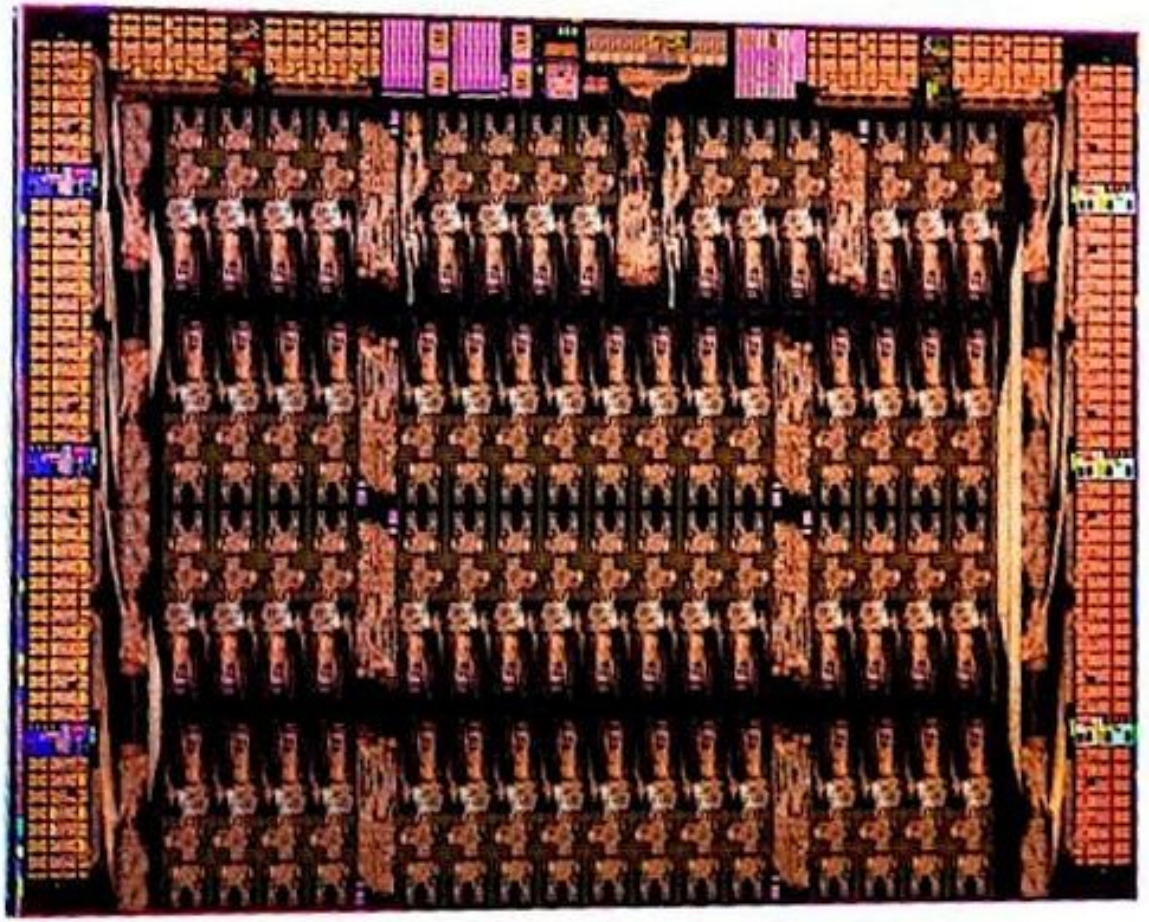

Tesla K20 Family: 3x Faster Than Fermi

Tesla K20X

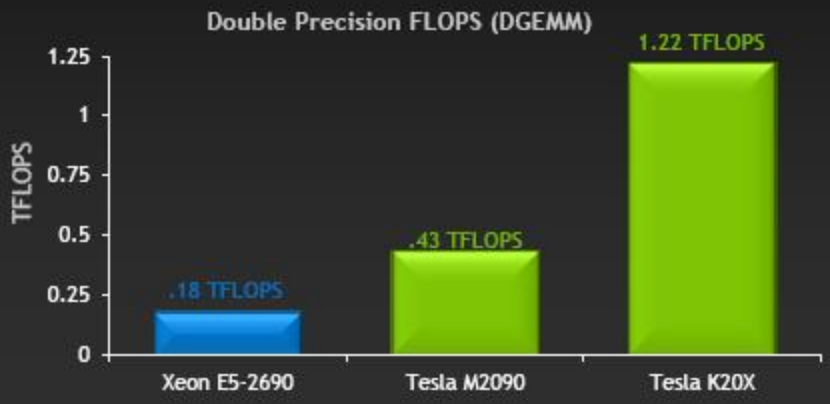

	Tesla K20X	Tesla K20
# CUDA Cores	2688	2496
Peak Double Precision	1.32 TF	1.17 TF
Peak DGEMM	1.22 TF	1.10 TF
Peak Single Precision	3.95 TF	3.52 TF
Peak SGEMM	2.90 TF	2.61 TF
Memory Bandwidth	250 GB/s	208 GB/s
Memory size	6 GB	5 GB
Total Board Power	235W	225W

Instruments

creativec.com

creativec.com

Other bookmarks

CONTACT

CREATIVEC

COMPUTE FASTER

HOME STELLETTO VR GROTTO BUY ESTELLA

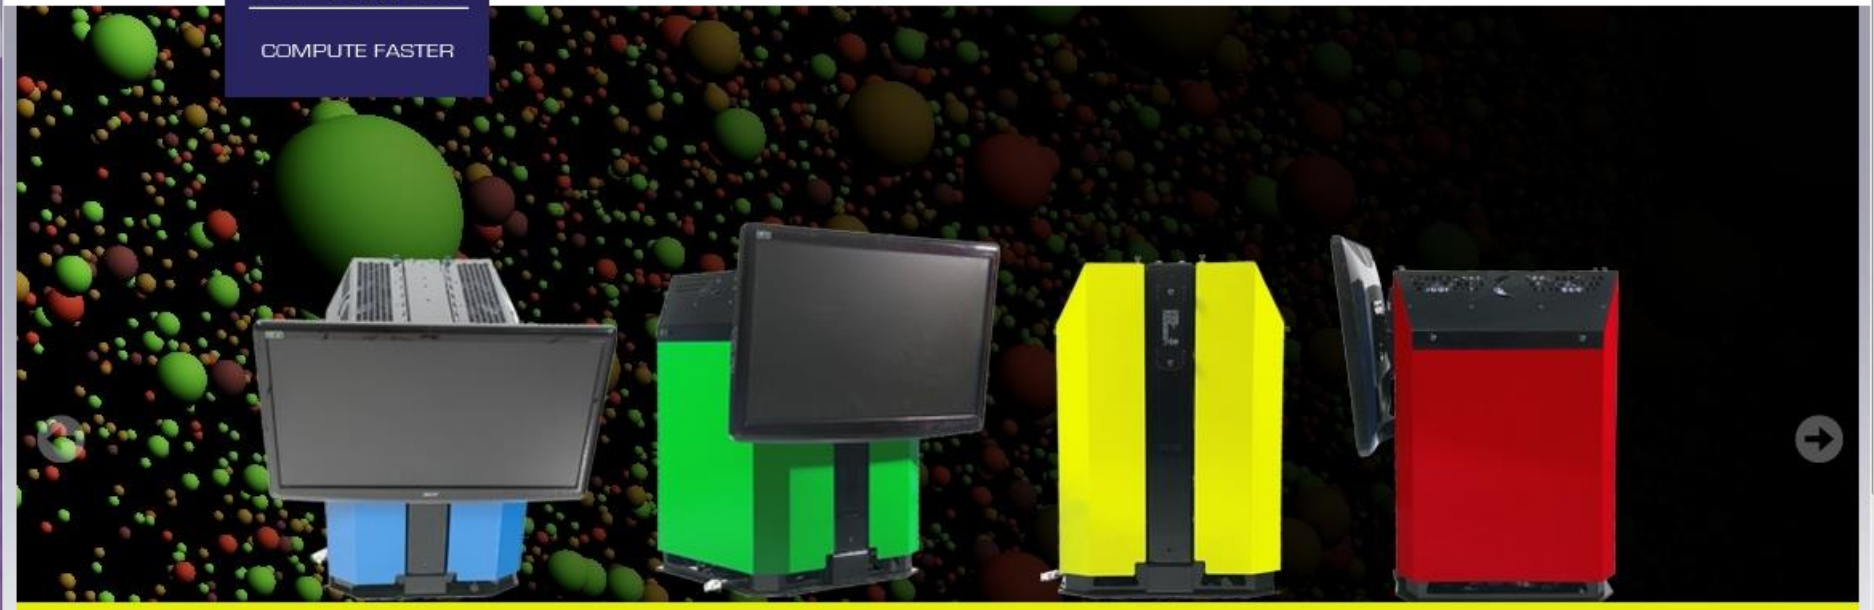

DEVELOP. COMPUTE. VISUALIZE. INTERACT.

Bring your simulations to life with CreativeC!

STELLETTO_AMD3.pdf - Adobe Reader

File Edit View Window Help

1 / 2 40% Tools Sign Comment

STELLETTO CODE DEVELOPMENT PLATFORM

Tesla K20 Family: 3x Faster Than Fermi

Tesla K20X

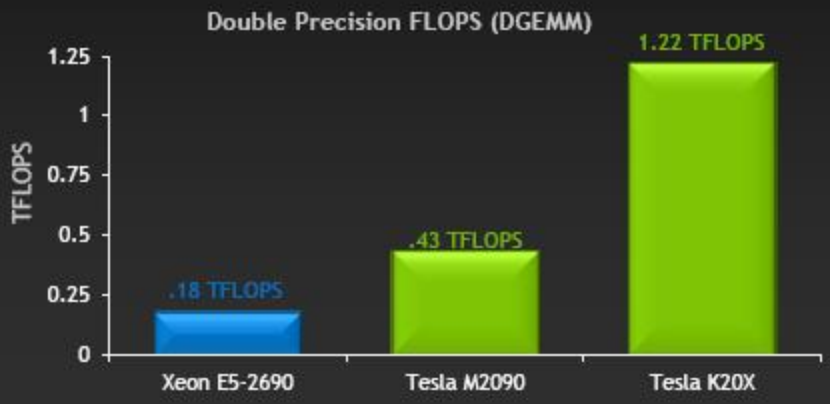

	Tesla K20X	Tesla K20
# CUDA Cores	2688	2496
Peak Double Precision	1.32 TF	1.17 TF
Peak DGEMM	1.22 TF	1.10 TF
Peak Single Precision	3.95 TF	3.52 TF
Peak SGEMM	2.90 TF	2.61 TF
Memory Bandwidth	250 GB/s	208 GB/s
Memory size	6 GB	5 GB
Total Board Power	235W	225W

3 GPU Comparison

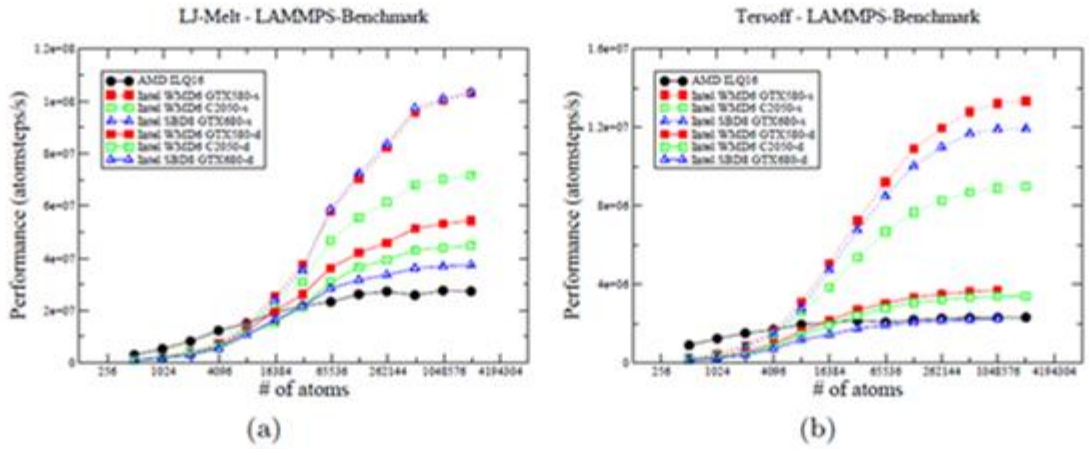

90s

These graphs show the ability of Stelletto to profile different hybrid architectures. This work is preliminary, internal and not meant to show results, but meant to show the ability to get results.

With the use of Accelerators, systems response gets much more complicated.

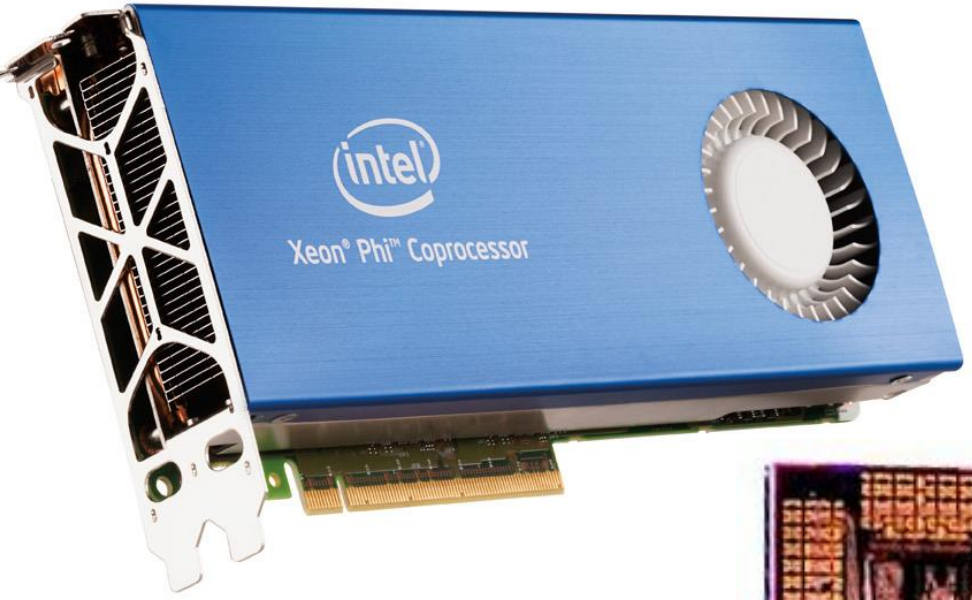
intel
Xeon® Phi™ Coprocessor

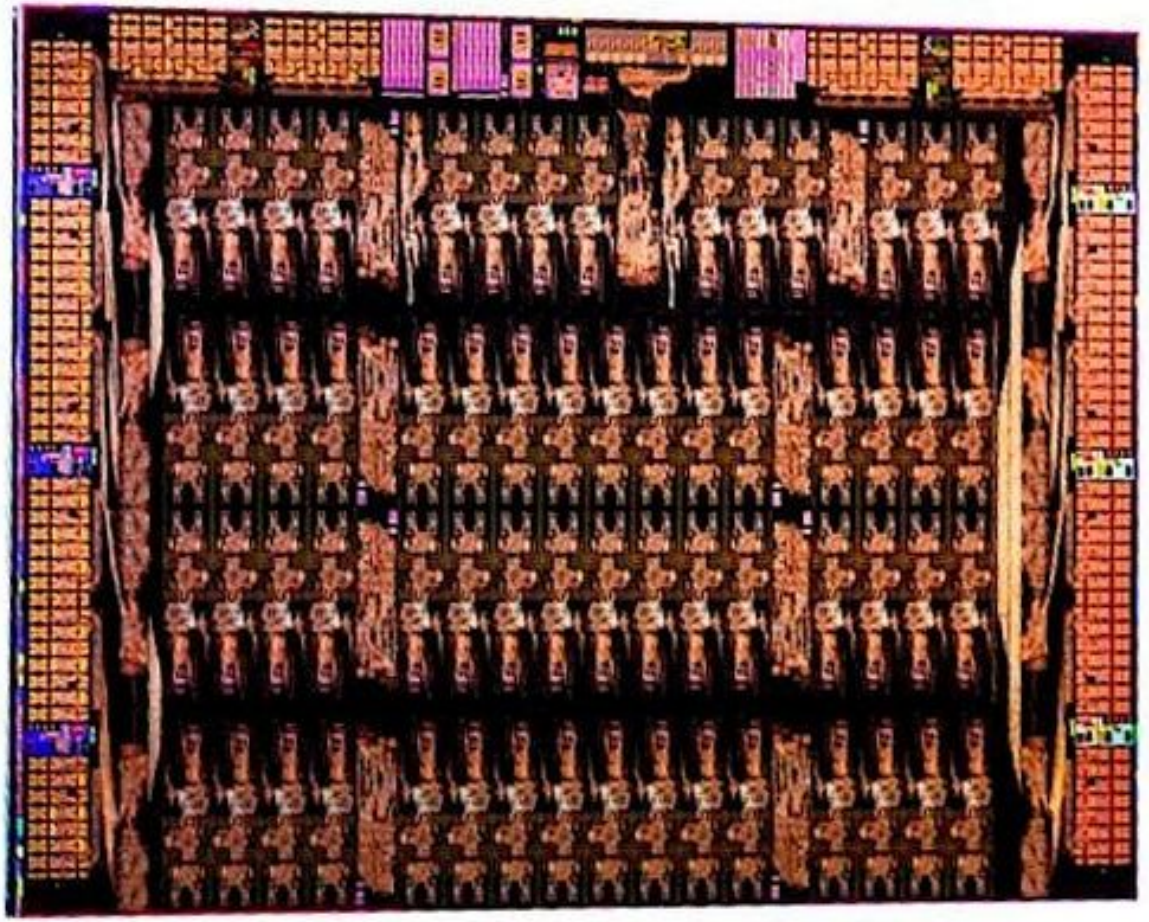

MedeA Instrument

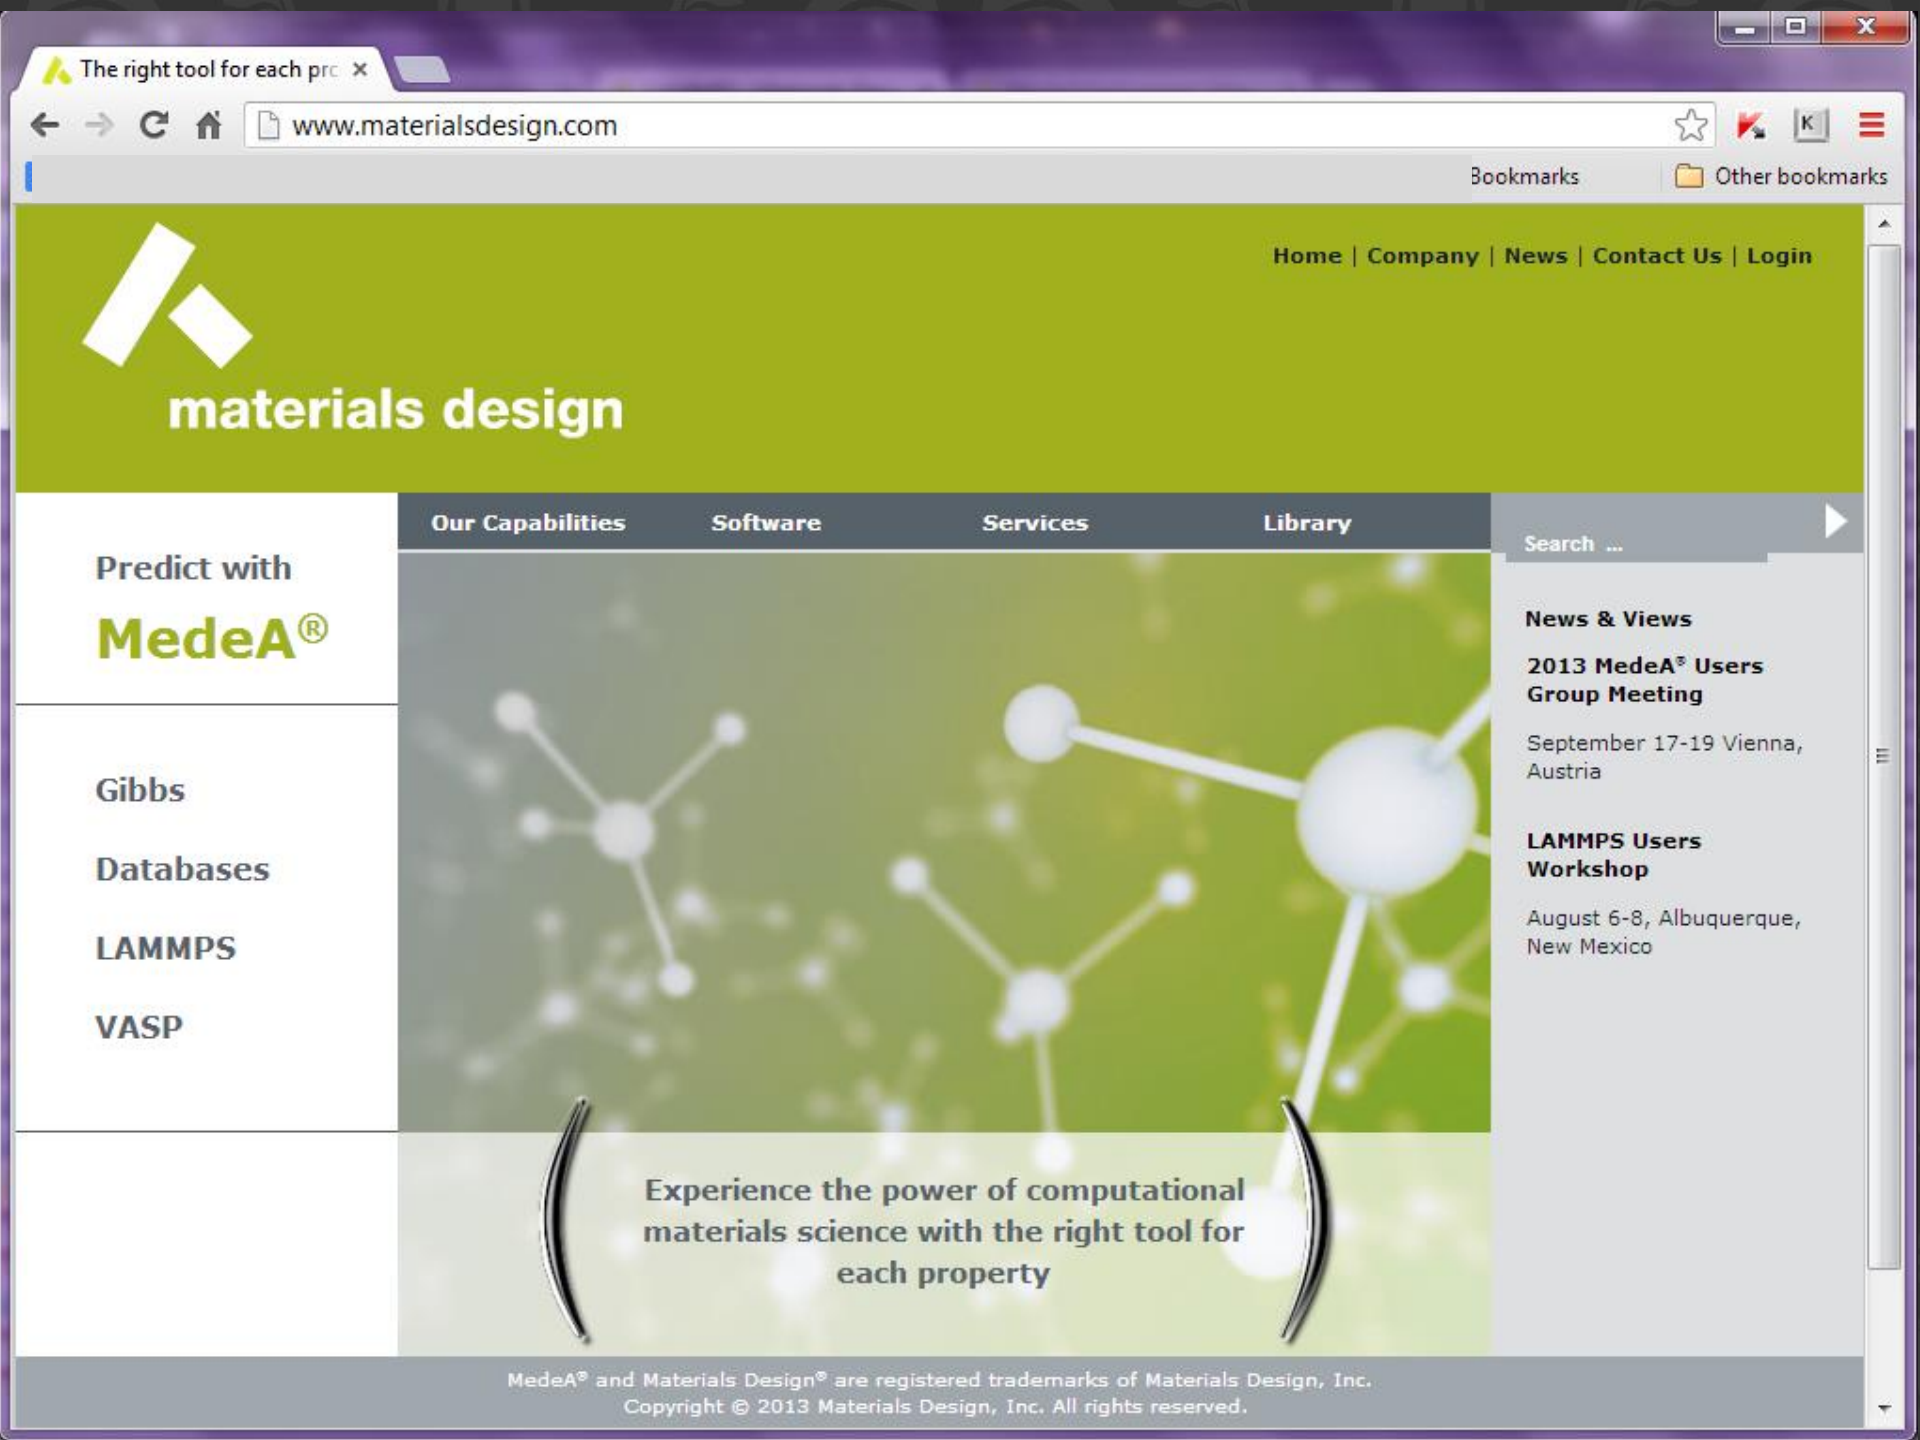
The right tool for each pro
www.materialsdesign.com
Bookmarks
Other bookmarks
Home | Company | News | Contact Us | Login
materials design
Our Capabilities
Software
Services
Library
Search ...
Predict with
MedeA®
Gibbs
Databases
LAMMPS
VASP
News & Views
2013 MedeA® Users Group Meeting
September 17-19 Vienna, Austria
LAMMPS Users Workshop
August 6-8, Albuquerque, New Mexico
Experience the power of computational materials science with the right tool for each property
MedeA® and Materials Design® are registered trademarks of Materials Design, Inc.
Copyright © 2013 Materials Design, Inc. All rights reserved.

Gibbs

LAMMPS

Embedded Atom Potentials

Thermal conductivity

Diffusion

Viscosity

JobServer TaskServer

Forcefield

MedeA® JEMS

Tip of the day

Our Customers can see more

Registered customers have access to our publications.

In addition to the abstract, you can read the full publication in pdf format.
Interested in diffusion of hydrogen in Ni?

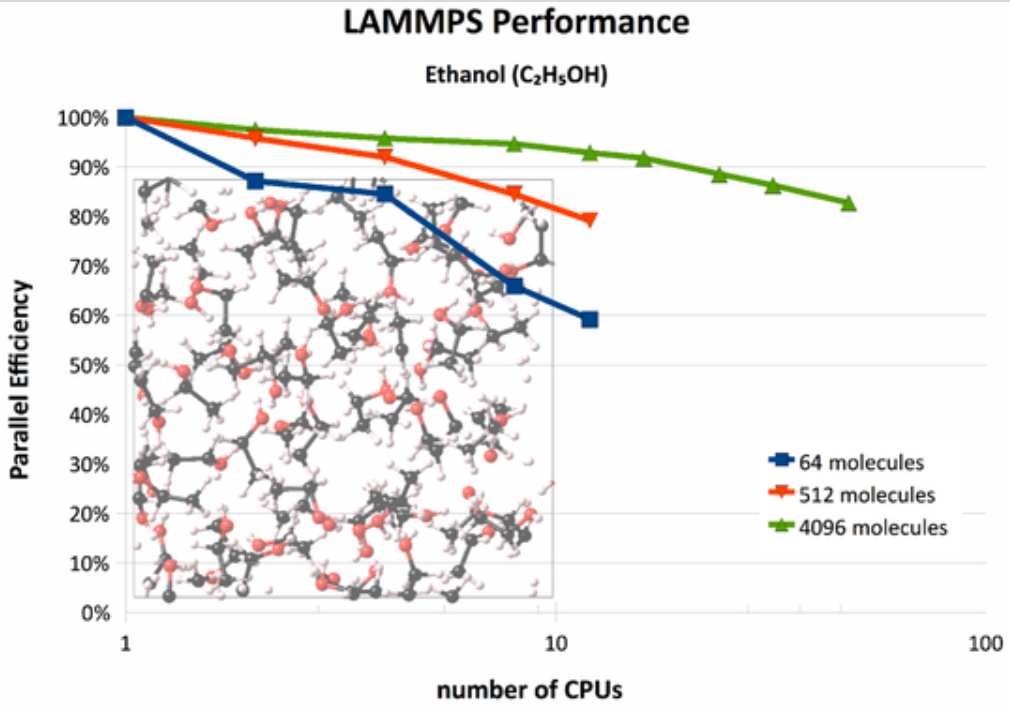

Key Benefits of MedeA®-LAMMPS

- Hides the details of getting a molecule or solid and its associated forcefield parameters into the format needed by LAMMPS
- Powerful flowchart interface lets you easily set up complex calculations
  - Flowcharts include stages that are computational–minimization, NVT dynamics, NPT dynamics, etc. – as well as those that manipulate or build structures, such as setting the density to that computed by a previous stage or building a larger box of fluid
  - Any number of stages can be chained together to perform a detailed, reproducible calculation. For example, you may wish initially equilibrate the system, then adjust the pressure, and finally run a long simulation to extract properties
  - Flow charts from any previous calculation can be edited and run again, on the same molecules or a different ones

Stress-Strain Behavior of a Cu Nanowire Simulated with MedeA-LAMMPS

Molecular Dynamics Simulation of Crystalline C_{60} Buckminsterfullerene with MedeA and LAMMPS

Atomistic Simulations of Microelectronic Materials: Prediction of Mechanical, Thermal, and Electrical Properties

Validation of molecular dynamics simulation

Atomistic study of edge and screw ⟨c + a⟩ dislocations in magnesium

materials design

Services
Training & Support
Consulting & Contract Research
MedeA® Instrument
Tip of the day
Publications

Find relevant publications (with full abstract and more) on the right side of many pages or **on our searchable publications list**

MedeA® Instrument

The MedeA® Instrument provides a powerful, integrated platform for atomistic simulations. State-of-the-art hardware and design bring affordable, high-performance computing out of the datacenter and into your office.

The MedeA® Environment comes preinstalled, completely wired, with a working queuing system, including all the modules you need, tested by our team of experts saving you time and hassle.

The MedeA® Instrument has been carefully designed to be quiet enough for your office. It plugs into a standard wall socket, and does not need special cooling – if you are comfortable, then it is too.

The MedeA® Instrument provides over 1 TFlop of peak computing power - more power than the largest supercomputer in the world about 15 years ago!

…and finally, to "avoid stupid in or stupid out"…

(…and to get meaningful work on a small system…)