





Large scale GPU based HPC in Computational Chemistry



Ananda Bhattacharjee

NVIDIA-India



*“ The more important fundamental laws and facts of physical science have all been discovered...
our future discoveries must be looked for in the sixth place of decimals. ”*

Albert Michelson, 1894 (prior to Quantum Physics)

Three Pillars of Research

$$\left[\frac{-\hbar^2}{2m} \nabla^2 + V \right] \Psi = i\hbar \frac{\partial}{\partial t} \Psi$$

Theory



Simulation




Experimentation



The Era of Accelerated Computing is Here



Era of
Accelerated Computing



Era of
Distributed Computing



Era of
Vector Computing




1980

1990

2000

2010


2020




Accelerated Computing Multi-core plus Many-cores



CPU
Optimized for
Serial Tasks




GPU Accelerator
Optimized for Many
Parallel Tasks



3-10X+ Comp Thruput
7X Memory Bandwidth
5x Energy Efficiency

CPU Vs GPU – Low latency or high throughput


CPU



CPU

- Optimized for low-latency access to cached data sets.
- Control logic for out-of-order and speculative execution

GPU




GPU

- Optimized for data-parallel, throughput computation
 - Architecture tolerant of memory latency
- More transistors dedicated to computation


Performance gap continues to grow



Peak Double Precision FLOPS




Peak Memory Bandwidth






How GPU Acceleration Works



Where Is Heterogeneous Computing Today?




TITAN: World's 2nd Fastest Supercomputer

18,688 Tesla K20X GPUs

27 Petaflops Peak, 17.59 Petaflops on Linpack

90% of Performance from GPUs



World's First Whole H1N1 Virus Simulation

More accurate & complete model

Furthers understanding of drug interactions



Mole-8.5
GPU Supercomputer
at CAS-IPE




中国科学院
CHINESE ACADEMY OF SCIENCES

Gordon Bell Prize Winner

Achieved using NVIDIA GPUs

Special Achievement in Scalability and Time-to Solution



“Peta-scale Phase-Field Simulation for Dendritic Solidification on the TSUBAME 2.0 Supercomputer”

-- Shimokawabe et. al.

Tsubame 2.0

Tokyo Institute of Technology



4,224 Tesla GPUs +
2,816 x86 CPUs


Flagship Scientific Applications on Titan



Material Science (WL-LSMS)
Role of material disorder, statistics, and fluctuations in nanoscale materials and systems.




Climate Change (CAM-SE)
Answer questions about specific climate change adaptation and mitigation scenarios; realistically represent features like precipitation patterns/statistics and tropical storms.




Biofuels (LAMMPS)
A multiple capability molecular dynamics code.


Astrophysics (NRDF)
Radiation transport – critical to astrophysics, laser fusion, combustion, atmospheric dynamics, and medical imaging.



Combustion (S3D)
Combustion simulations to enable the next generation of diesel/bio-fuels to burn more efficiently.




Nuclear Energy (Denovo)
Unprecedented high-fidelity radiation transport calculations that can be used in a variety of nuclear energy and technology applications.




Tesla Kepler Family

World's Fastest and Most Efficient HPC Accelerators



	GPUs	Single Precision Peak (SGEMM)	Double Precision Peak (DGEMM)	Memory Size	Memory Bandwidth (ECC off)	System Solution
Weather & Climate, Physics, BioChemistry, CAE, Material Science	K20X	3.95 TF (2.90 TF)	1.32 TF (1.22 TF)	6 GB	250 GB/s	Server only
	K20	3.52 TF (2.61 TF)	1.17 TF (1.10 TF)	5 GB	208 GB/s	Server + Workstation
Image, Signal, Video, Seismic	K10	4.58 TF	0.19 TF	8 GB	320 GB/s	Server only

GPU Acceleration Across All Platforms



Tesla Data Center & Workstation GPU Products & Platforms

M



Tesla M-series GPUs
K10, K20 and K20X

C



Tesla C-series GPUs
K20




OEM Integrated CPU-GPU
Servers & Blades



Workstations
2 to 4 Tesla GPUs

A multi-user GPU cluster



- “Multi-user mode”:
 - Multiple users share cluster
 - Access to GPU nodes is allocated using job scheduler
- Typically requires at least one “head node” which doesn’t take part in compute, runs cluster services

GPU Management & Monitoring

NVIDIA Systems Management Interface (nvidia-smi)

Products

Features

All GPUs

- List of GPUs
- Product ID
- GPU Utilization
- PCI Address to Device Enumeration

Server products

- Exclusive use mode
- ECC error count & location
- GPU temperature
- Unit fan speeds
- PSU voltage/current
- LED state
- Serial number
- Firmware version

Utilization

% GPU Utilization, % Memory utilization

```
[user@cuda-linux ~]$ nvidia-smi -q
Timestamp                               : wed JUN 9 10:01:01 2010
Unit 0:
  Product Name                           : NVIDIA Tesla SXYZ
  Product ID                              : 123-45678-012
  Serial Number                           : 0123456789012
  Firmware Ver                            : X.Y
GPU 0:
  Product Name                            : Tesla C2050
  PCI ID                                  : 6d110de
  Temperature                             : 63 C
  ECC errors                              :
  Single bit                              : 0
  Double bit                              : 0
  Total                                   : 0
  Aggregate single bit                    : 0
  Aggregate double bit                    : 10
  Aggregate total                         : 10
Fan Tachs:
  #00: 263 Status: NORMAL
  #01: 263 Status: NORMAL
  #02: 263 Status: NORMAL
...
PSU:
  Voltage                                 : 12.37 V
  Current                                 : 12.07 A
LED:
  State                                  : AMBER
```



LSF, HPC, Cluster Manager



Bright Cluster Manager



ROCKS+MOAB



PBS Professional



NVML Plugin for GPUs



Univa Grid Engine

Job Scheduling & Cluster Management

Torque with NVIDIA GPUs

- Reports information on GPUs per-node, uses that data for scheduling decisions
- Job requests can set GPU compute mode

```
[adeconinck@psglogin ~]$ pbsnodes wm030
```

```
wm030
```

```
state = free  
np = 12
```

```
properties = batch,qdr1b,wm,m2090,m2090x1,1b  
ntype = cluster
```

```
status = reftime=1359671852,varattr=,jobs=,state=free,netload=2135037617,gres=,loadave=0.41,ncpus=1  
2,physmem=49552084kb,availmem=55934124kb,totmem=57560444kb,idletime=79806,nusers=0,nsessions=0,uname=Lin  
ux wm030 2.6.32-220.23.1.el6.x86_64 #1 SMP Tue Jun 12 11:20:15 EDT 2012 x86_64,opsys=linux
```

```
mom_service_port = 15002  
mom_manager_port = 15003
```

```
gpus = 1
```

```
gpu_status = gpu[0]=gpu_id=0000:03:00.0;gpu_pci_device_id=277942494;gpu_pci_location_id=0000:03:00.  
0;gpu_product_name=Tesla M2090;gpu_display=Enabled;gpu_memory_total=6143 MB;gpu_memory_used=10 MB;gpu_mo  
de=Default;gpu_state=Unallocated;gpu_utilization=0%;gpu_memory_utilization=0%;gpu_ecc_mode=Disabled,driv  
er ver=304.64,timestamp=Thu Jan 31 14:37:32 2013
```

Node name
and state

Node status
info

GPU count
and status
info

3 Ways to Program GPUs

Applications

Libraries

“Drop-in”
Acceleration

OpenACC
Directives

Easily Accelerate
Applications

Programming
Languages


Maximum
Flexibility

GPU Accelerated Libraries

“Drop-in” Acceleration for your Applications



Linear Algebra

FFT, BLAS,
SPARSE, Matrix




Numerical & Math

RAND, Statistics




Data Struct. & AI

Sort, Scan, Zero Sum



Visual Processing

Image & Video



OpenACC

Open Programming Standard for Parallel Computing

“PGI OpenACC will enable programmers to easily develop portable applications that maximize the performance and power efficiency benefits of the hybrid CPU/GPU architecture of Titan.”

--Buddy Bland, Titan Project Director, Oak Ridge National Lab

“OpenACC is a technically impressive initiative brought together by members of the OpenMP Working Group on Accelerators, as well as many others. We look forward to releasing a version of this proposal in the next release of OpenMP.”


--Michael Wong, CEO OpenMP Directives Board



OpenACC Members



OpenACC Directives



Portable compiler hints

Compiler parallelizes code

Designed for multicore CPUs & many core GPUs / Accelerators

Programming Languages

Numerical analytics ▶

MATLAB, Mathematica

Fortran ▶

CUDA Fortran

C ▶

CUDA C

C++ ▶

CUDA C++


Python ▶

PyCUDA

C# ▶

GPU.NET

Papers mentioning parallel programming languages.
Data according to Google Scholar (Feb. 2014)



(c) Simon McIntosh-Smith 2014

MATLAB and Parallel Computing Toolbox


Industry standard language for algorithm development & data analysis



GPU Value

- Allows practical analysis of large data
- Scales from GPU workstations (*Parallel Computing Toolbox*) up to GPU clusters (*MATLAB Distributed Computing Server*)
- Significant acceleration for spectral analysis, linear algebra, and stochastic simulations, etc.

Highlights

- GPU accelerated native MATLAB operations
- GPU support for several MATLAB Toolboxes, including *Signal Processing*, *Phased Array System*, *Communications System*, and *Neural Network*
- Integration with user CUDA kernels in MATLAB
- MATLAB Compiler support (GPU acceleration in stand-alone applications)





TESLA

Quantum Chemistry Module




Overview of Life & Material Accelerated Apps



- MD: All key codes are available
 - CHARMM, DESMOND, DL_POLY, GROMACS, LAMMPS, OpenMM, NAMD
 - GPU only codes: AMBER, ACEMD, HOOMD-Blue
 - Great multi-GPU performance
 - Focus: scaling to large numbers of GPUs / nodes



- QC: All key codes are ported/optimizing:
 - Active GPU acceleration projects:
 - Abinit, BigDFT, CP2K, GAMESS, Gaussian, GPAW, NWChem, Quantum Espresso, VASP & more



- GPU only code: TeraChem
- Bioinformatics – actively recruiting
- Analytical instruments – targeted recruiting

Individual_1_haplo1	AACGATTATCGCAATAACGAGGATTATCCCAGTTA
Individual_1_haplo2	AACGATTATCGCAATGACGAGGATTATCTCAGTTA
Individual_2_haplo1	AACGACTATCGCAATAACGAGGATTATCCCAGTTA
Individual_2_haplo2	AACGATTATCGCAATAACGAGGATTATCCCAGTTA
Individual_3_haplo1	AACGACTATCGCAATAACGAGGATTATCCCAGTTA
Individual_3_haplo2	AACGATTATCGCAATGACGAGGATTATCTCAGTTA
Individual_4_haplo1	AACGATTATCGCAATAACGAGGATTATCCCAGTTA
Individual_4_haplo2	AACGATTATCGCAATGACGAGGATTATCTCAGTTA

↑ ↑ ↑ ↑

Molecular Dynamics (MD) Applications



Application	Features Supported	GPU Perf	Release Status	Notes/Benchmarks
AMBER	PMEMD Explicit Solvent & GB Implicit Solvent, REMD, aMD	> 360 ns/day JAC NVE on 2X K40s (peer-2-peer)	Released, AMBER 14 Multi-GPU, multi-node	http://ambermd.org/gpus/benchmarks.htm#Benchmarks
CHARMM	Implicit (5x), Explicit (2x) Solvent via OpenMM	2x C2070 equals 32-35x X5667 CPUs	Released, C37b1 Single & Multi-GPU in single node	http://www.charmm.org/news/c37b1.html#postjump Native GPU port in development @ NREL
DESMOND	Bonded, pair, excluded interactions; Van der Waals, electrostatic, non-bonded far interactions	D.E. Shaw White Paper coming	Released Version 3.4.0/0.7.2	Academic: https://www.deshawresearch.com/resources_desmond.html For-Profit: http://www.schrodinger.com/productpage/14/3/
GROMACS	Explicit (native port), Implicit (via OpenMM)	Explicit-water RNase, vsites, Core i7-3930 + 1x K20: 300ns/day, Explicit-water villin, vsites, Core i7-3930 + 1x K20: 850ns/day	Released Multi-GPU, multi-node	Release 4.6.5; 1 st Multi-GPU support www.gromacs.org Version 5.0 beta2 released
LAMMPS	Lennard-Jones, Gay-Berne, Tersoff & <u>many</u> more potentials	3.5-18x on ORNL Titan	Released Multi-GPU, multi-node	http://lammps.sandia.gov/bench.html#desktop and http://lammps.sandia.gov/bench.html#titan
NAMD	Full electrostatics with PME and most simulation features	4.0 ns/day F1-ATPase on 1x K20X	Released 100M atom capable Multi-GPU, multi-node	NAMD 2.9 http://www.ks.uiuc.edu/Research/namd/ NAMD 2.10 beta1 in nightly build

GPU Perf compared against Multi-core x86 CPU socket.
GPU Perf benchmarked on GPU supported features
and may be a kernel to kernel perf comparison

New/Additional MD Applications Ramping



Application	Features Supported	GPU Perf	Release Status	Notes
ACEMD	<u>Written for use only on GPUs</u>	150 ns/day DHFR on 1x K20	Released Single and Multi-GPUs	Production biomolecular dynamics (MD) software specially optimized to run on GPUs http://www.acellera.com/
DL_POLY	Two-body Forces, Link-cell Pairs, Ewald SPME forces, Shake VV	4x	Release V 4.04 Multi-GPU, multi-node	Source only, Results Published http://www.stfc.ac.uk/CSE/randd/ccg/software/DL_POLY/25526.aspx
Folding@Home	Powerful distributed computing molecular dynamics system; implicit solvent and folding	Depends upon number of GPUs	Released GPUs and CPUs	http://folding.stanford.edu GPUs get 4X the points of CPUs
GPUGrid.net	High-performance all-atom biomolecular simulations; explicit solvent and binding	Depends upon number of GPUs	Released	http://www.gpugrid.net/
HALMD	Simple fluids and binary mixtures (pair potentials, high-precision NVE and NVT, dynamic correlations)	Up to 66x on 2090 vs. 1 CPU core	Released, Version 0.2.1 Single GPU	http://halmd.org/benchmarks.html#supercooled-binary-mixture-kob-andersen
HOOMD-Blue	<u>Written for use only on GPUs</u>	Kepler 2X faster than Fermi	Released Multi-GPU, multi-node	http://codeblue.umich.edu/hoomd-blue/ and http://vimeo.com/85452639
mdcore	TBD	TBD	Released, Version 0.1.7	http://mdcore.sourceforge.net/download.html
OpenMM	Implicit and explicit solvent, custom forces	Implicit: 127-213 ns/day Explicit: 18-55 ns/day DHFR	Released, Version 6.0.1 Multi-GPU	library and application for MD; supports custom forces, https://simtk.org/home/openmm

GPU Perf compared against Multi-core x86 CPU socket.
GPU Perf benchmarked on GPU supported features and may be a kernel to kernel perf comparison

Quantum Chemistry Applications



Application	Features Supported	GPU Perf	Release Status	Notes
ABINIT	Local Hamiltonian, non-local Hamiltonian, LOBPCG algorithm, diagonalization / orthogonalization	1.3-2.7X	Released; Version 7.4.1 Multi-GPU support	www.abinit.org
ACES III	Integrating scheduling GPU into SIAL programming language and SIP runtime environment	10X on kernels	Under development Multi-GPU support	http://www.olcf.ornl.gov/wp-content/training/electronic-structure-2012/deumens_ESaccel_2012.pdf
ADF	Fock Matrix, Hessians	TBD	Available Q1 2014 Multi-GPU support	www.scm.com
BigDFT	DFT; Daubechies wavelets, part of Abinit	2-5X	Released, Version 1.7.0 Multi-GPU support	http://bigdft.org
Casino	Code for performing quantum Monte Carlo (QMC) electronic structure calculations for finite and periodic systems	TBD	Under development Multi-GPU support	http://www.tcm.phy.cam.ac.uk/~mdt26/casino.html
CASTEP	TBD	TBD	Under development	http://www.castep.org/Main/HomePage
CP2K	DBCSP (sparse matrix multiply library)	2-7X	Released Multi-GPU support	http://www.olcf.ornl.gov/wp-content/training/ascc_2012/friday/ACSS_2012_VandeVon_dele_s.pdf
GAMESS-US	Libqc with Rys Quadrature Algorithm, Hartree-Fock, MP2 and CCSD	1.3-1.6X, 2.3-2.9x HF	Released Multi-GPU support	http://www.msg.ameslab.gov/gamess/index.html

GPU Perf compared against Multi-core x86 CPU socket.
GPU Perf benchmarked on GPU supported features
and may be a kernel to kernel perf comparison

Quantum Chemistry Applications



Application	Features Supported	GPU Perf	Release Status	Notes
GAMESS-UK	(ss ss) type integrals within calculations using Hartree Fock <i>ab initio</i> methods and density functional theory. Supports organics & inorganics.	8x	Released, Version 7.0 Multi-GPU support	http://www.ncbi.nlm.nih.gov/pubmed/21541963
Gaussian	Joint PGI, NVIDIA & Gaussian Collaboration	TBD	Under development Multi-GPU support	http://www.gaussian.com/g_press/nvidia_press.htm
GPAW	Electrostatic poisson equation, orthonormalizing of vectors, residual minimization method (rmm-diiis)	8x	Released Multi-GPU support	https://wiki.fysik.dtu.dk/gpaw/devel/projects/gpu.html , Samuli Hakala (CSC Finland) & Chris O'Grady (SLAC)
Jaguar	Investigating GPU acceleration	TBD	Under development Multi-GPU support	Schrodinger, Inc. http://www.schrodinger.com/kb/278 http://www.schrodinger.com/productpage/14/7/32/
LATTE	CU_BLAS, SP2 algorithm	TBD	Released Multi-GPU support	http://on-demand.gputechconf.com/gtc/2013/presentations/S3195-Fast-Quantum-Molecular-Dynamics-in-LATTE.pdf
MOLCAS	CU_BLAS support	1.1x	Released, Version 7.8 Single GPU; Additional GPU support coming in Version 8	www.molcas.org
MOLPRO	Density-fitted MP2 (DF-MP2), density fitted local correlation methods (DF-RHF, DF-KS), DFT	1.7-2.3X projected	Under development Multiple GPU	www.molpro.net Hans-Joachim Werner

GPU Perf compared against Multi-core x86 CPU socket.
GPU Perf benchmarked on GPU supported features
and may be a kernel to kernel perf comparison

Quantum Chemistry Applications



Application	Features Supported	GPU Perf	Release Status	Notes
MOPAC2012	Pseudodiagonalization, full diagonalization, and density matrix assembling	3.8-14X	Released MOPAC2013 available Q1 2014 Single GPU	Academic port. http://openmopac.net
NWChem	Triples part of Reg-CCSD(T), CCSD & EOMCCSD task schedulers	2.75-3.75X	Released, Version 6.3 Multiple GPUs	Development GPGPU benchmarks: www.nwchem-sw.org And http://www.olcf.ornl.gov/wp-content/training/electronic-structure-2012/Krishnamoorthy-ESCMA12.pdf
Octopus	Full GPU support for ground-state, real-time calculations; Kohn-Sham Hamiltonian, orthogonalization, subspace diagonalization, poisson solver, time propagation	1.5-8X	Released, Version 4.1.0	http://www.tddft.org/programs/octopus/
ONETEP	TBD	TBD	Under development	http://www2.tcm.phy.cam.ac.uk/onetep/
PEtot	Density functional theory (DFT) plane wave pseudopotential calculations	6-10X	Released Multi-GPU	First principles materials code that computes the behavior of the electron structures of materials
Q-CHEM	RI-MP2	8x-14x	Released, Version 4.0.1 Multi-GPU support	http://www.q-chem.com/doc_for_web/qchem_manual_4.0.pdf

GPU Perf compared against Multi-core x86 CPU socket.
GPU Perf benchmarked on GPU supported features
and may be a kernel to kernel perf comparison

Quantum Chemistry Applications



Application	Features Supported	GPU Perf	Release Status	Notes
QMCPACK	Main features	3-4x	Released Multiple GPUs	NCSA University of Illinois at Urbana-Champaign http://cms.mcc.uiuc.edu/qmcpack/index.php/GPU_version_of_QMCPACK
Quantum Espresso/PWscf	PWscf package: linear algebra (matrix multiply), explicit computational kernels, 3D FFTs	2.5-3.5x	Released, Version 5.0 Multiple GPUs	Created by Irish Centre for High-End Computing http://www.quantum-espresso.org/index.php and http://www.quantum-espresso.org/
TeraChem	" <u>Full GPU-based solution</u> "	44-650X vs. GAMESS CPU version	Released, Version 1.5k Multi-GPU/single node	Completely redesigned to exploit GPU parallelism. YouTube: http://youtu.be/EJODzk6RFxE?hd=1 and http://www.olcf.ornl.gov/wp-content/training/electronic-structure-2012/Luehr-ESCMA.pdf
VASP	Hybrid Hartree-Fock DFT functionals including exact exchange	2x 2 GPUs comparable to 128 CPU cores	Available on request Multiple GPUs	By Carnegie Mellon University http://arxiv.org/pdf/1111.0716.pdf
WL-LSMS	Generalized Wang-Landau method	3x with 32 GPUs vs. 32 (16-core) CPUs	Under development Multi-GPU support	NICS Electronic Structure Determination Workshop 2012: http://www.olcf.ornl.gov/wp-content/training/electronic-structure-2012/Eisenbach_OakRidge_February.pdf

GPU Perf compared against Multi-core x86 CPU socket.
GPU Perf benchmarked on GPU supported features
and may be a kernel to kernel perf comparison

Thank you .

anandab@nvidia.com

9731133558