Quantum Chemistry (QC) on GPUs

March 2018

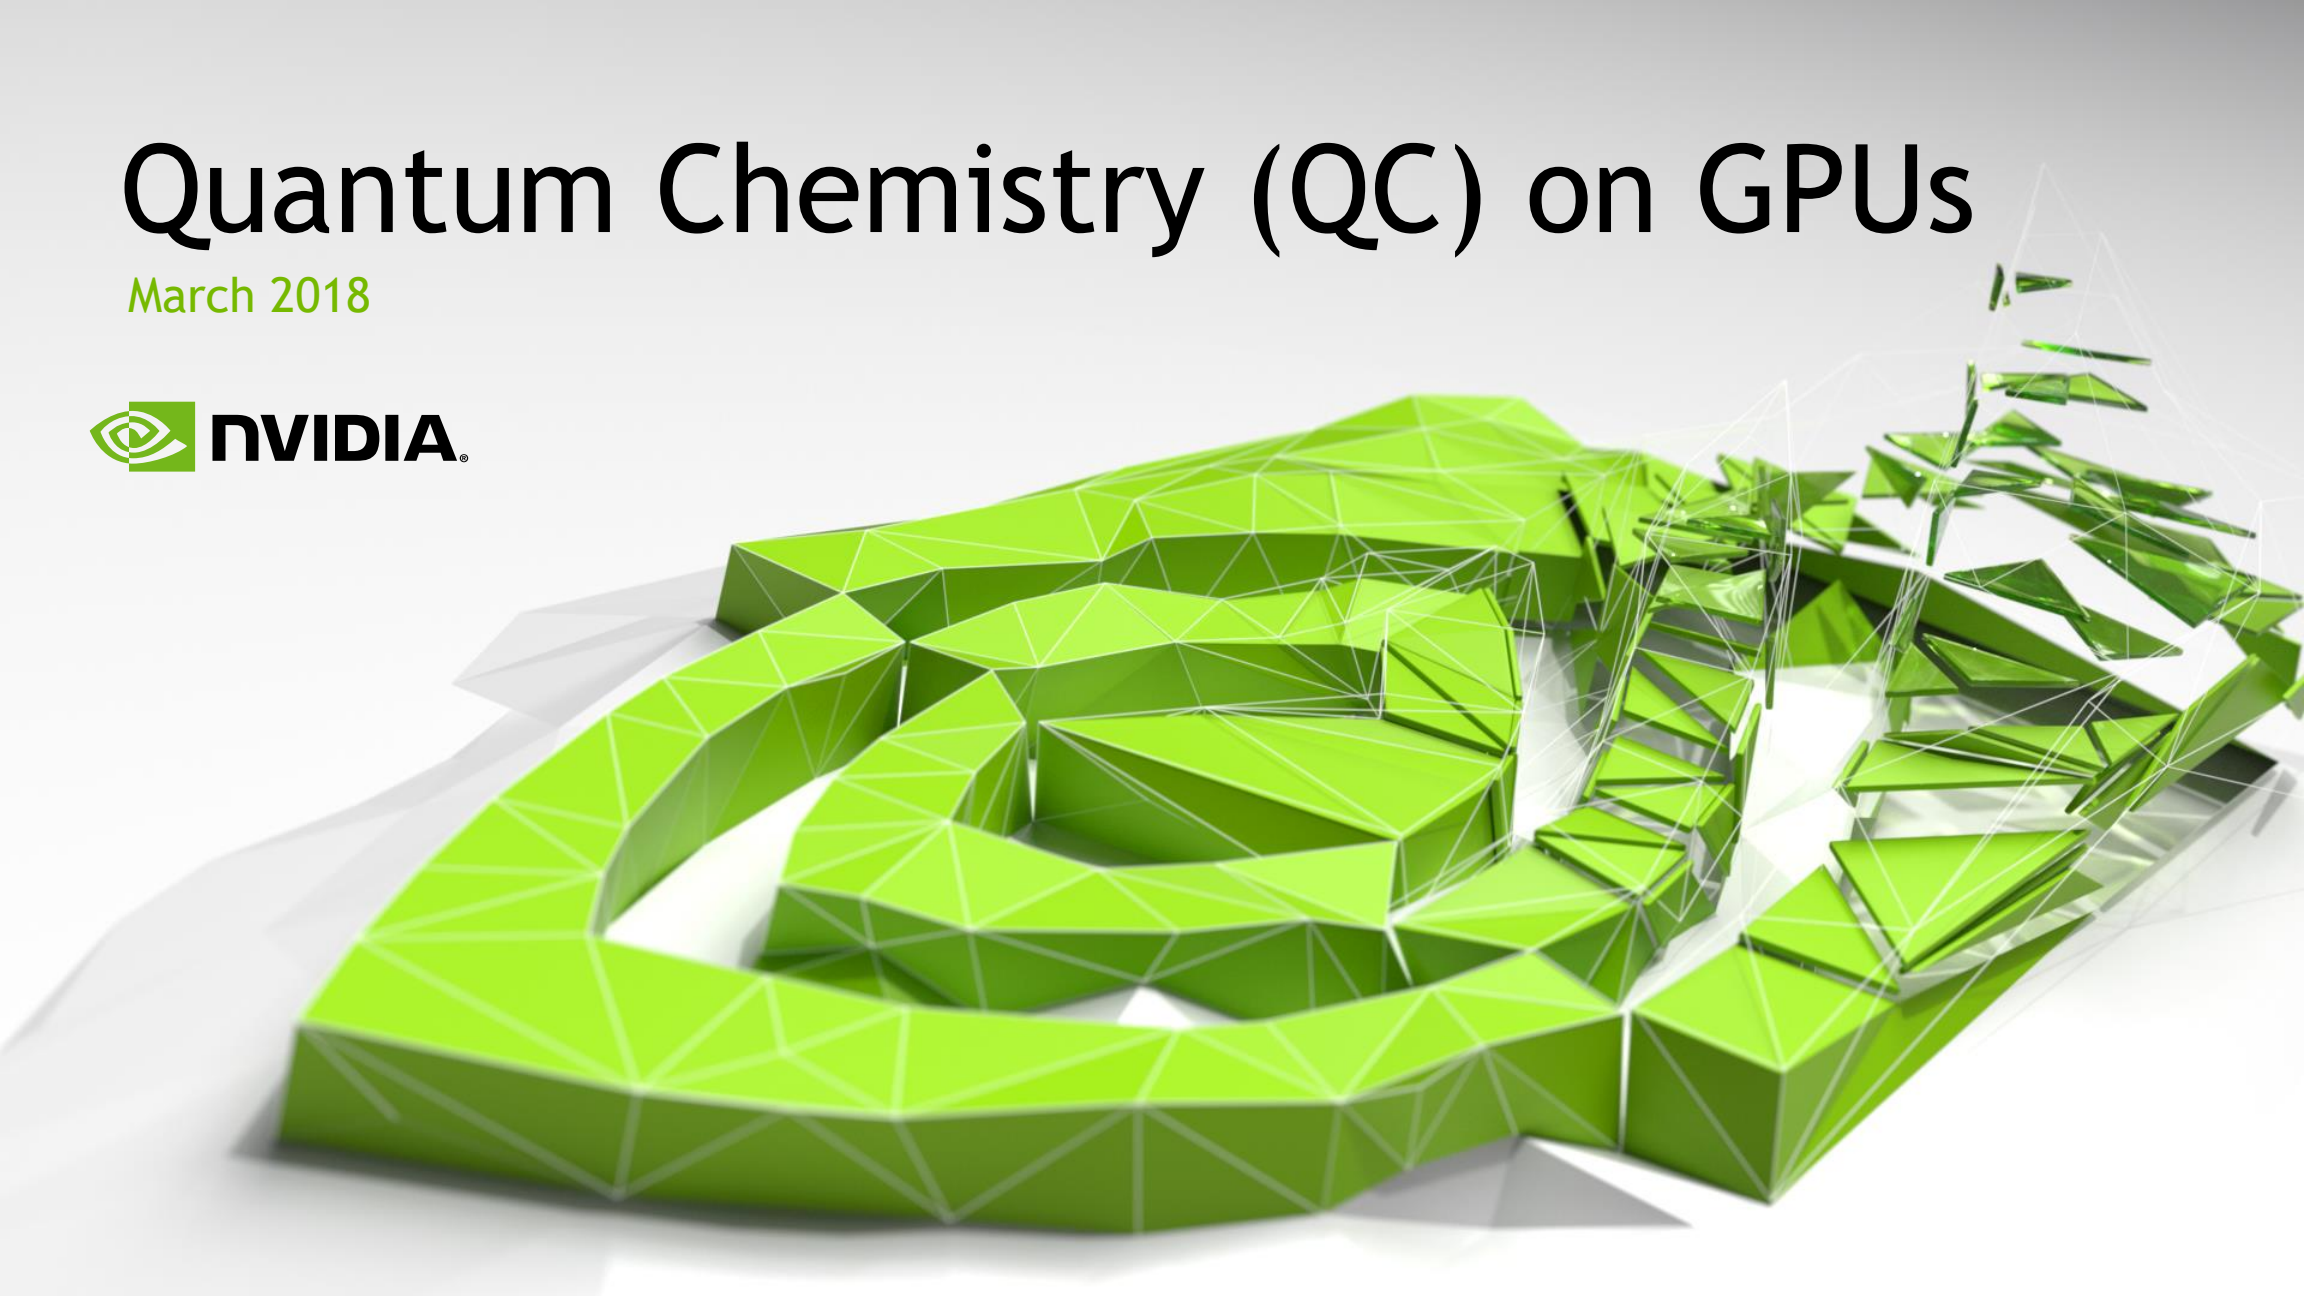

Overview of Life & Material Accelerated Apps

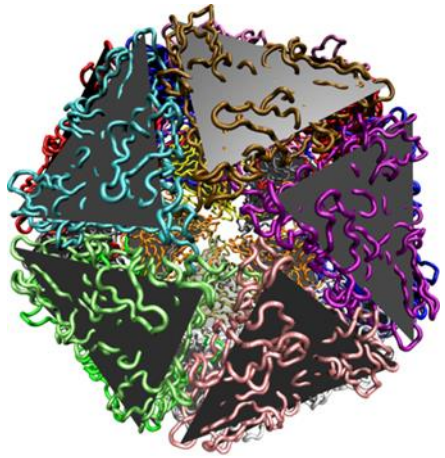

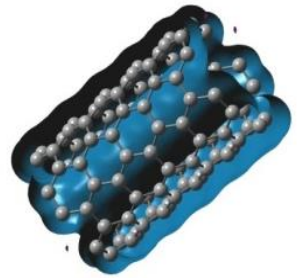

MD: All key codes are GPU-accelerated

- Great multi-GPU performance
- Focus on dense (up to 16) GPU nodes &/or large # of GPU nodes
- ACEMD*, AMBER (PMEMD)*, BAND, CHARMM, DESMOND, ESPResso, Folding@Home, GPUgrid.net, GROMACS, HALMD, HTMD, HOOMD-Blue*, LAMMPS, Lattice Microbes*, mdcore, MELD, miniMD, NAMD, OpenMM, PolyFTS, SOP-GPU* & more

QC: All key codes are ported or optimizing

- Focus on using GPU-accelerated math libraries, OpenACC directives
- GPU-accelerated and available today:
  - ABINIT, ACES III, ADF, BigDFT, CP2K, GAMESS, GAMESS-UK, GPAW, LATTE, LSDalton, LSMS, MOLCAS, MOPAC2012, NWChem, OCTOPUS*, PEtot, QUICK, Q-Chem, QMCPack, Quantum Espresso/PWscf, QUICK, TeraChem*
- Active GPU acceleration projects:
  - CASTEP, GAMESS, Gaussian, ONETEP, Quantum Supercharger Library*, VASP & more

green* = application where >90% of the workload is on GPU

MD vs. QC on GPUs

“Classical” Molecular Dynamics	Quantum Chemistry (MO, PW, DFT, Semi-Emp)
Simulates positions of atoms over time; chemical-biological or chemical-material behaviors	Calculates electronic properties; ground state, excited states, spectral properties, making/breaking bonds, physical properties
Forces calculated from simple empirical formulas (bond rearrangement generally forbidden)	Forces derived from electron wave function (bond rearrangement OK, e.g., bond energies)
Up to millions of atoms	Up to a few thousand atoms
Solvent included without difficulty	Generally in a vacuum but if needed, solvent treated classically (QM/MM) or using implicit methods
Single precision dominated (FP32)	Double precision is important (FP64)
Uses cuFFT, CUDA	Uses cuBLAS, cuFFT, Tensor/Eigen Solvers, OpenACC
GeForce (Workstations), Tesla (Servers)	Tesla recommended
ECC off	ECC on

Accelerating Discoveries

Using a supercomputer powered by the Tesla Platform with over 3,000 Tesla accelerators, University of Illinois scientists performed the first all-atom simulation of the HIV virus and discovered the chemical structure of its capsid — “the perfect target for fighting the infection.”

Without gpu, the supercomputer would need to be 5x larger for similar performance.

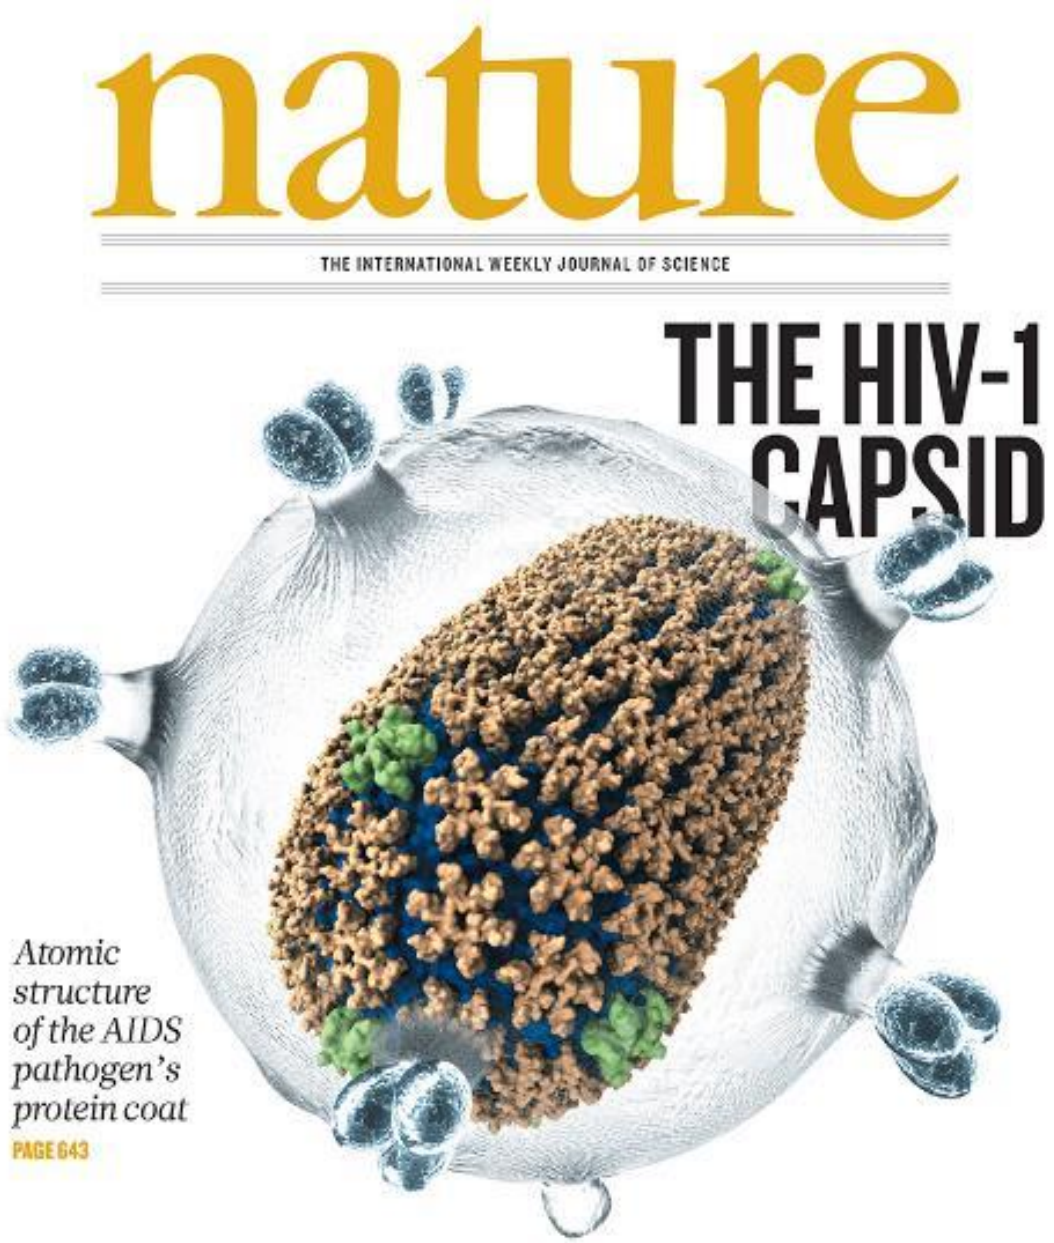

GPU-Accelerated Quantum Chemistry Apps

Green Lettering Indicates Performance Slides Included

- Abinit
- ACES III
- ADF
- BigDFT
- CP2K
- DIRAC
- GAMESS-US
- Gaussian
- GPAW
- FHI-AIMS
- LATTE
- LSDalton
- MOLCAS
- Mopac2012
- NWChem
- Octopus
- ONETEP
- Petot
- Q-Chem
- QMCPACK
- Quantum Espresso
- Quantum SuperCharger Library
- RMG
- TeraChem
- UNM
- VASP
- WL-LSMS

GPU Perf compared against dual multi-core x86 CPU socket.

ABINIT

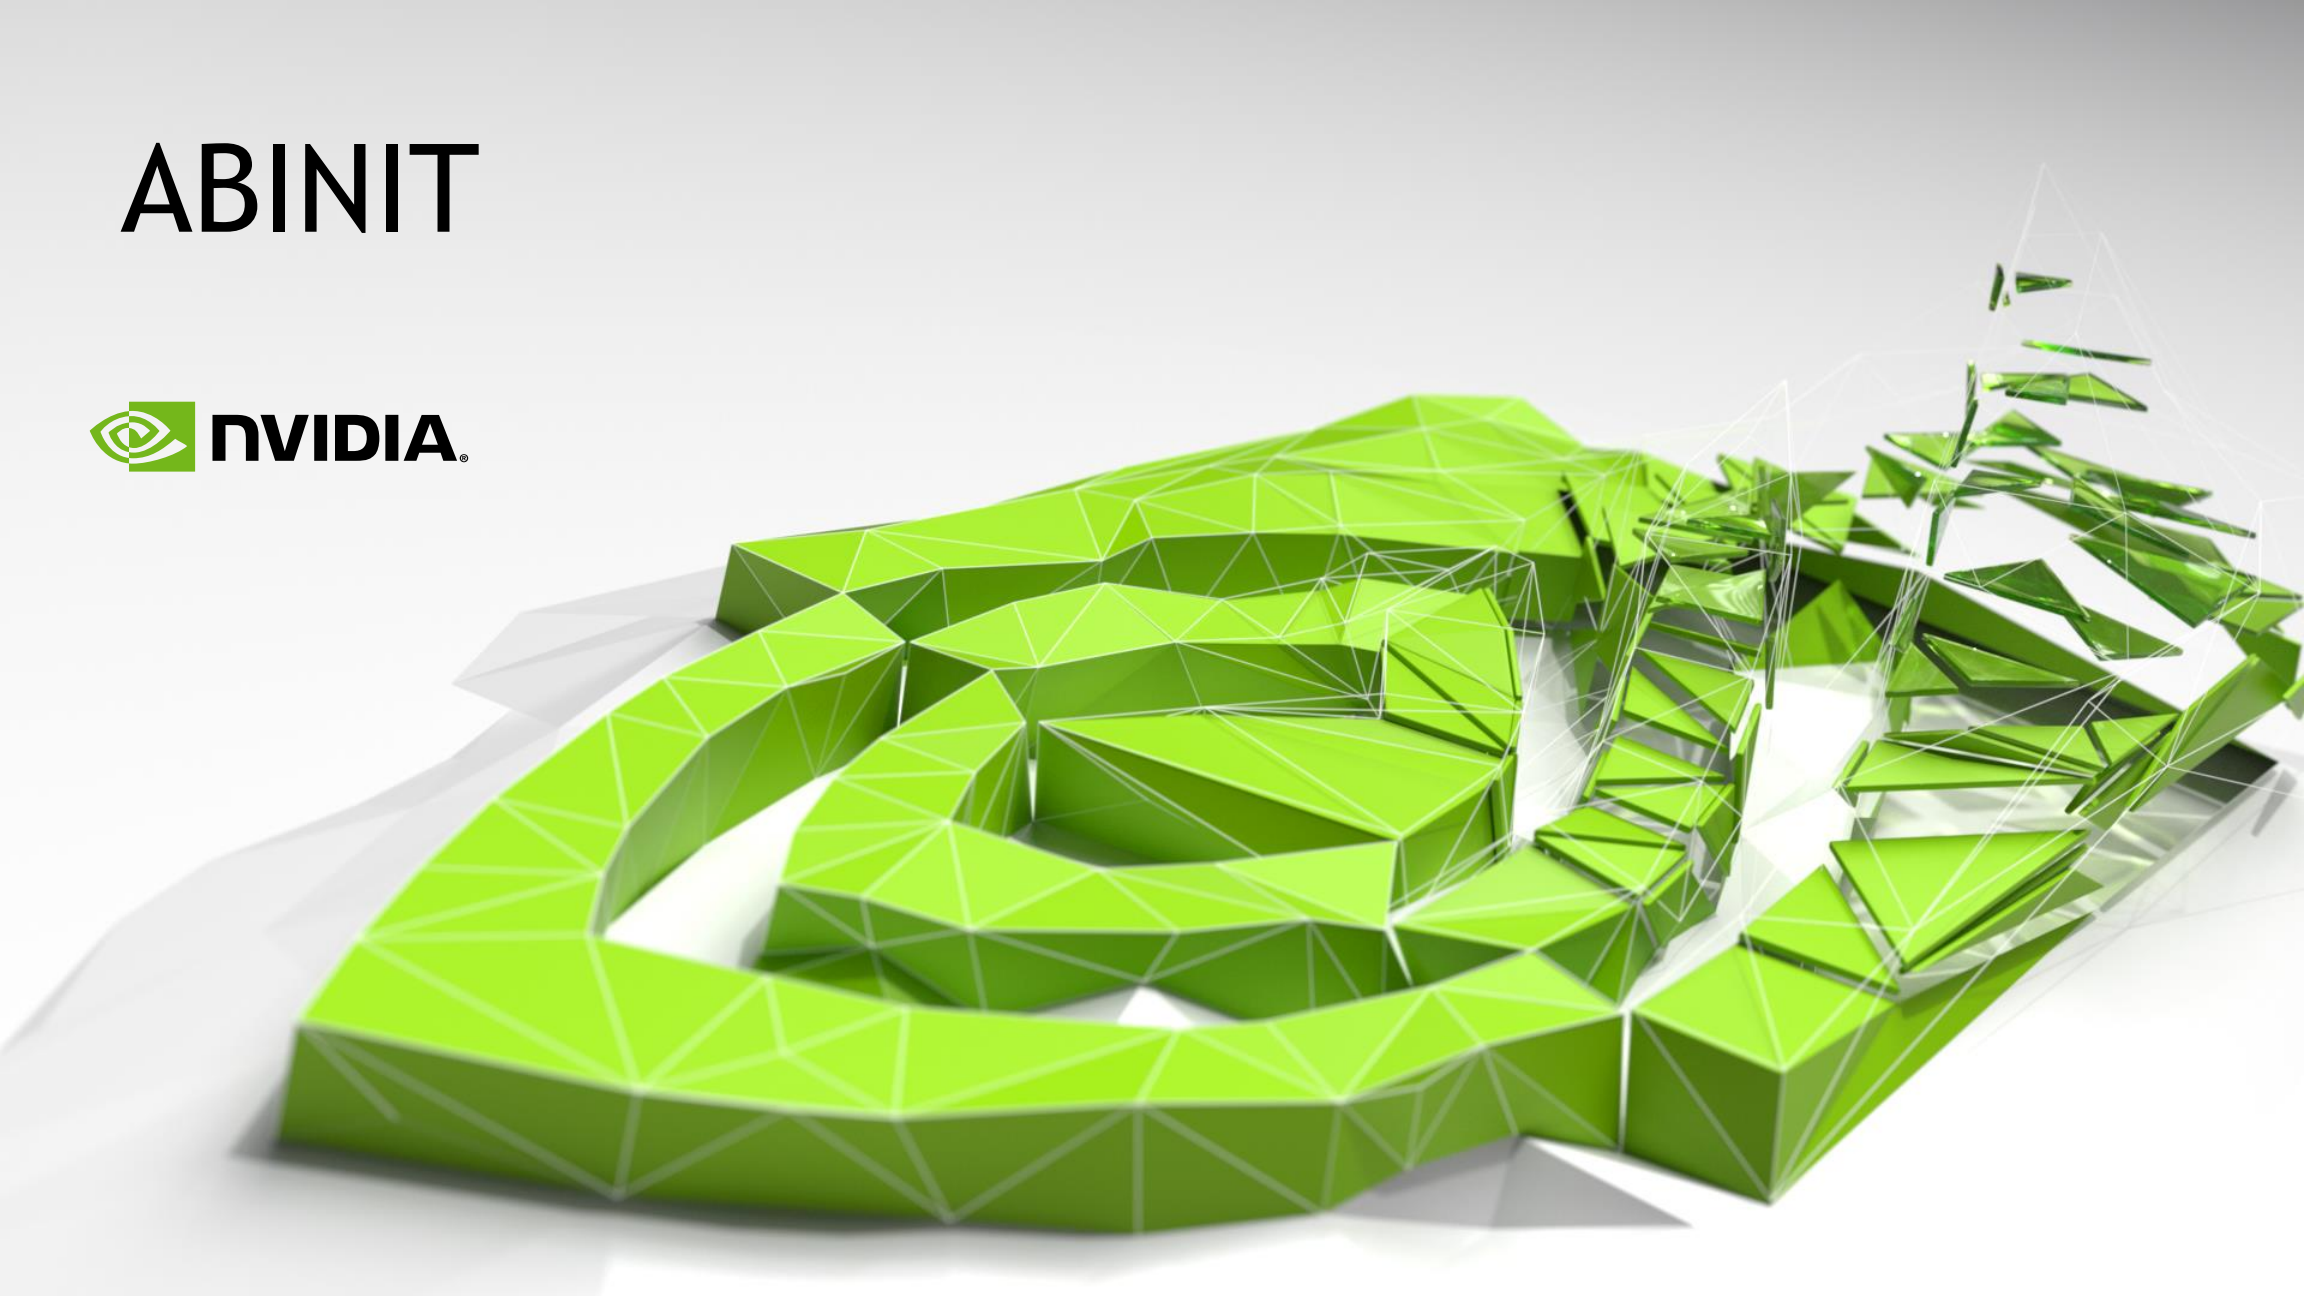

ABINIT on GPUS

- Speed in the parallel version:
  - For ground-state calculations, GPUs can be used. This is based on CUDA+MAGMA
  - For ground-state calculations, the wavelet part of ABINIT (which is BigDFT) is also very well parallelized : MPI band parallelism, combined with GPUs

BigDFT

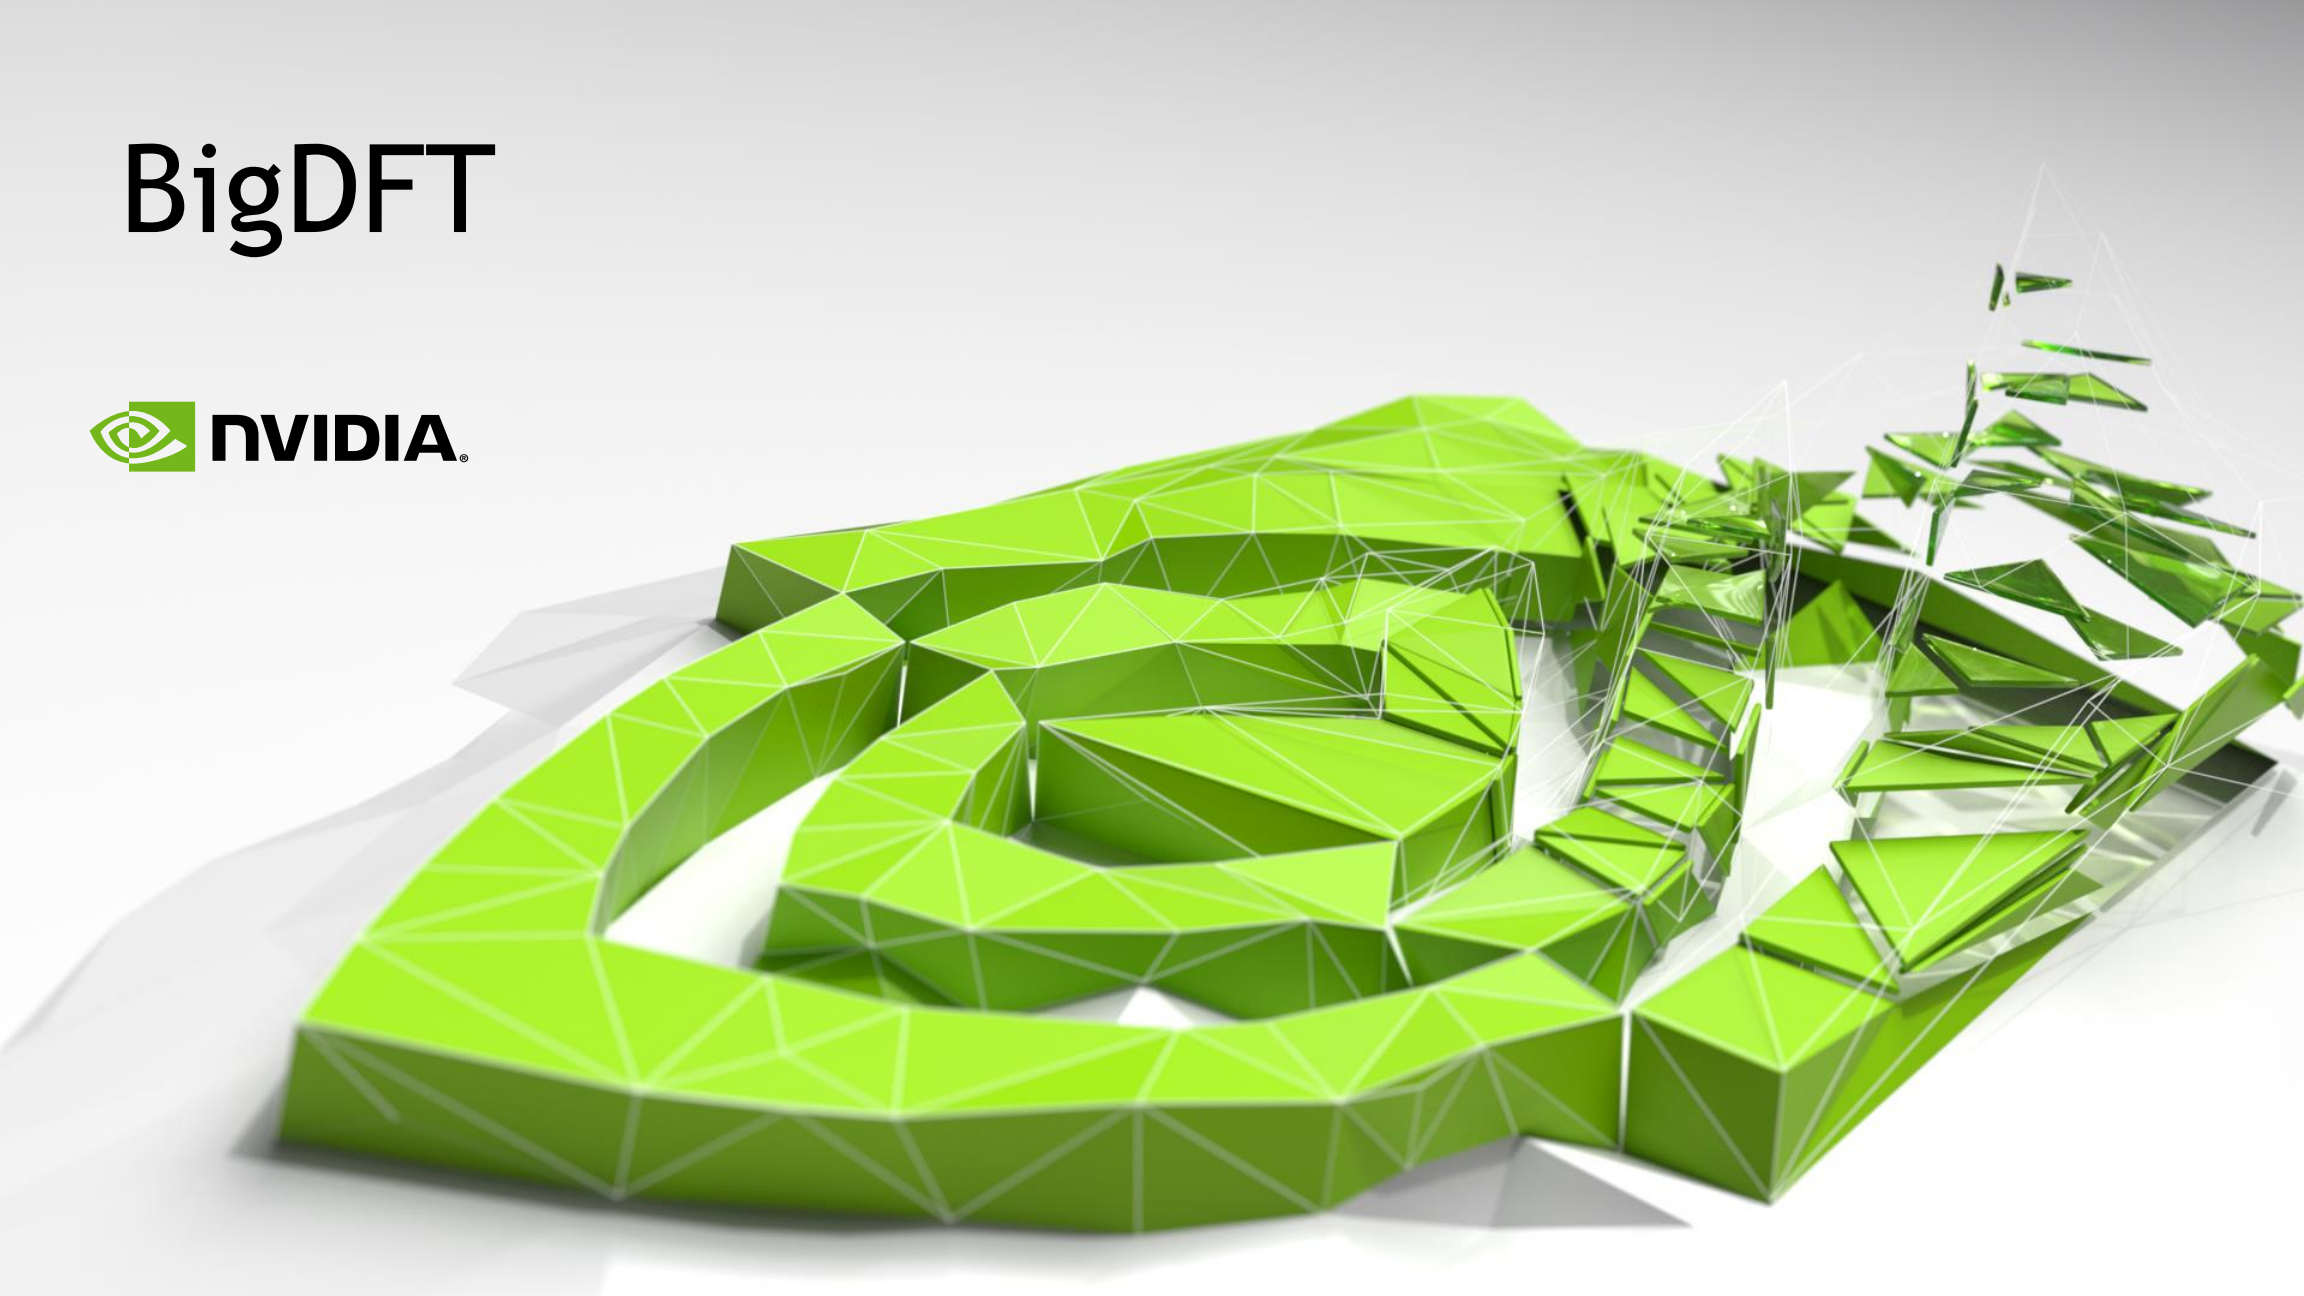

BigDFT
http://bigdft.org

Multiscale Modelling Methods for Applications in Materials Science CECAM JÜLICH, GERMANY

Introduction to Electronic Structure Calculations with BigDFT

Thierry Deutsch, Damien Caliste, Luigi Genovese

L_Sim - CEA Grenoble

17 September 2013

Courtesy of BigDFT team @ CEA

BigDFT version 1.7: capabilities

BigDFT
http://bigdft.org

http://bigdft.org

- Free, surface and periodic boundary conditions
- Geometry optimizations (with constraints)
- Born-Oppenheimer Molecular Dynamics
- Saddle point searches (Nudged-Elastic Band Method)
- Vibrations
- External electric fields
- Unoccupied KS orbitals
- Collinear and Non-collinear magnetism
- All XC functionals of the ABINIT package
- Hybrid functionals
- Empirical van der Waals interactions (many flavors)
- Also available within the ABINIT package

Courtesy of BigDFT team @ CEA

BigDFT version 1.7: capabilities

BigDFT
http://bigdft.org

http://bigdft.org

- Free, surface and periodic boundary conditions
- Geometry optimizations (with constraints)
- Born-Oppenheimer Molecular Dynamics
- Saddle point searches (Nudged-Elastic Band Method)
- Vibrations
- External electric fields
- Unoccupied KS orbitals
- Collinear and Non-collinear magnetism
- All XC functionals of the ABINIT package
- Hybrid functionals
- Empirical van der Waals interactions (many flavors)
- Also available within the ABINIT package

Courtesy of BigDFT team @ CEA

GPU-ported operations in BigDFT (double precision)

BigDFT
http://bigdft.org

Convolutions (OpenCL rewritten)

GPU speedups between 10 and 20 can be obtained for different sections

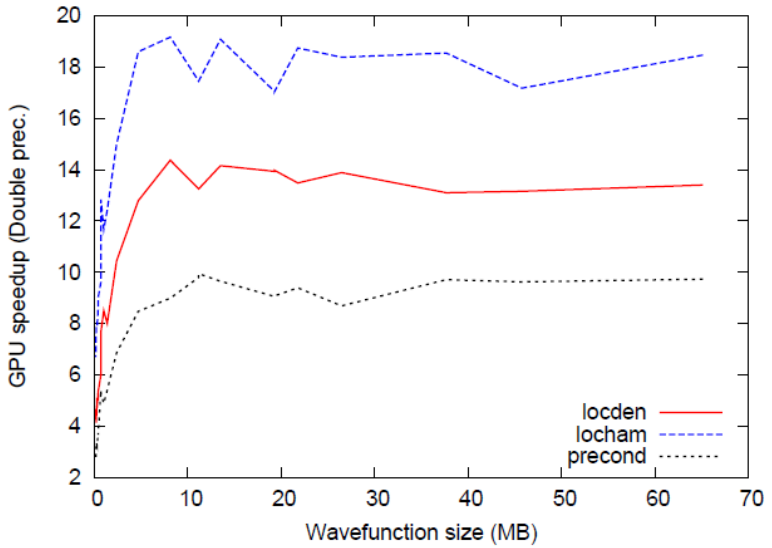

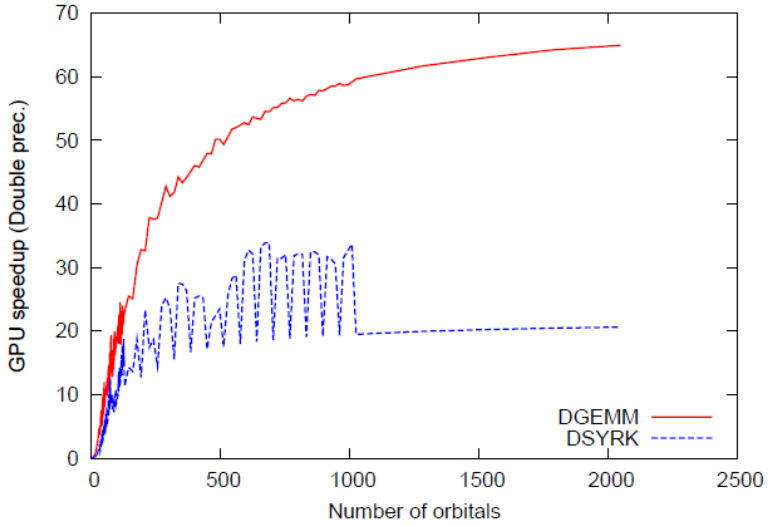

Linear algebra (CUBLAS library)

The interfacing with CUBLAS is immediate, with considerable speedups

Courtesy of BigDFT team @ CEA

BigDFT code on Hybrid architectures

BigDFT code can run on hybrid CPU/GPU supercomputers
In multi-GPU environments, double precision calculations

No Hot-spot operations

Different code sections can be ported on GPU
up to 20x speedup for some operations,
7x for the full parallel code

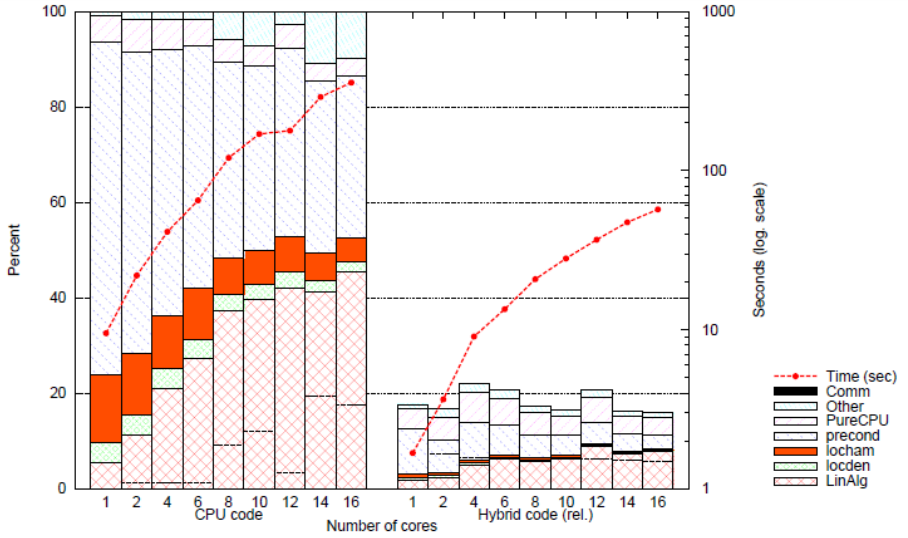

Courtesy of BigDFT team @ CEA

Hands on

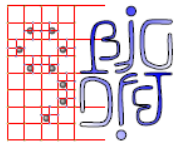

BigDFT
http://bigdft.org

See
http://bigdft.org/Wiki/index.php?title=Category:Tutorials

- First runs with BigDFT
- Basis-set convergence
- Acceleration example on different platforms:
  Kohn-Sham DFT Operation with GPU acceleration

Courtesy of BigDFT team @ CEA

GAMESS

February 2018

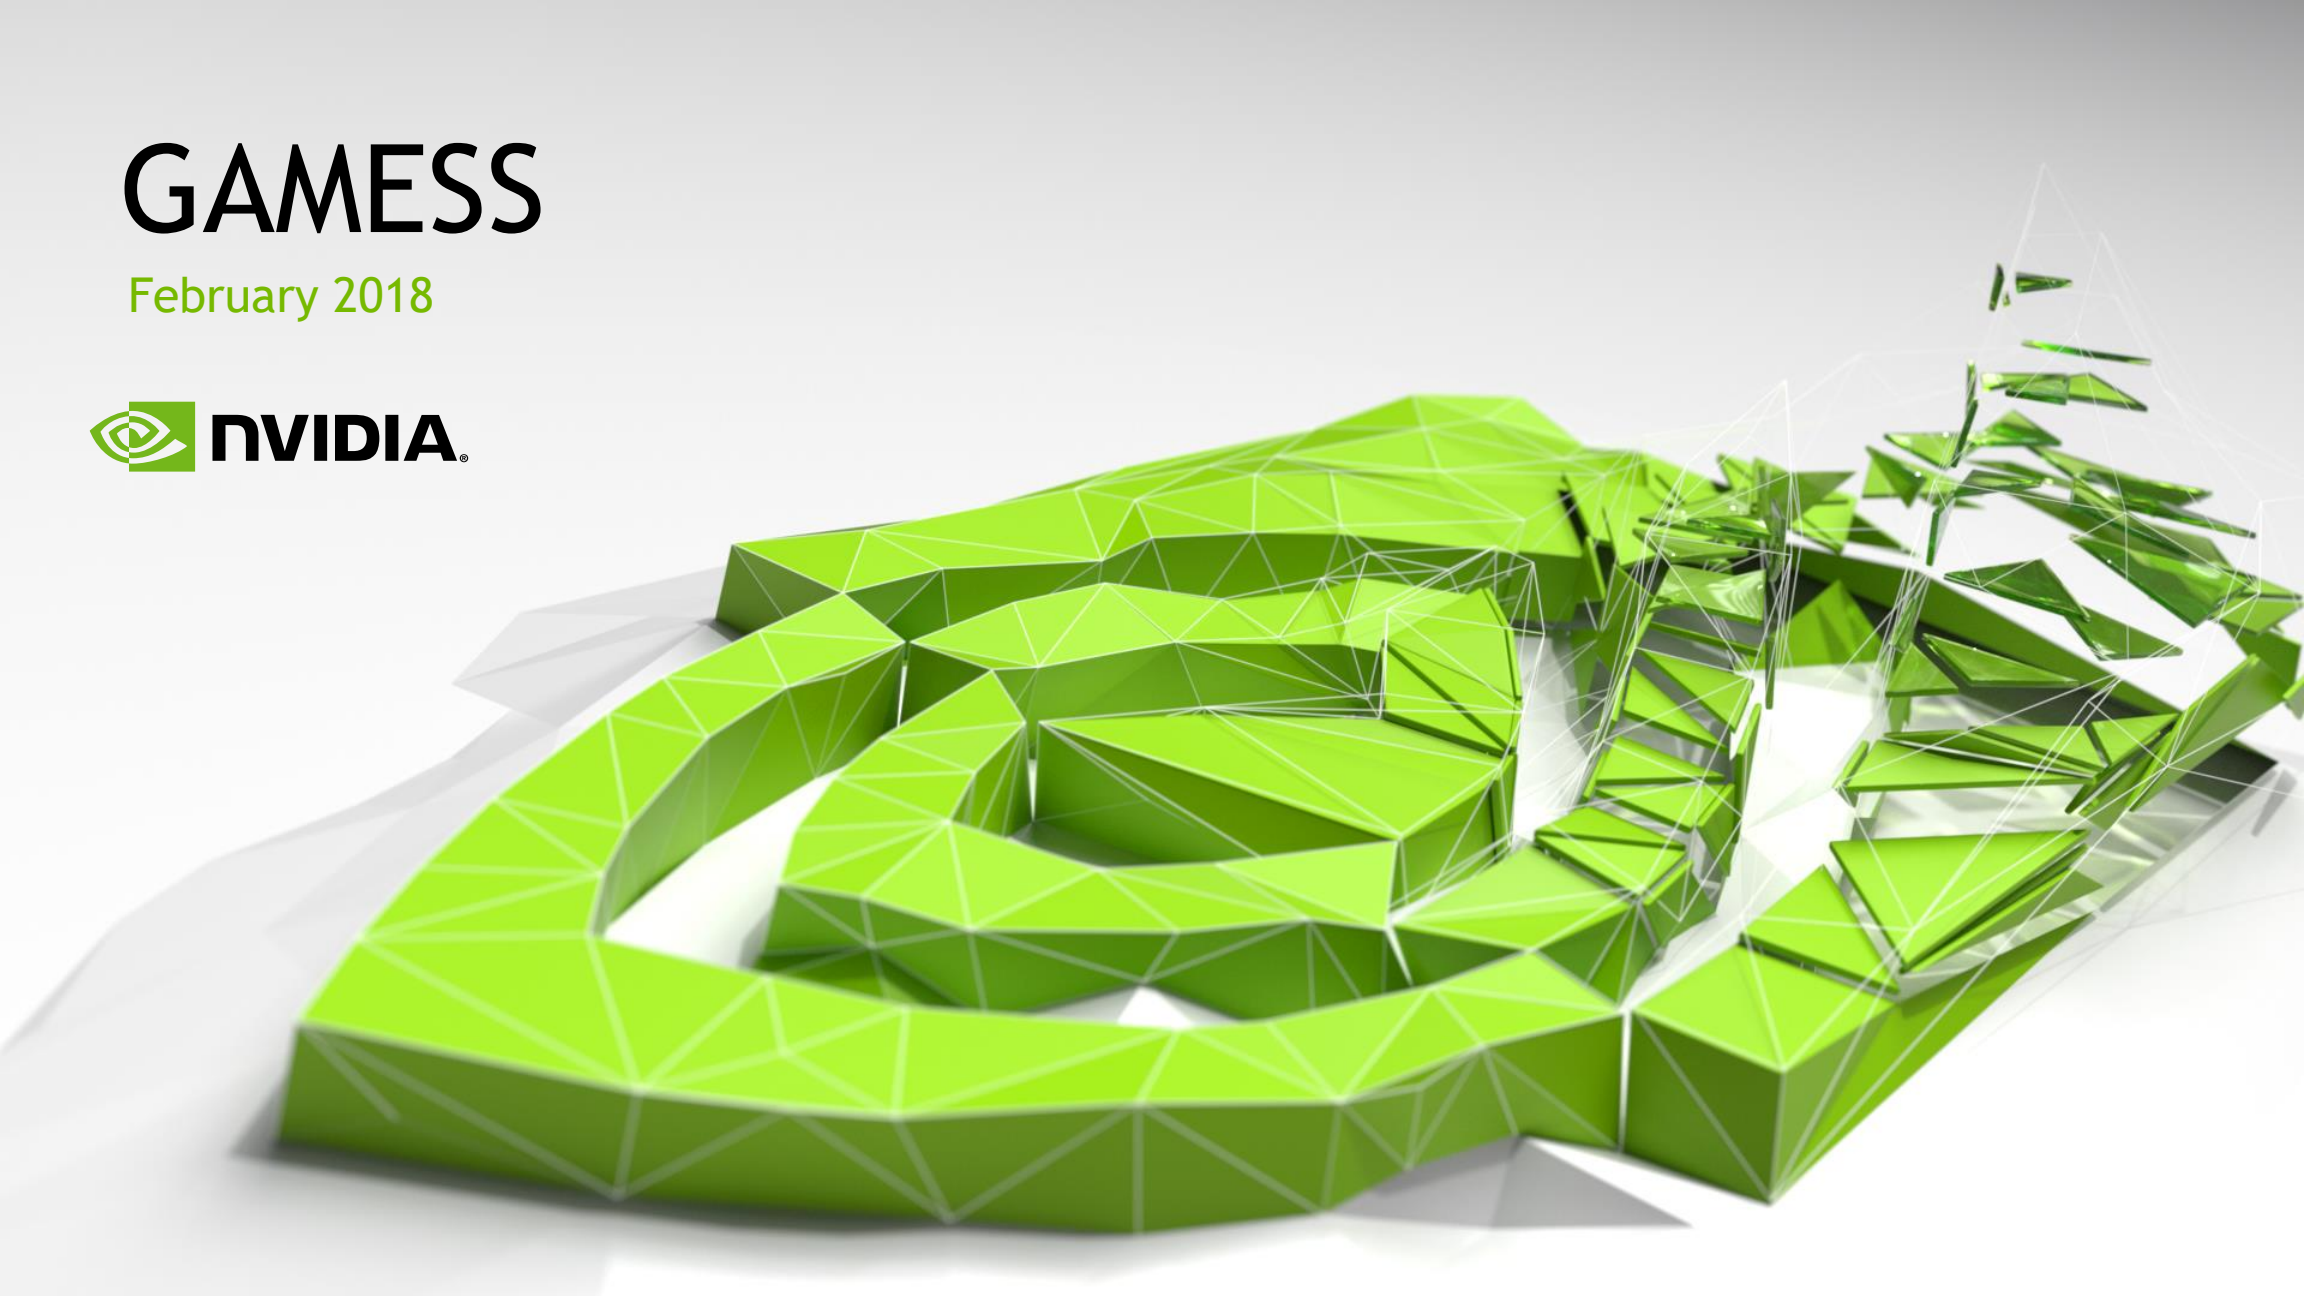

GAMESS valinomycin rimp2 energy
on V100 vs P100 (PCIe 16GB)

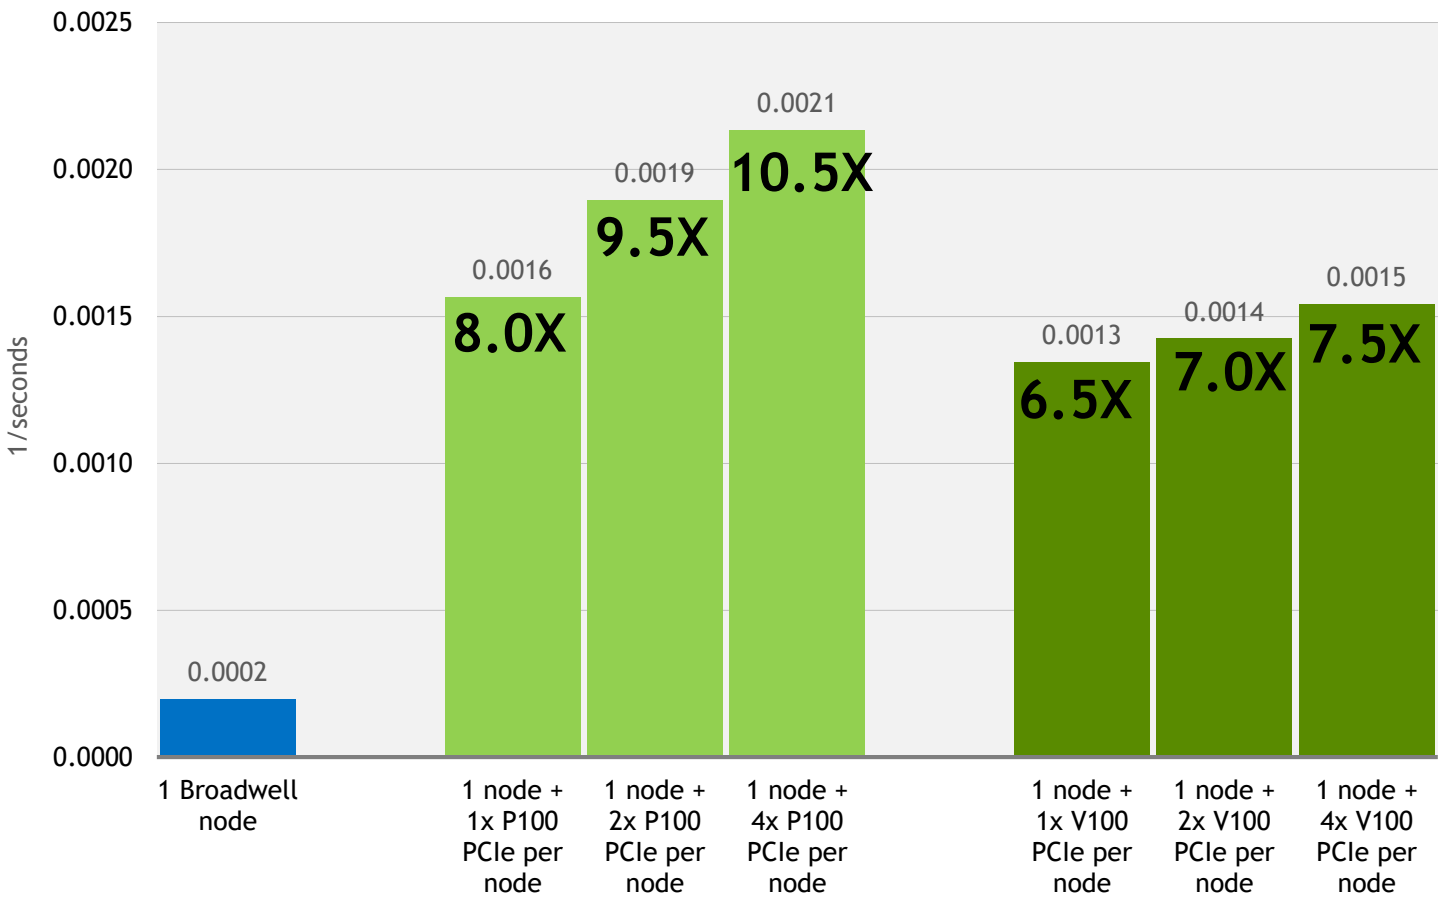

(Untuned on Volta)
Running GAMESS

The blue node contains Dual Intel Xeon E5-2690 v4@2.6GHz [3.5GHz Turbo] (Broadwell) CPUs

The green nodes contain Dual Intel Xeon E5-2690 v4@2.6GHz [3.5GHz Turbo] (Broadwell) CPUs + Tesla P100 PCIe (16GB) or V100 PCIe (16GB) GPUs

Gaussian 16

March 2018

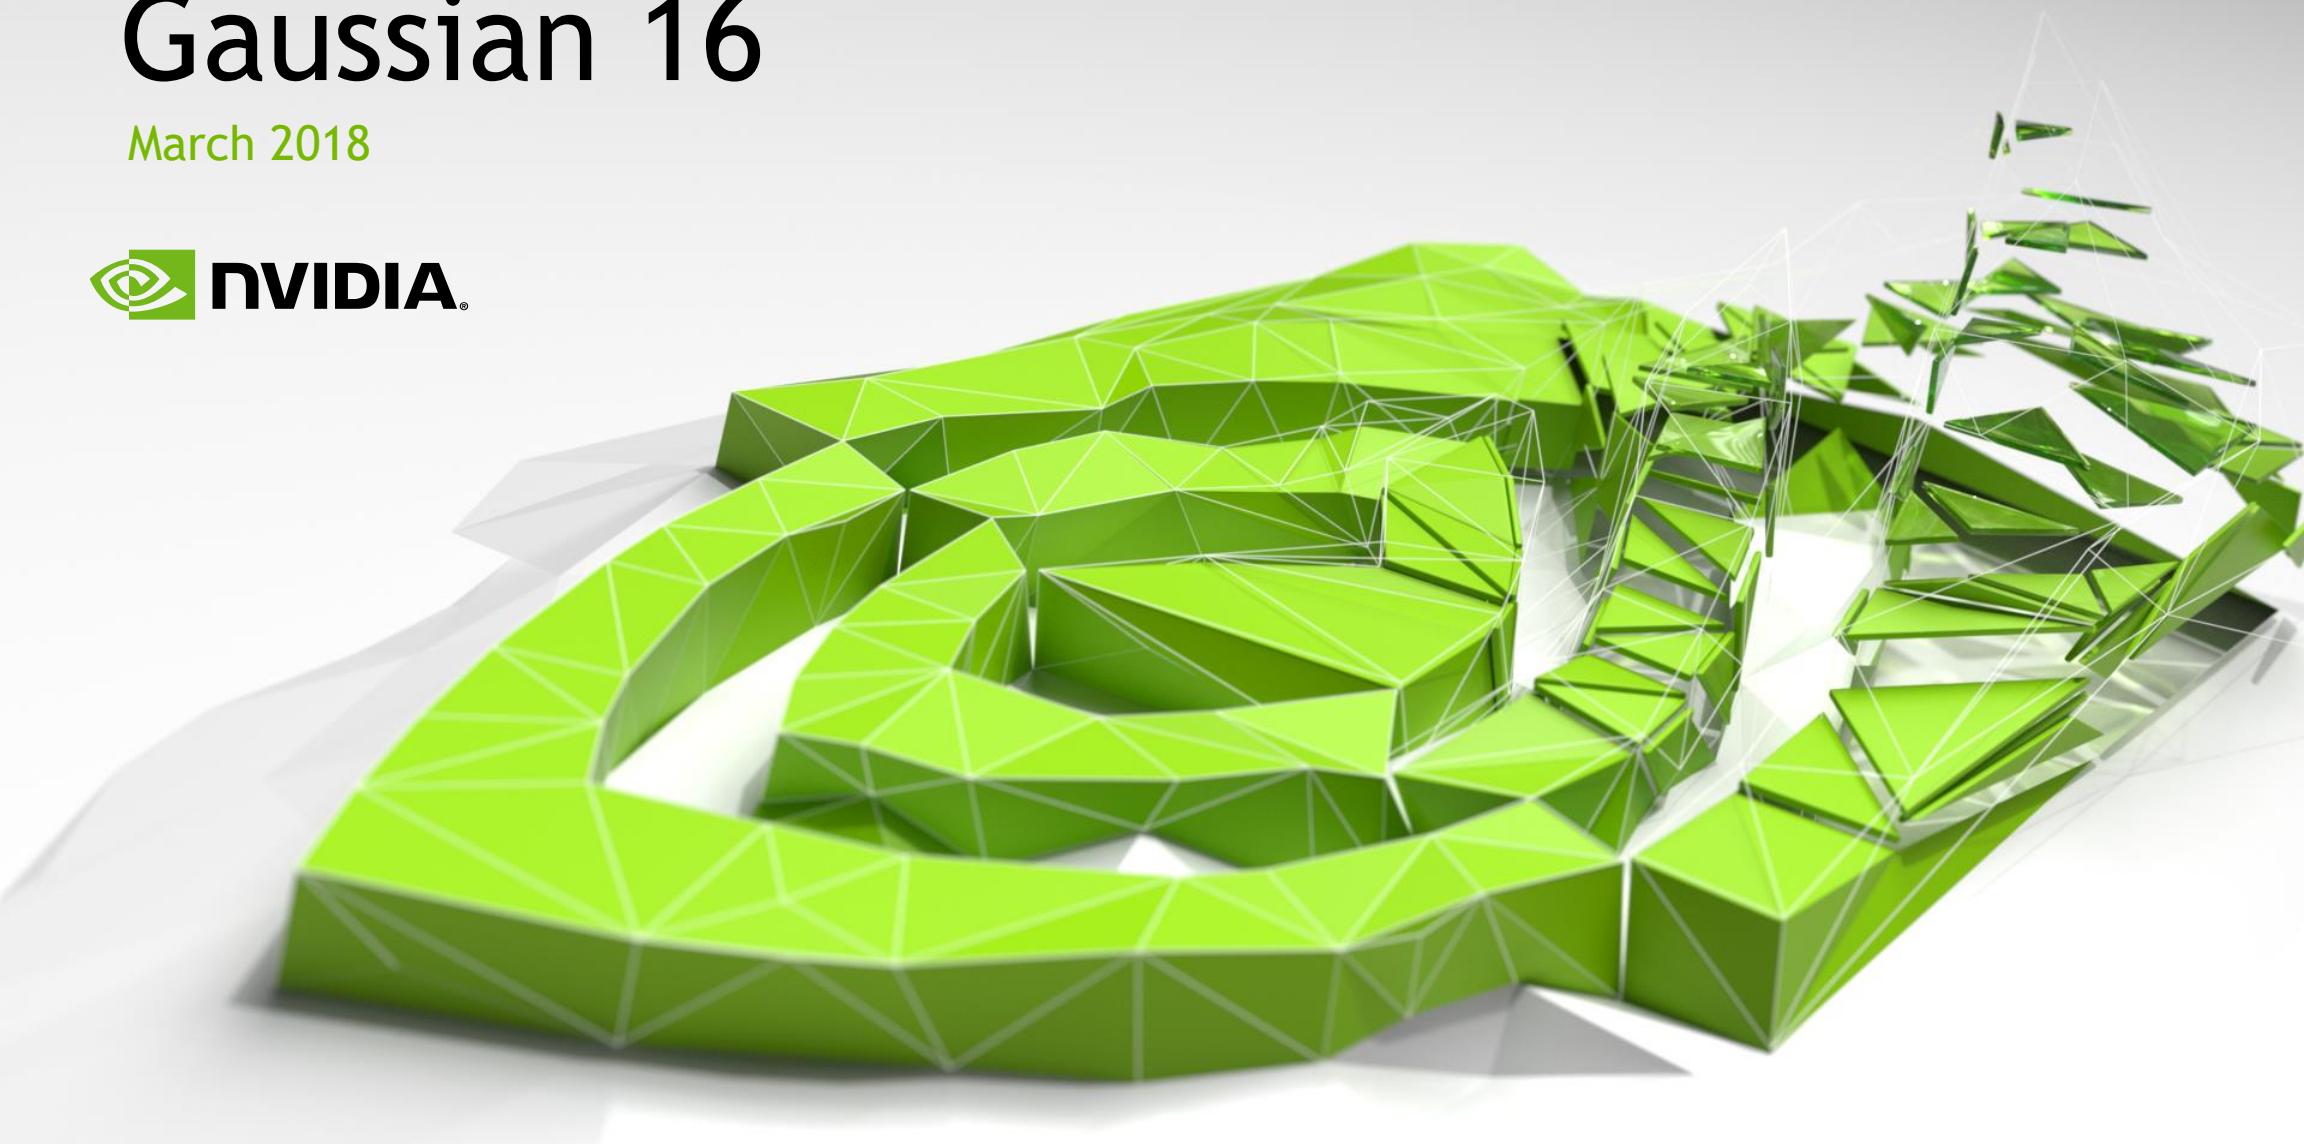

GAUSSIAN 16

A Leading Computation Chemistry Code

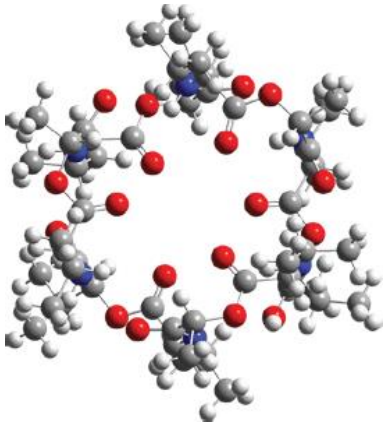

Valinomycin
Frequency. APFD 6-311+G(2d,p)
7X speedup on 8X P100 GPUs

Hardware: HPE Apollo 6500 server with dual Intel Xeon E5-2680 v4 CPUs (2.40GHz; 14 cores/chip, 28 cores total), 256GB memory and 8 Tesla P100 GPU boards (autoboost clocks). Gaussian source code compiled with PGI Accelerator Compilers (17.7) with OpenACC (2.5 standard).

Mike Frisch, Ph.D.
President and CEO
Gaussian, Inc.

"Using OpenACC allowed us to continue development of our fundamental algorithms and software capabilities simultaneously with the GPU-related work. In the end, we could use the same code base for SMP, cluster/ network and GPU parallelism. PGI's compilers were essential to the success of our efforts."

PGI

Effects of using K80s

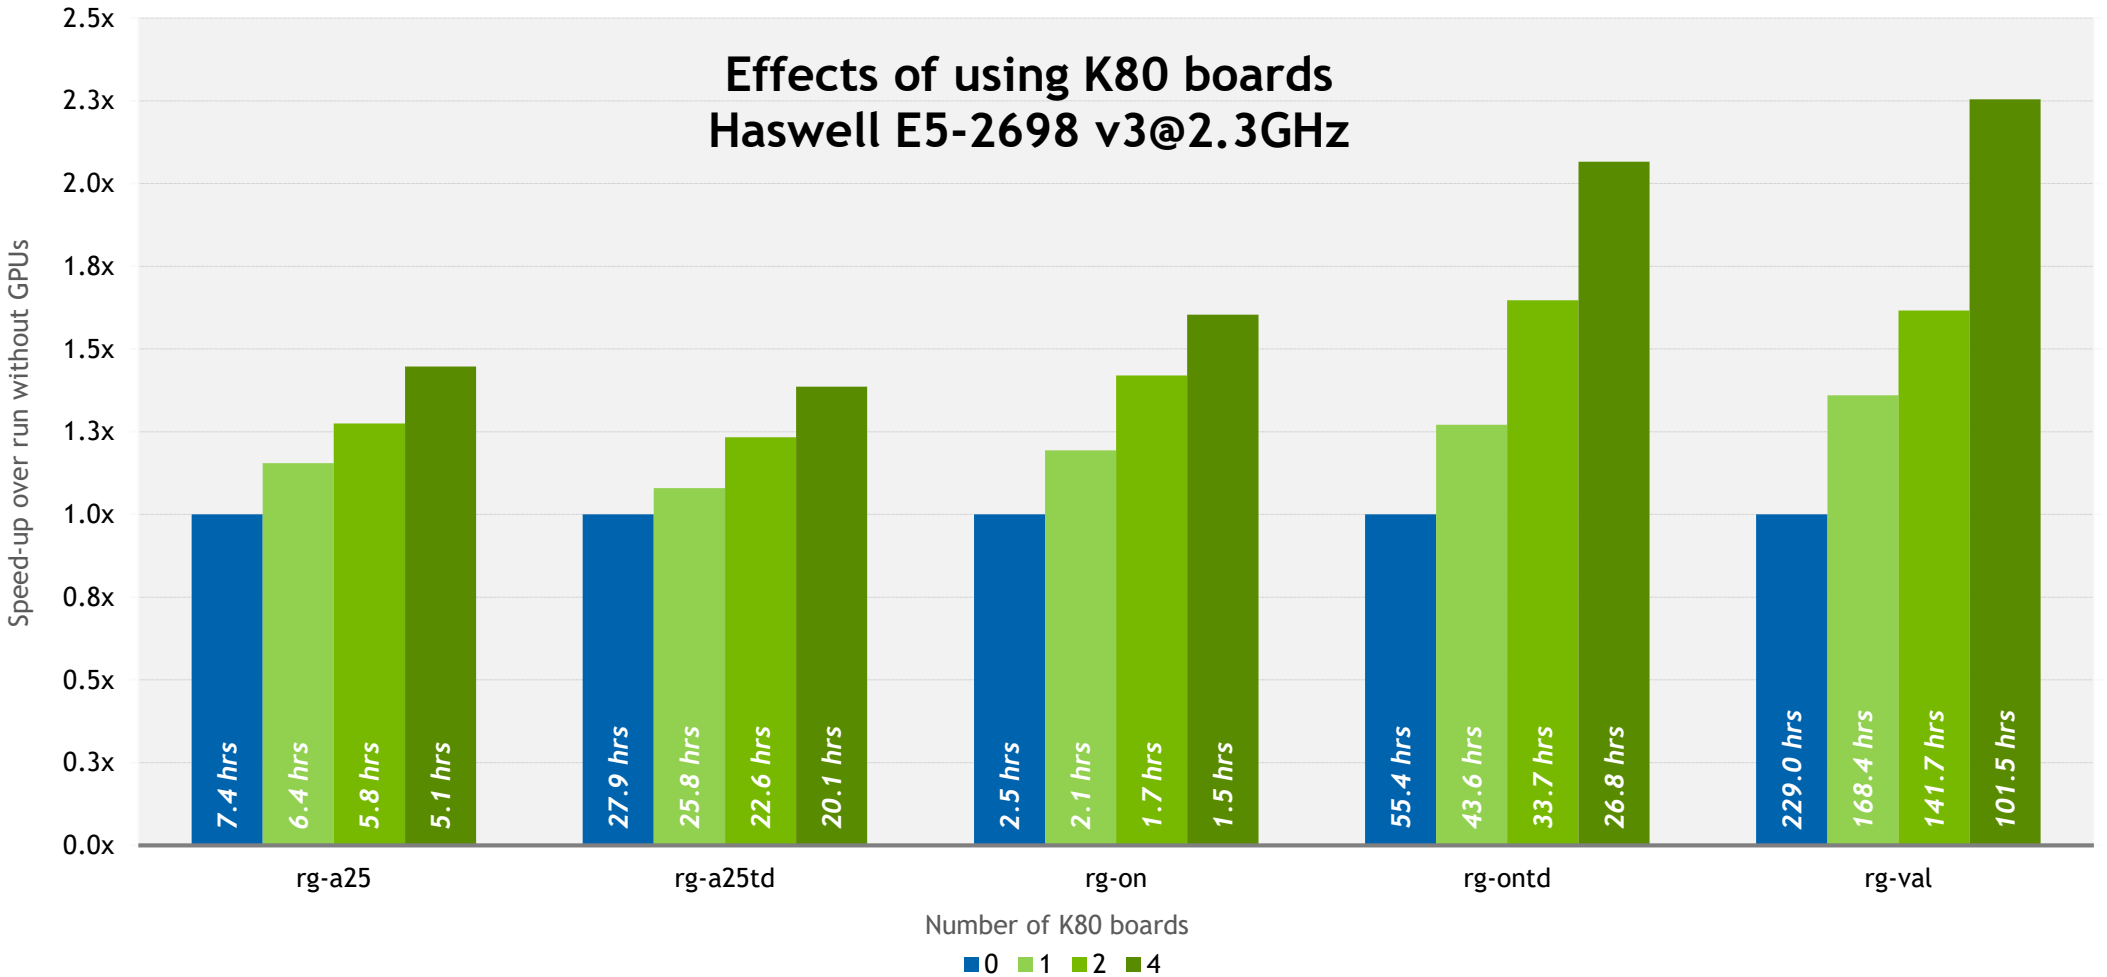

GAUSSIAN 16 P100 PERFORMANCE

Dual-socket 28-core Broadwell vs 8x P100 GPUs

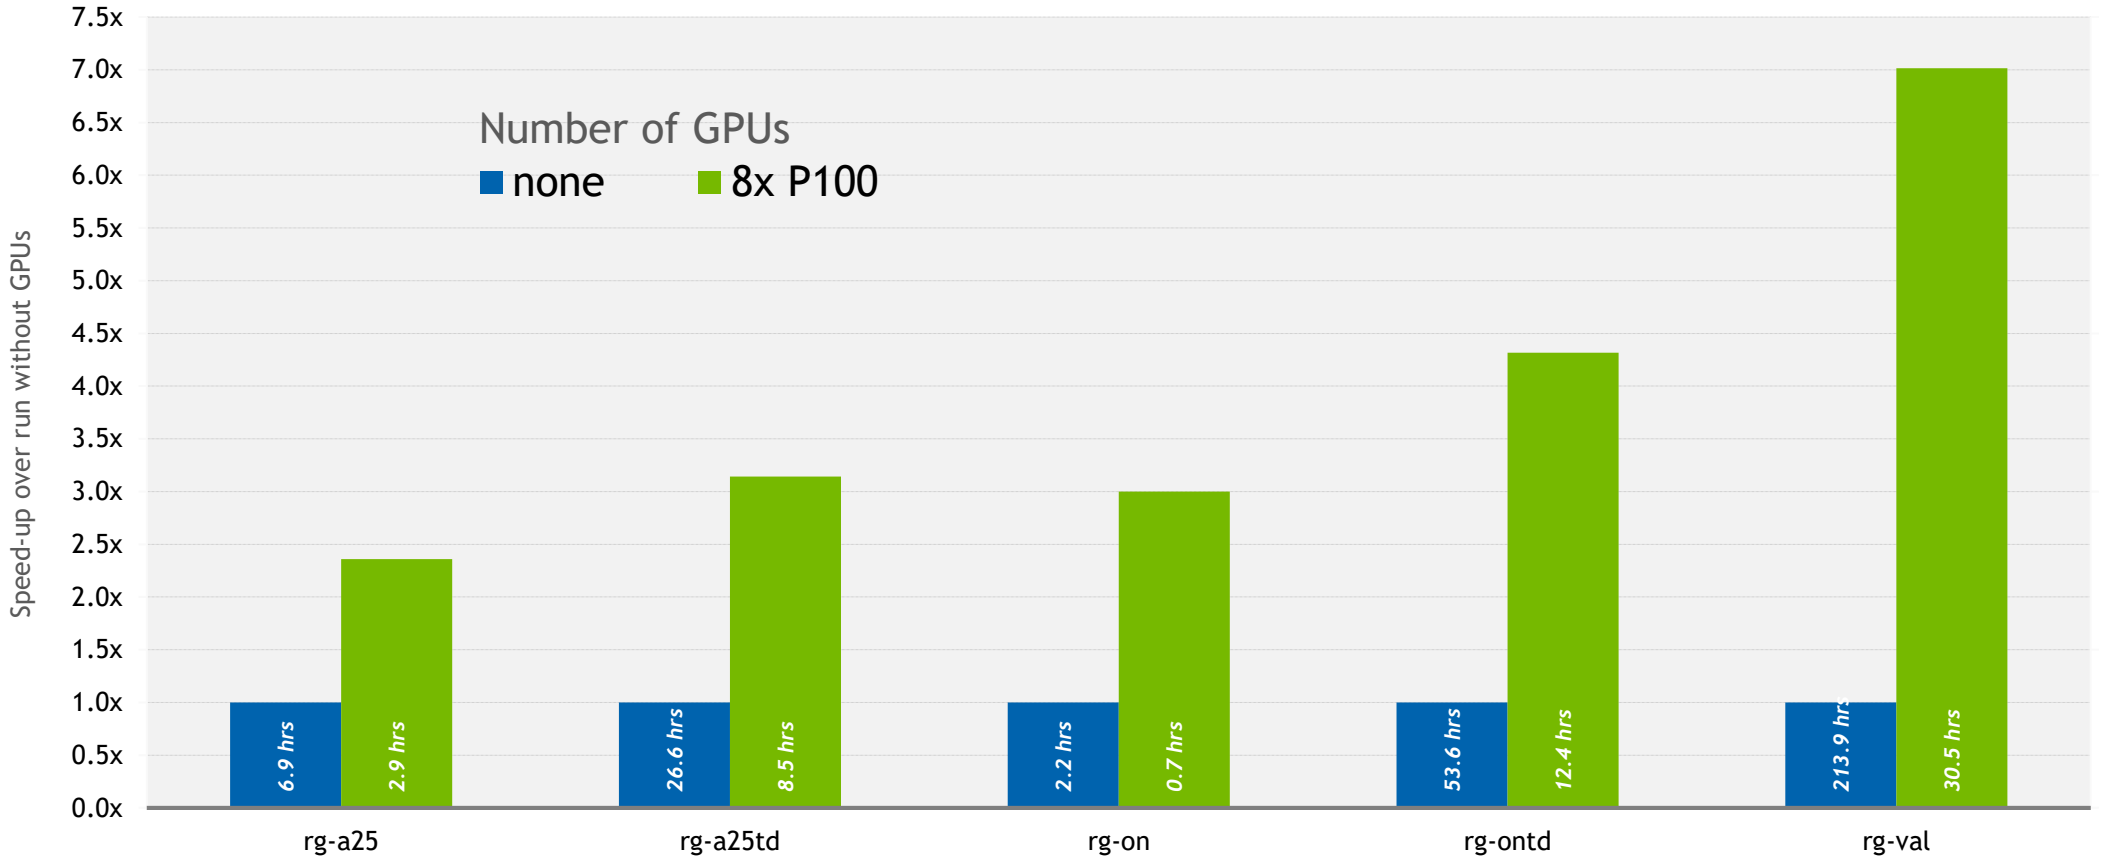

GPU-ACCELERATED GAUSSIAN 16 AVAILABLE

100% PGI OpenACC Port (no CUDA)

- Gaussian is a Top Ten HPC (Quantum Chemistry) Application.
- **80-85% of use cases are GPU-accelerated** (Hartree-Fock and DFT: energies, 1st derivatives (gradients) and 2nd derivatives for ground & excited states). More functionality to come.
- K40, K80, P100 support; B.01 release.
- **No pricing difference** between Gaussian CPU and GPU versions.
- Existing Gaussian 09 **customers under maintenance contract get (free) upgrade**.
- Existing non-maintenance customers required to pay upgrade fee.
- To get the bits or to ask about the upgrade fee, please contact Gaussian, Inc.'s Jim Hess, Operations Manager; jhess@gaussian.com.

rg-a25 on K80s

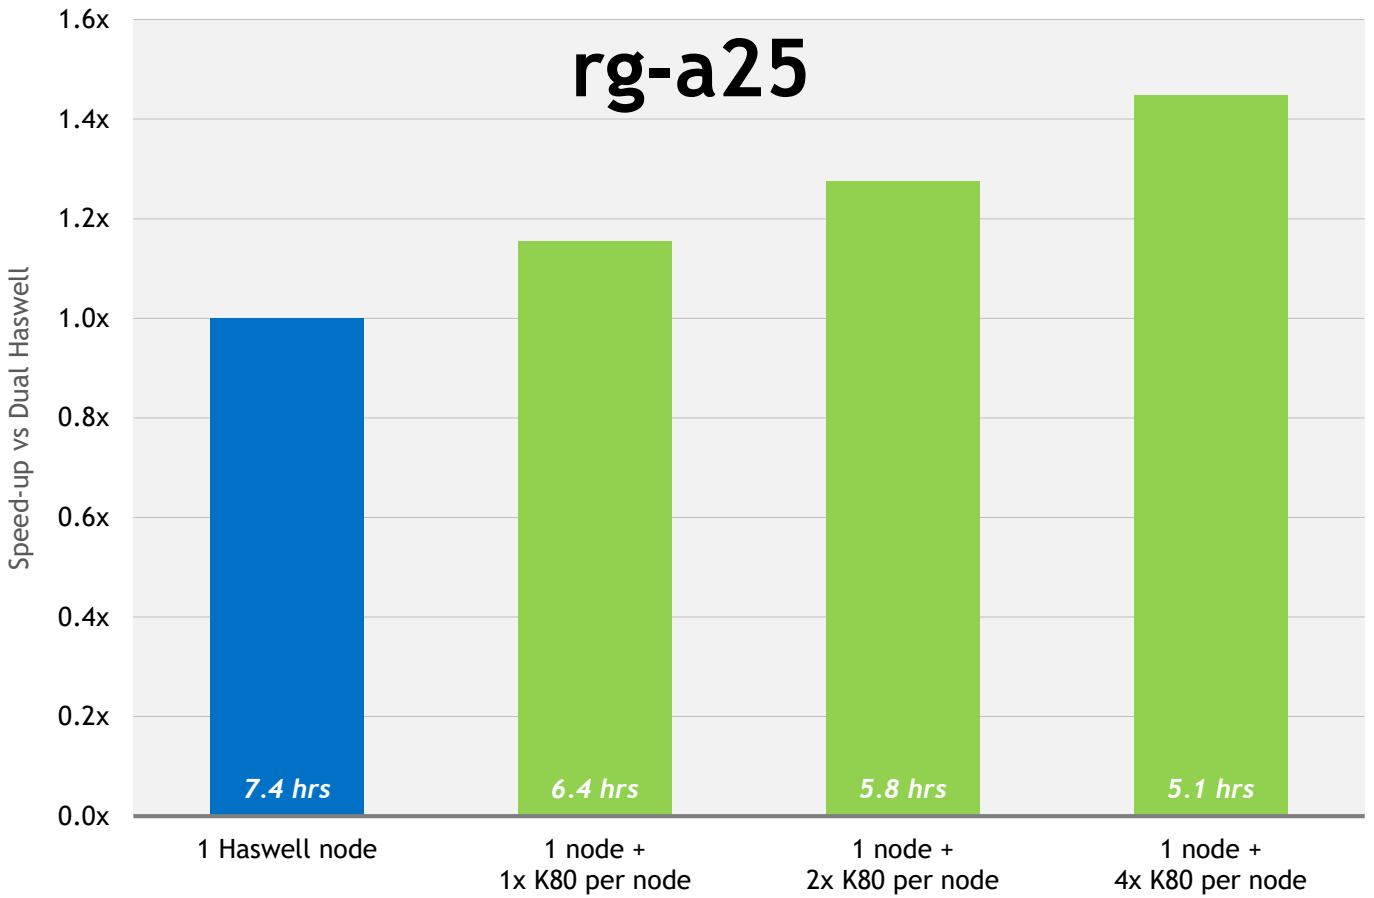

Running Gaussian version 16

The blue node contains Dual Intel Xeon E5-2698 v3@2.3GHz (Haswell) CPUs

The green nodes contain Dual Intel Xeon E5-2698 v3@2.3GHz (Haswell) CPUs + Tesla K80 (autoboost) GPUs

Alanine 25. Two steps: Force and Frequency. APFD 6-31G*
nAtoms = 259, nBasis = 2195

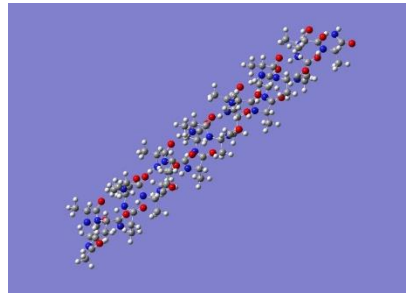

rg-a25td on K80s

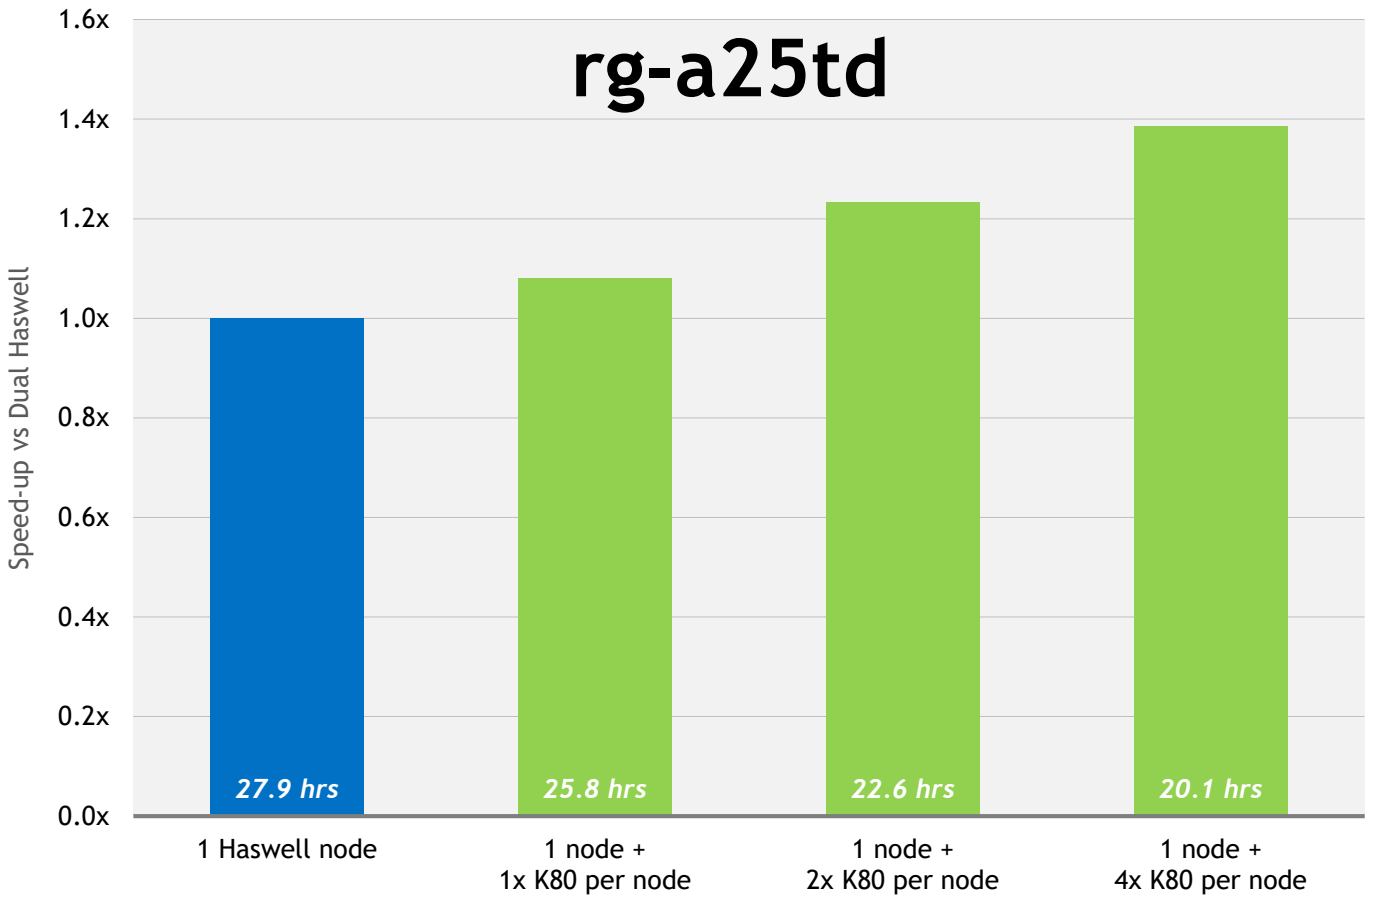

Running Gaussian version 16

The blue node contains Dual Intel Xeon E5-2698 v3@2.3GHz (Haswell) CPUs

The green nodes contain Dual Intel Xeon E5-2698 v3@2.3GHz (Haswell) CPUs + Tesla K80 (autoboost) GPUs

Alanine 25. Two Time-Dependent (TD) steps: Force and Frequency. APFD 6-31G* nAtoms = 259, nBasis = 2195

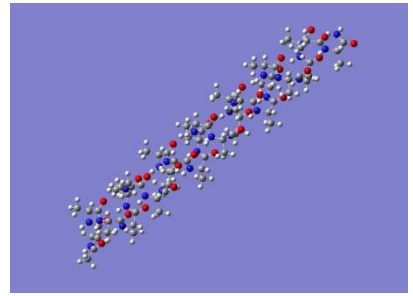

rg-on on K80s

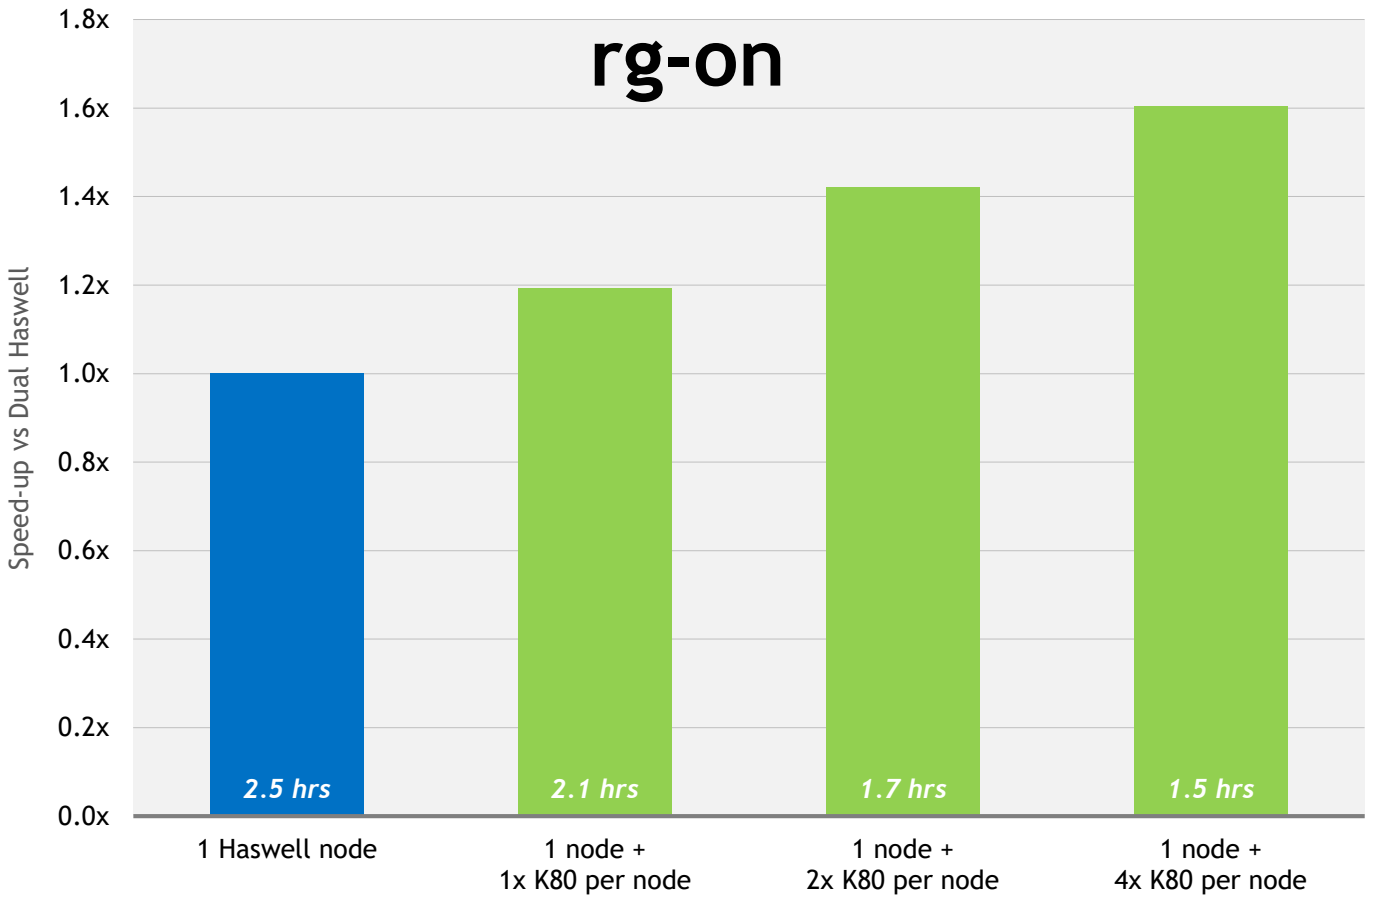

Running Gaussian version 16

The blue node contains Dual Intel Xeon E5-2698 v3@2.3GHz (Haswell) CPUs

The green nodes contain Dual Intel Xeon E5-2698 v3@2.3GHz (Haswell) CPUs + Tesla K80 (autoboost) GPUs

GFP ONIOM. Two steps: Force and Frequency. APFD/6-311+G(2d,p):amber=softfirst)=embed
nAtoms = 3715 (48/3667), nBasis = 813

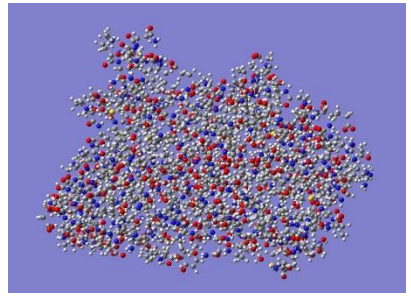

rg-ontd on K80s

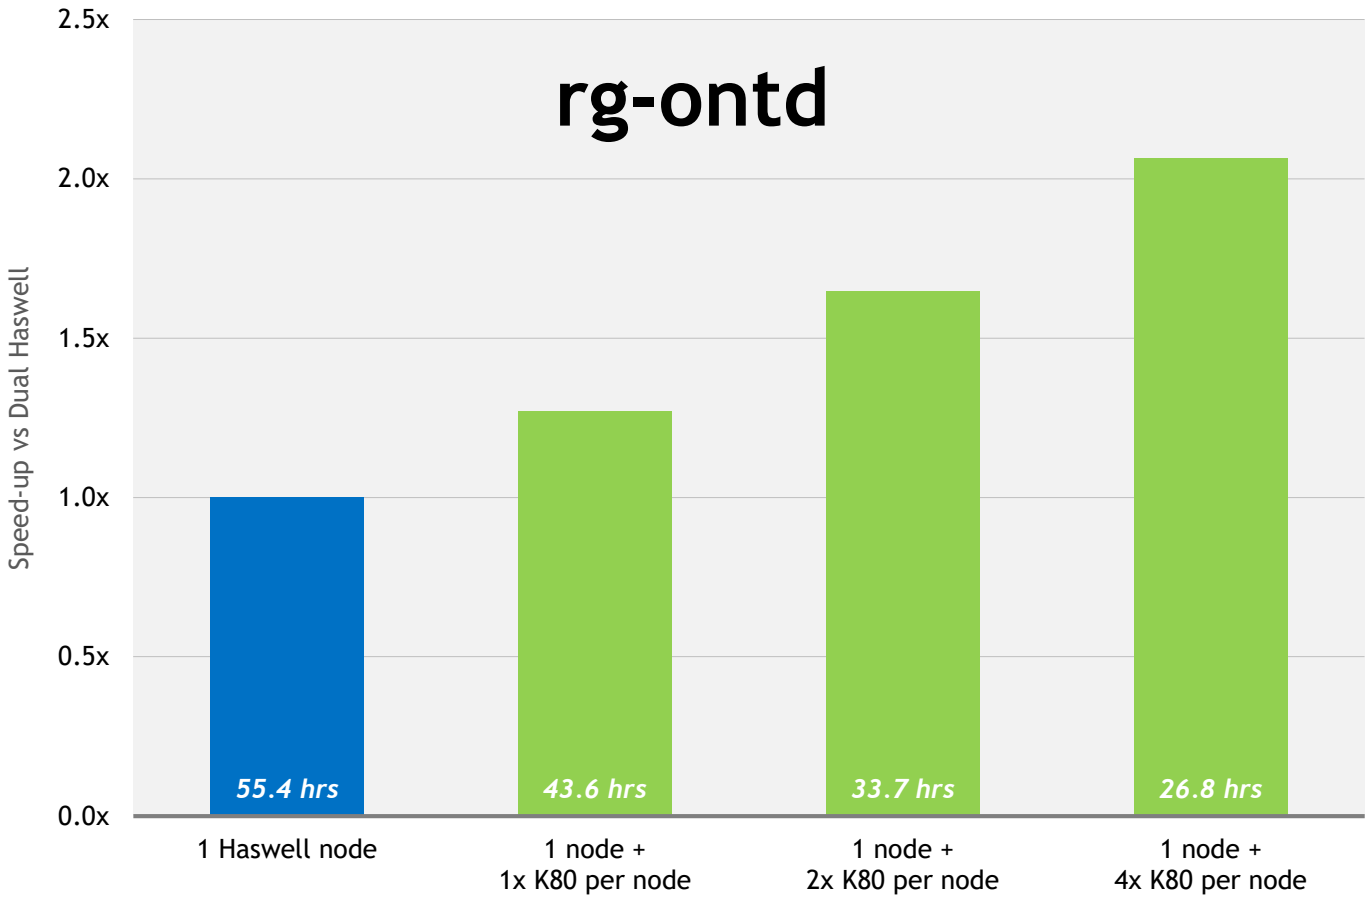

Running Gaussian version 16

The blue node contains Dual Intel Xeon E5-2698 v3@2.3GHz (Haswell) CPUs

The green nodes contain Dual Intel Xeon E5-2698 v3@2.3GHz (Haswell) CPUs + Tesla K80 (autoboost) GPUs

GFP ONIOM. Two Time-Dependent (TD) steps: Force and Frequency. APFD/6-311+G(2d,p):amber=softfirst)=embed nAtoms = 3715 (48/3667), nBasis = 813

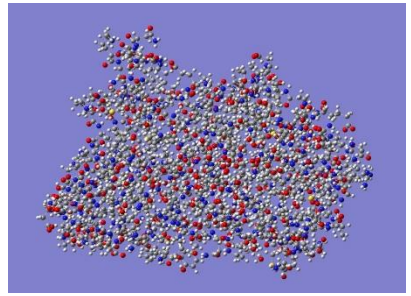

rg-val on K80s

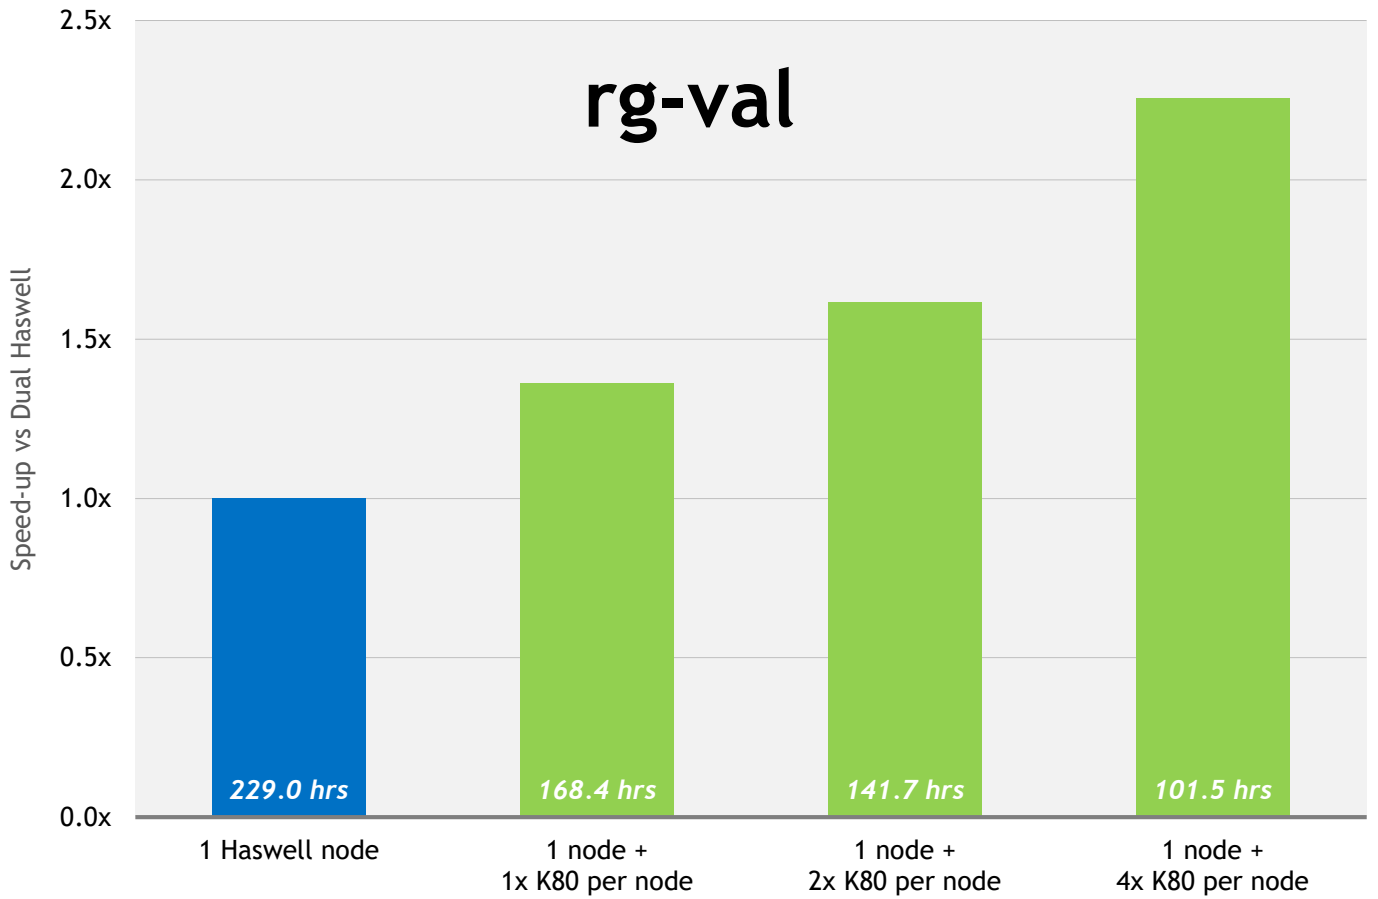

Running Gaussian version 16

The blue node contains Dual Intel Xeon E5-2698 v3@2.3GHz (Haswell) CPUs

The green nodes contain Dual Intel Xeon E5-2698 v3@2.3GHz (Haswell) CPUs + Tesla K80 (autoboost) GPUs

Valinomycin. Two steps: Force and Frequency. APFD 6-311+G(2d,p) nAtoms = 168, nBasis = 2646

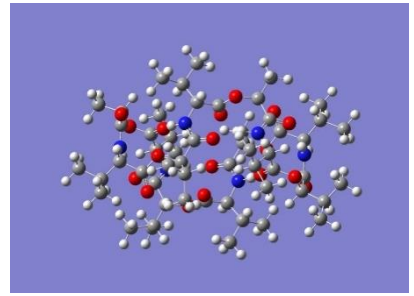

Gaussian 16 Supported Platforms

- 4-way collaboration; Gaussian, Inc., NVIDIA (& PGI) and HPE
- HPE, NVIDIA and PGI is the development platform
- All released/certified x86_64 versions of Gaussian 16 use the PGI compilers
- Certified versions of Gaussian 16 use Intel only for Itanium, XLF for some IBM platforms, Fujitsu compilers for some SPARC-based machines and PGI for the rest (including some Apple products)
- GINC is collaborating with IBM, PGI (and NVIDIA) to release an OpenPower version of Gaussian that also uses the PGI compiler
- See Gaussian Supported Platforms for more details: http://gaussian.com/g16/g16_plat.pdf

CLOSING REMARKS

Significant Progress has been made in enabling Gaussian on GPUs with OpenACC

OpenACC is increasingly becoming more versatile

Significant work lies ahead to improve performance

Expand feature set:

- PBC, Solvation, MP2, triples-Corrections

ACKNOWLEDGEMENTS

Development is taking place with:

Hewlett-Packard (HP) Series SL2500 Servers (Intel® Xeon® E5-2680 v2 (2.8GHz/10-core/25MB/8.0GT-s QPI/115W, DDR3-1866)

NVIDIA® Tesla® GPUs (V100)

PGI Accelerator Compilers (18.x) with OpenACC (2.5 standard)

GPAW

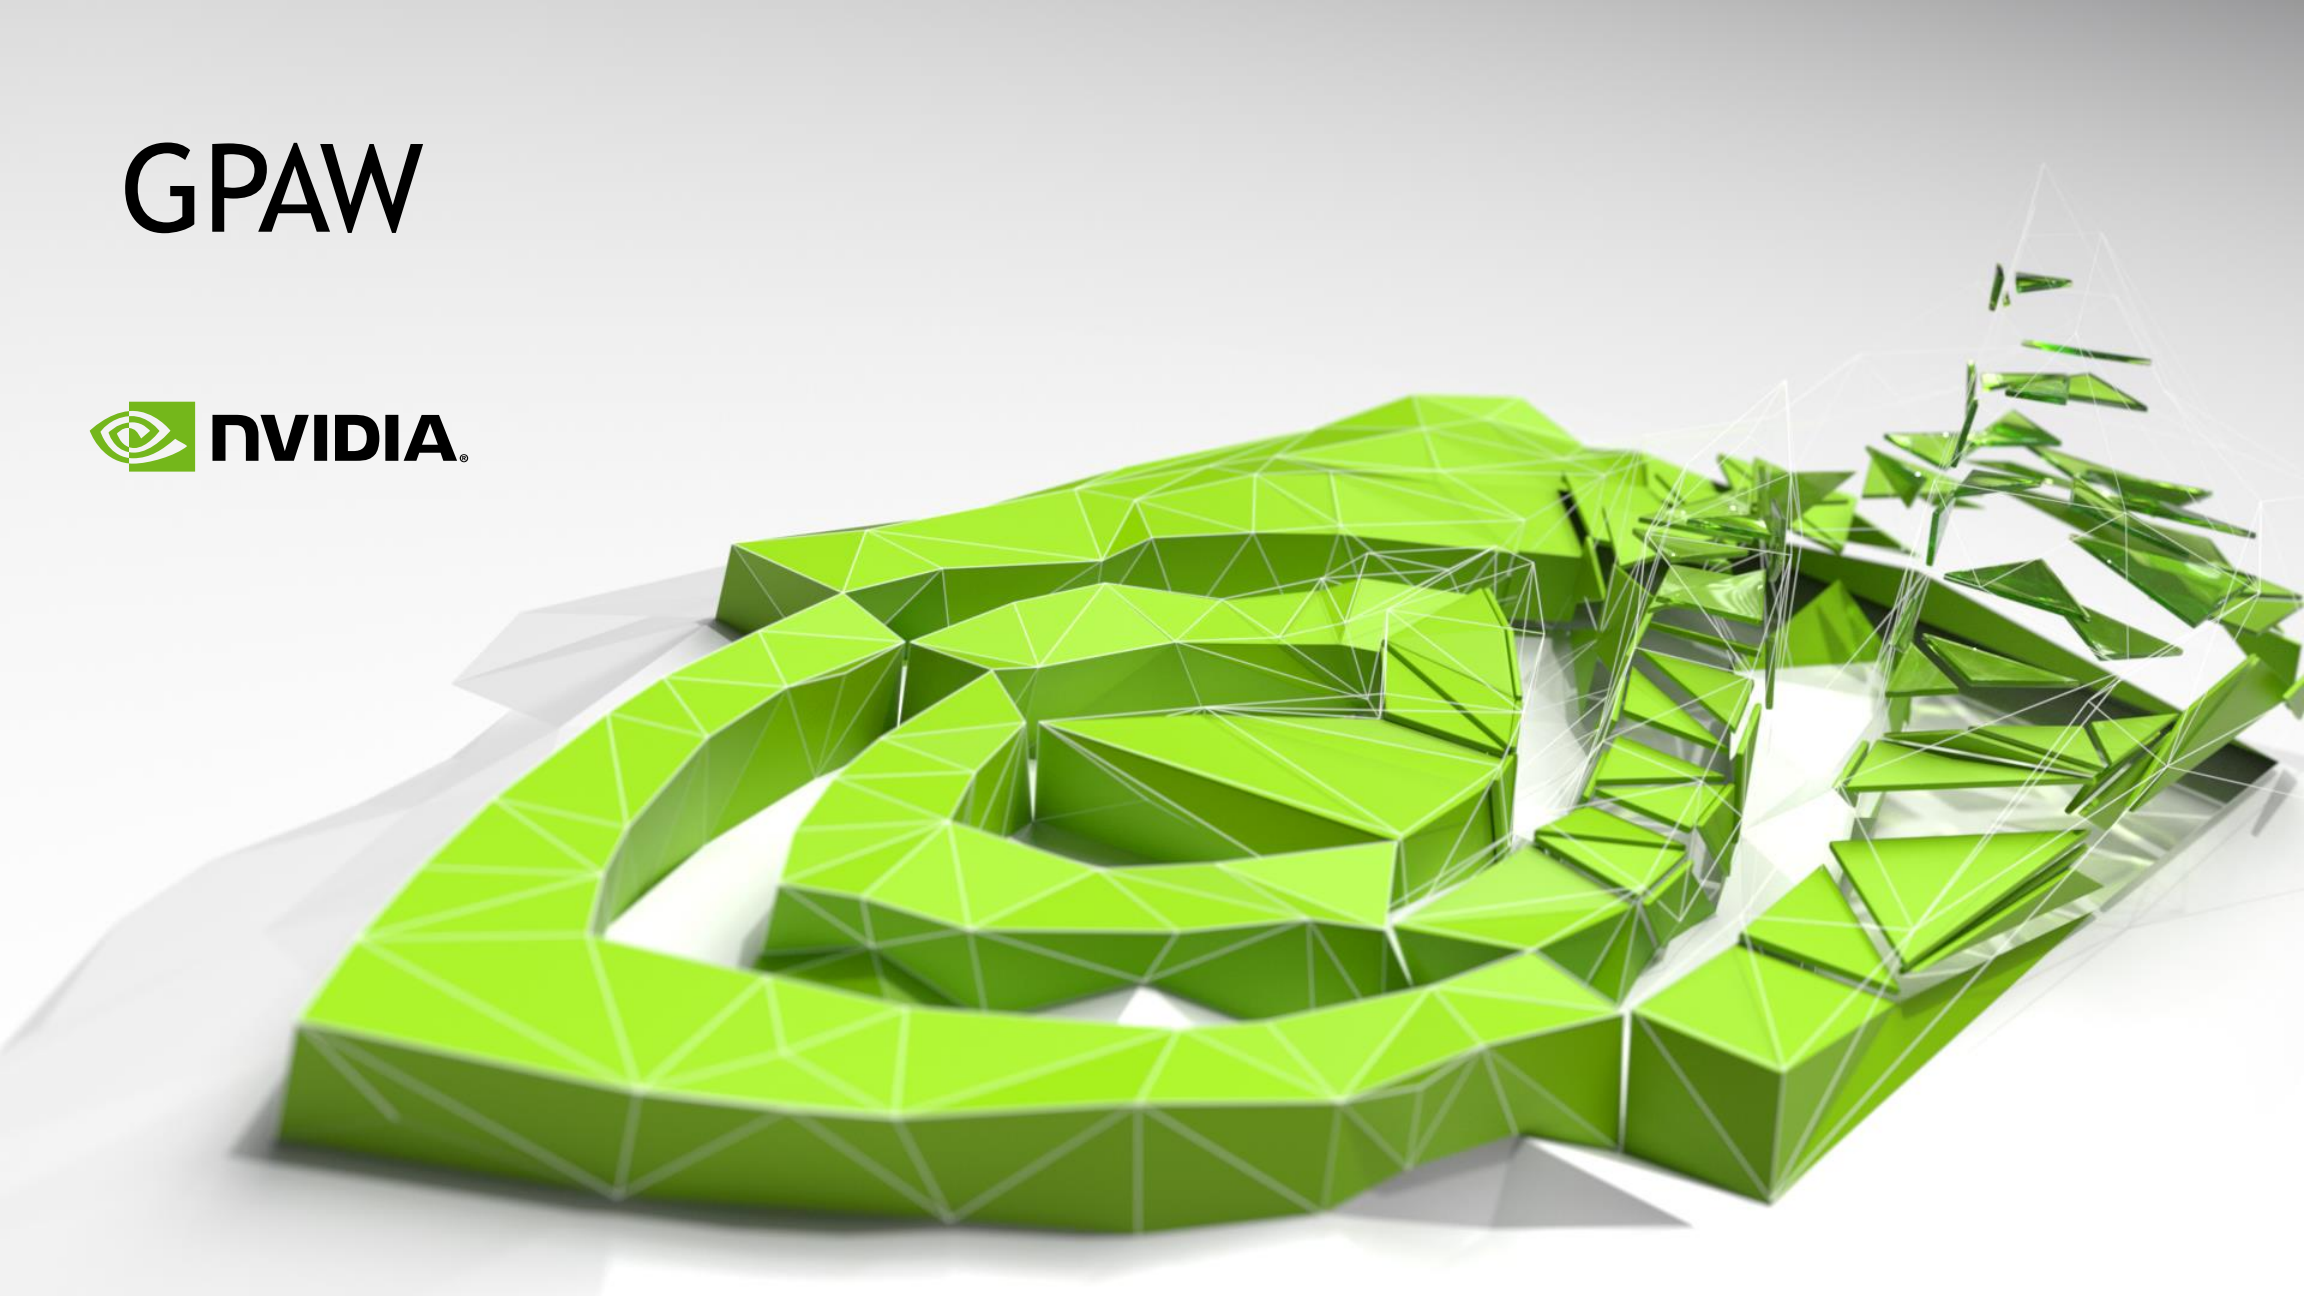

Increase Performance with Kepler

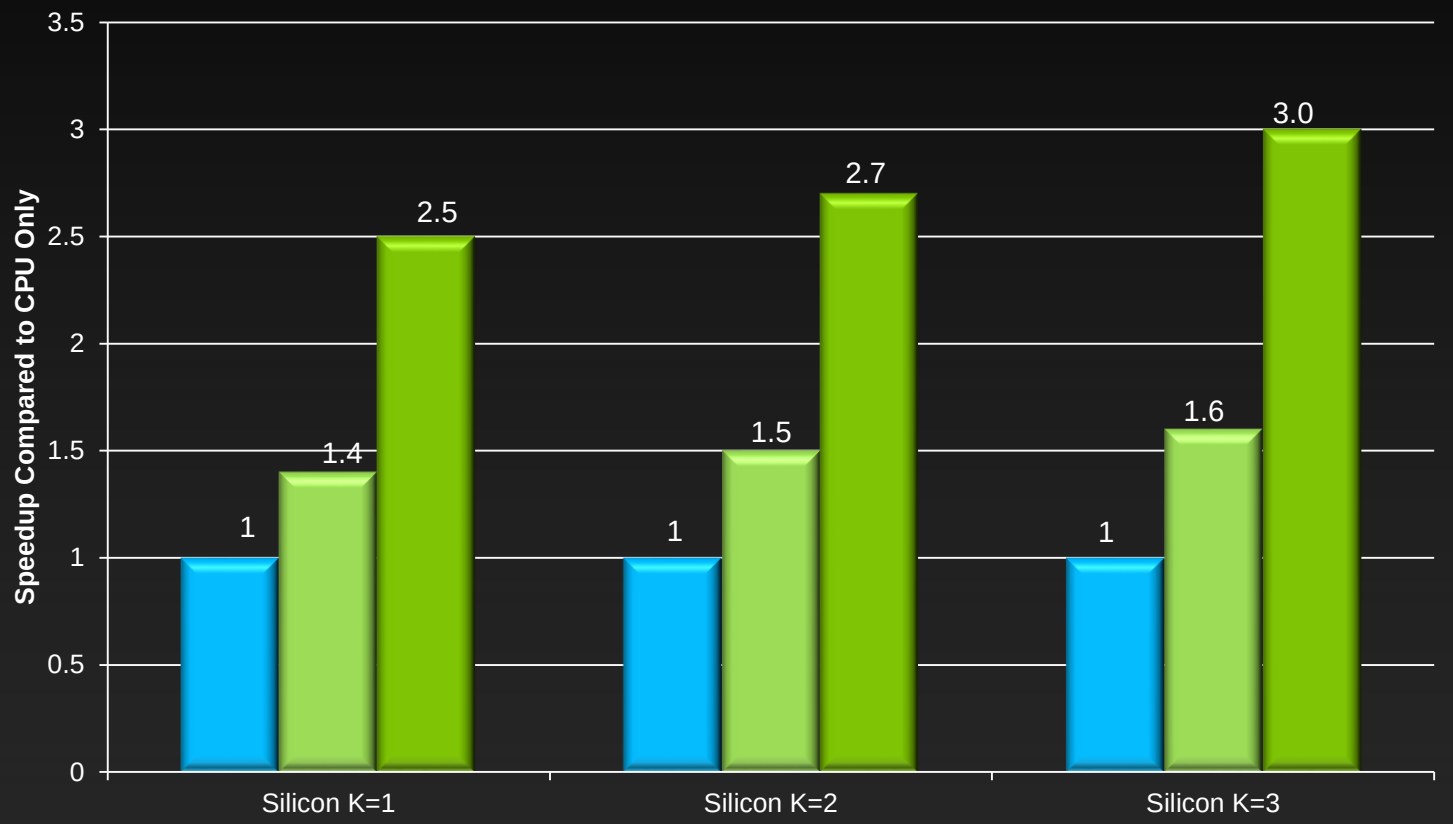

Running GPAW 10258

The blue nodes contain 1x E5-2687W CPU (8 Cores per CPU).

The green nodes contain 1x E5-2687W CPU (8 Cores per CPU) and 1x or 2x NVIDIA K20X for the GPU.

Increase Performance with Kepler

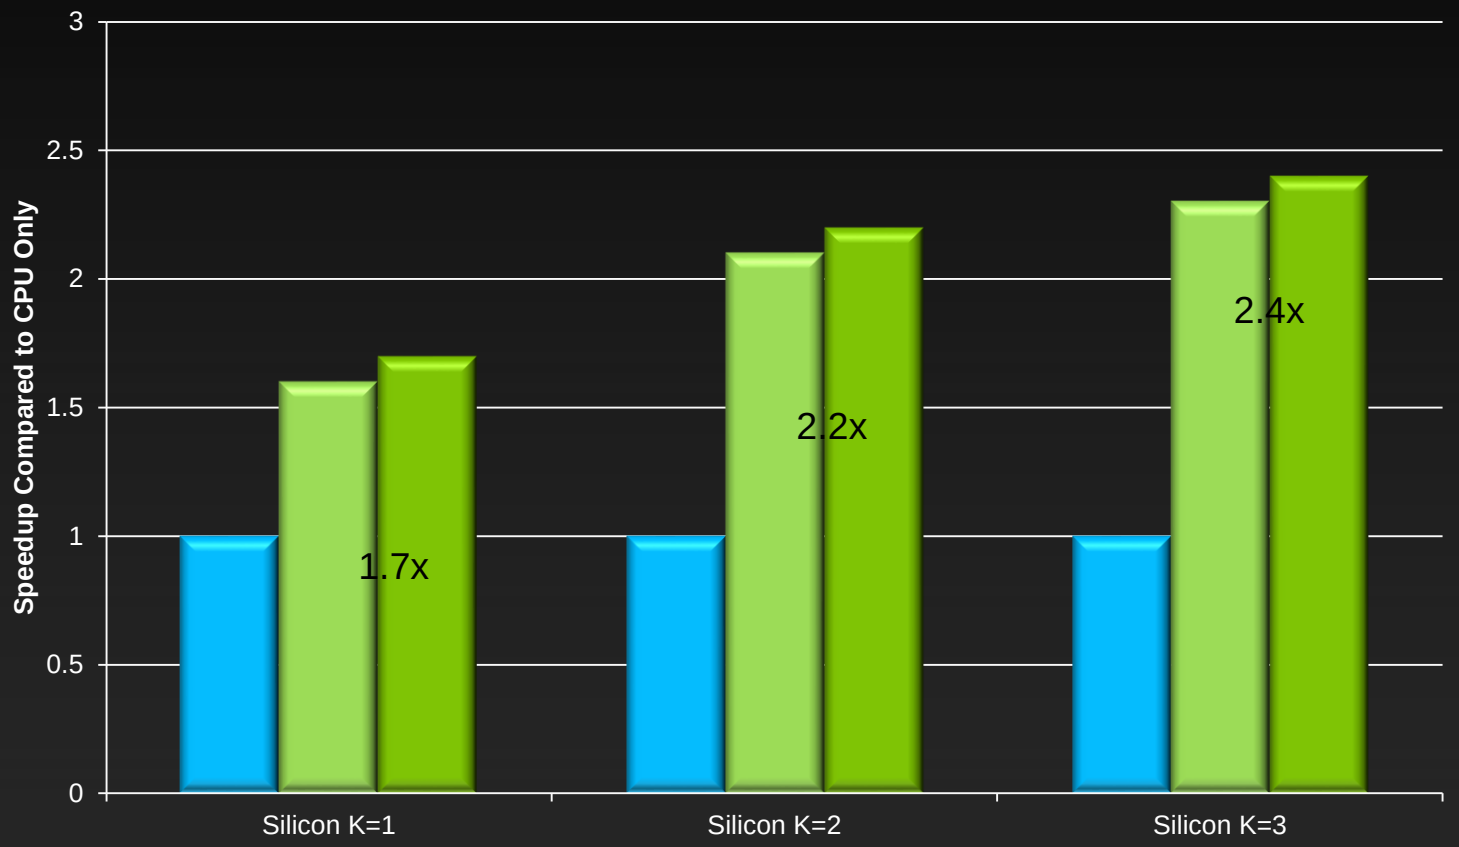

Running GPAW 10258

The blue nodes contain 1x E5-2687W CPU (8 Cores per CPU).

The green nodes contain 1x E5-2687W CPUs (8 Cores per CPU) and 2x NVIDIA K20 or K20X for the GPU.

Increase Performance with Kepler

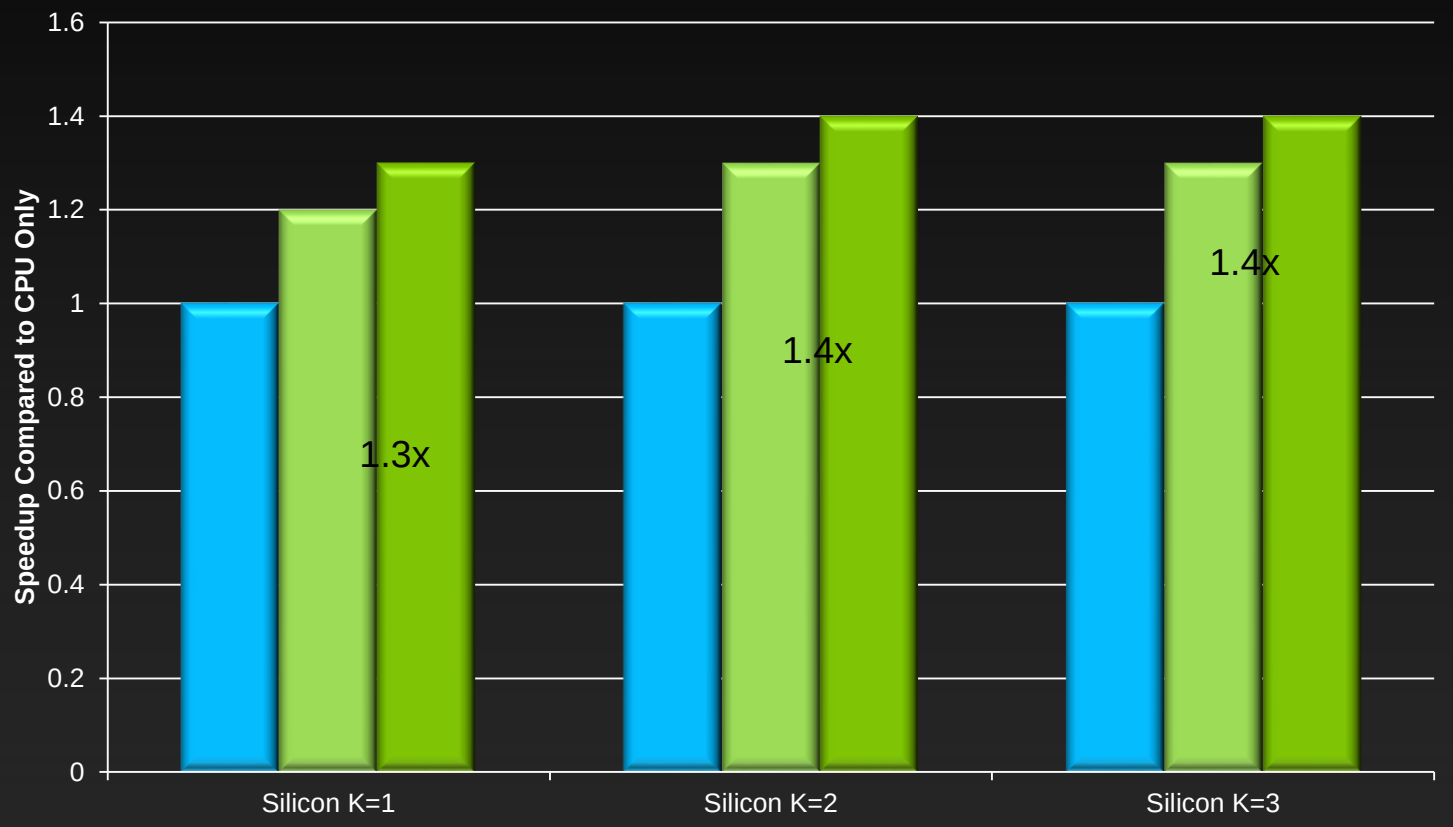

Running GPAW 10258

The blue nodes contain 2x E5-2687W CPUs (8 Cores per CPU).

The green nodes contain 2x E5-2687W CPUs (8 Cores per CPU) and 2x NVIDIA K20 or K20X for the GPU.

Multi-GPU Accelerated Large Scale Electronic Structure Calculations

Samuli Hakala
COMP Centre of Excellence
Department of Applied Physics
Aalto University School of Science
Email: samuli.hakala@aalto.fi

GPU Technology Conference, March 2013

Aalto University
School of Science

GPAW

- Density Functional Theory (DFT) program package for electronic structure calculations
- Time-Dependent Density Functional Theory (TDDFT) is implemented in the linear response and time propagation schemes
- Can use real-space grids, atom centered basis functions or plane waves
- Random Phase Approximation (RPA) also available
- Scales to thousands of cores and suitable for large scale calculations
- Open Source software licensed under GPL

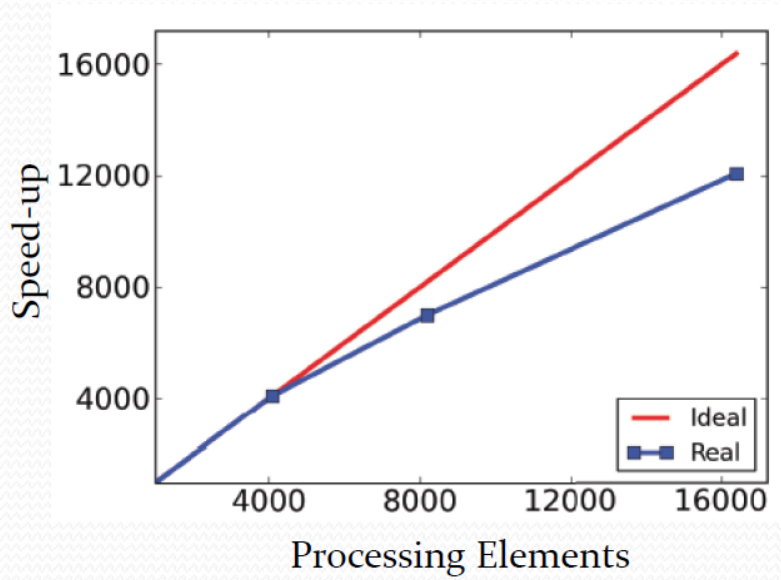

Ground state DFT calculation of 561 Au atom cluster on Blue Gene/P.

LibXC on GPUs

- A reusable library of >250 exchange-correlation functionals
- Used by 15 different codes (Abinit, GPAW, BigDFT, etc.)
- Can be a performance bottleneck for small systems
- Can "clone" existing functionals for GPU use with fairly minimal changes to existing LibXC code and parallelizes well over grid points
- More information:
  - https://confluence.slac.stanford.edu/display/SUNCAT/libxc+on+GPUs
- Work by Lin Li, Jun Yan, Christopher O'Grady (Stanford/SLAC)

Functional	Type	Speedup ((GPU+CPU)/CPU)
PW, PW Mode, OB PW, PW RPA	LDA Correlation	**23,23,23,37**
PBE, PBE sol, xPBE, PBE JRGX, RGE2, APBE	GGA Correlation	**56, 58, 58, 58, 58, 58**
RPBE	GGA Exchange	**95**
TPSS	MGGA Exchange	**51**

Ground State Performance

Bulk Silicon

- 95 atoms with periodic boundary conditions, 380 bands and 1 k-point. Grid size: 56x56x80.
- Time is in seconds per one SCF iteration.
- Intel Xeon X5650, NVIDIA Tesla M2070

Fullerene

- C60 molecule with 240 valence electrons. Grid size: 84x84x84
- Intel Xeon X5650, NVIDIA Tesla M2070

Si95	CPU	GPU	%	S-Up
Poisson Solver	1.8	0.13	1%	**14**
Orthonormalization	23	3.0	23%	**7.7**
Precondition	9.4	0.77	6%	**12**
RMM-DIIS other	32	3.2	25%	**10**
Subspace Diag	23	2.1	16%	**11**
Other	2.7	2.7	21%	**1.0**
Total (SCF-Iter)	93	13		9.7/7.7

C60	CPU	GPU	%	S-Up
	13	0.64	7%	**20**
	11	1.2	13%	**9.2**
	16	0.99	11%	**16**
	8.1	0.6	7%	**13**
	22	2.1	23%	**10**
	3.5	3.2	35%	**1.1**
	76	9.1		13/8.3

Multi-GPU Parallelization

- Parallelization is done with MPI
- Multiple GPUs can be used by domain decomposition or parallelization over k-points or spins
- Domain decomposition for the stencil operations involves exchanging boundary regions between neighboring nodes
- Communications between nodes require data movement: device memory → host memory → destinations node host memory → destinations node device memory.
- Overlaps receives, sends and computations in the middle part of the grid, BUT this causes issues with small grids
  - Small grids: Synchronous transfers
  - Medium grids: Asynchronous transfers
  - Large grids: Overlap calculations and asynchronous transfers
  - Combine of several wave functions and boundary regions into few large transfers

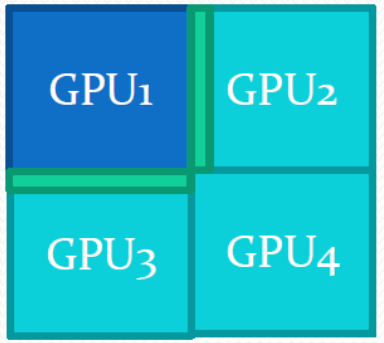

Weak Scalability (Carbon)

- The size of a carbon nanotube and the number of MPI tasks are varied from 80 atoms (240 states) to 320 atoms (1280 states) and 1 task to 12 tasks.
- Comparison between equal number of GPUs and CPU cores.
- CPU: Intel Xeon X5650 GPU: NVIDIA Tesla M2070
- Calculations performed on Vuori cluster at CSC

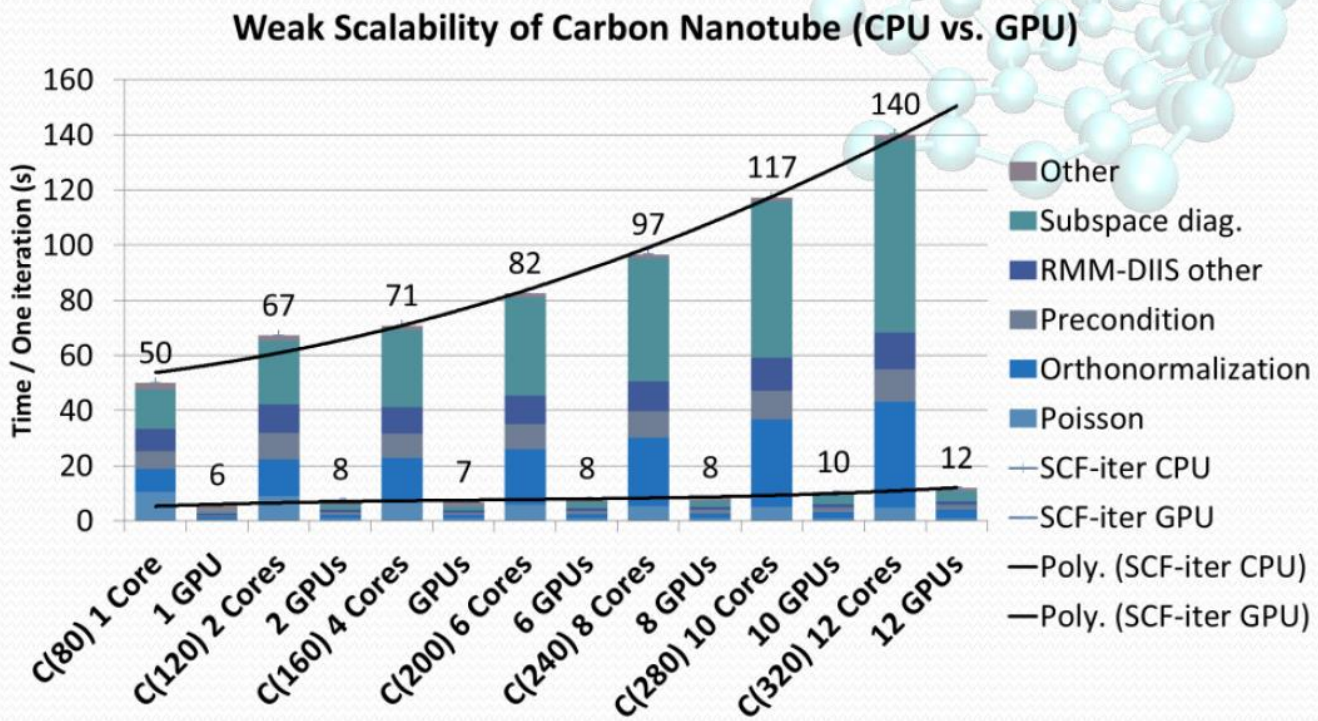

# MPI tasks	1	2	4	6	8	10	12
Speed-Up	8.8	8.7	10.5	10.2	11.5	11.3	11.9

Strong Scalability

- Bulk silicon with 1151 atoms with periodic boundary conditions, 4604 bands and 1 k-point in the Brillouin zone.
- The number of GPUs is increased from 64 to 256.
- Grid size: 164x164x108
- Speed-up comparison to 64 GPUs.
- NVIDIA Tesla M2090
- Calculations performed on CURIE cluster in France at GENCI/CEA

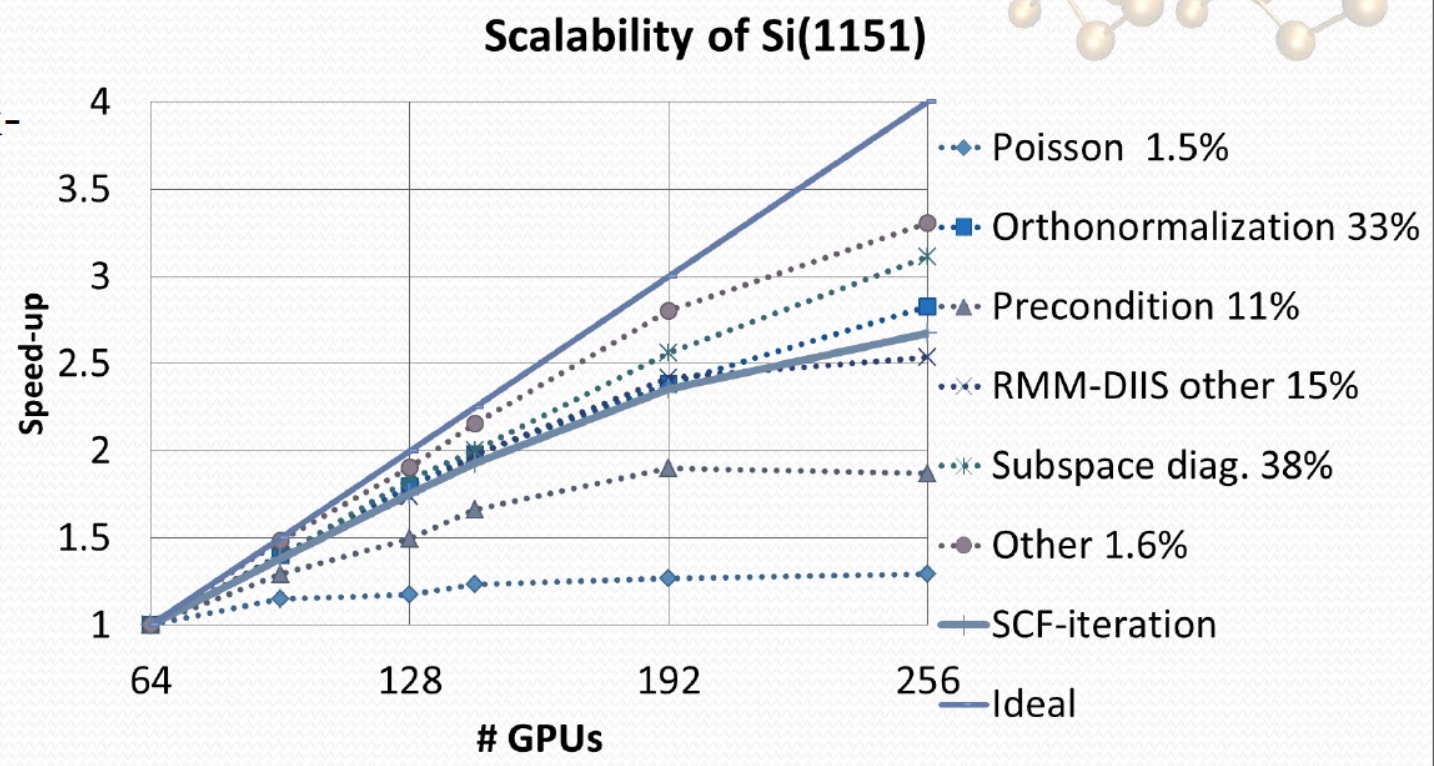

Weak Scalability (Silicon)

- The size of bulk silicon system and the number of MPI tasks are varied from 383 atoms (1532 bands) to 2046 atoms (8188 bands) and 8 task to 256 tasks with periodic boundary conditions.
- The largest system requires about 1.3TB of memory for calculations.
- CPU: Intel Xeon E5640 GPU: NVIDIA Tesla M2090

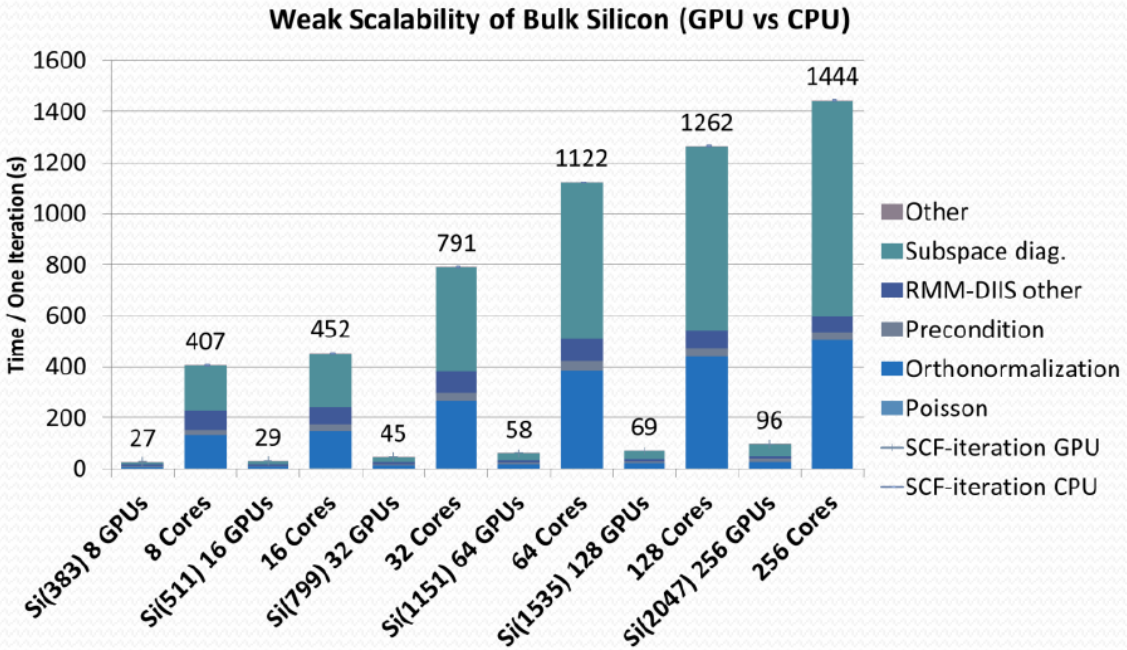

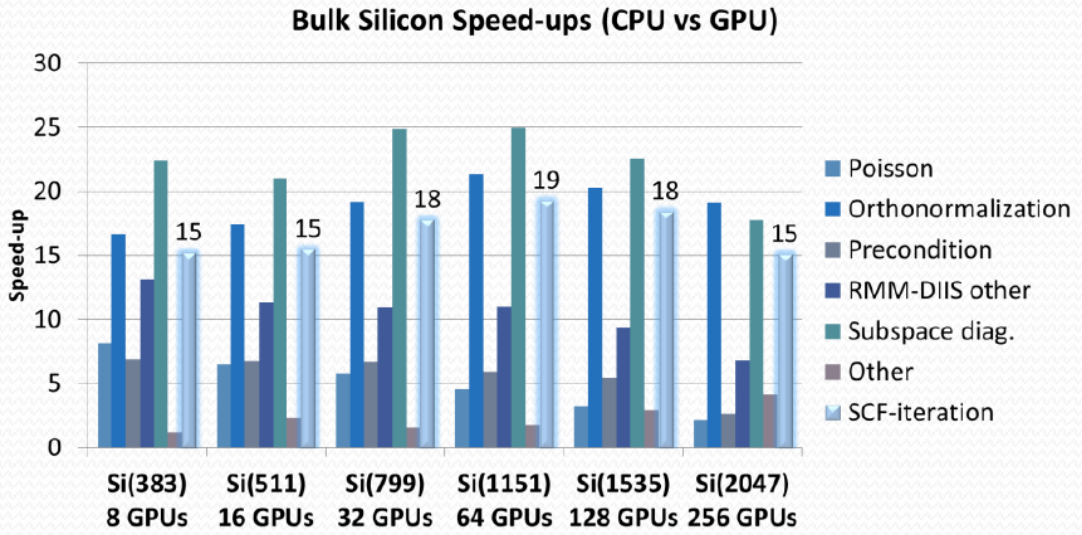

Random Phase Approximation

GPAW Random Phase Approximation (RPA) code:

- 6000 lines of python, 1000 lines of C/CUDA (and re-uses many GPAW functions)
- Better than DFT for correlated materials, but more computationally expensive
- Useful for oxides, Van der Waals systems, etc.

GPU Techniques:

- Use BLAS3 "zherk" instead of BLAS2 "zher"
- Batch FFTs
- GPU kernels parallelized over atoms/bands/projector-functions
- No thunking: all calculations on GPU

Preliminary ((GPU+CPU)/CPU) speedup for 202-electron N_2-on-Ru: 30x

Work by Jun Yan, Lin Li, Christopher O'Grady (Stanford/SLAC)

Summary

- We have accelerated the most numerically intensive parts of ground state DFT calculations
- Overall speed-ups in our tests varied from 8.8 to 19 depending on system size.
- Our multi-GPU implementation scales well even on large hybrid clusters.
- Code is available at GPAW Subversion repository.
- Acknowledgements to CSC and PRACE for computing resources

Hakala S., Havu V., Enkovaara J., Nieminen R. M. "Parallel Electronic Structure Calculations Using Multiple Graphics Processing Units (GPUs)" In: Manninen, P., Öster, P. (eds.) PARA 2012. LNCS, vol. 7782, pp. 63--76. Springer, Heidelberg (2013)

LSDALTON

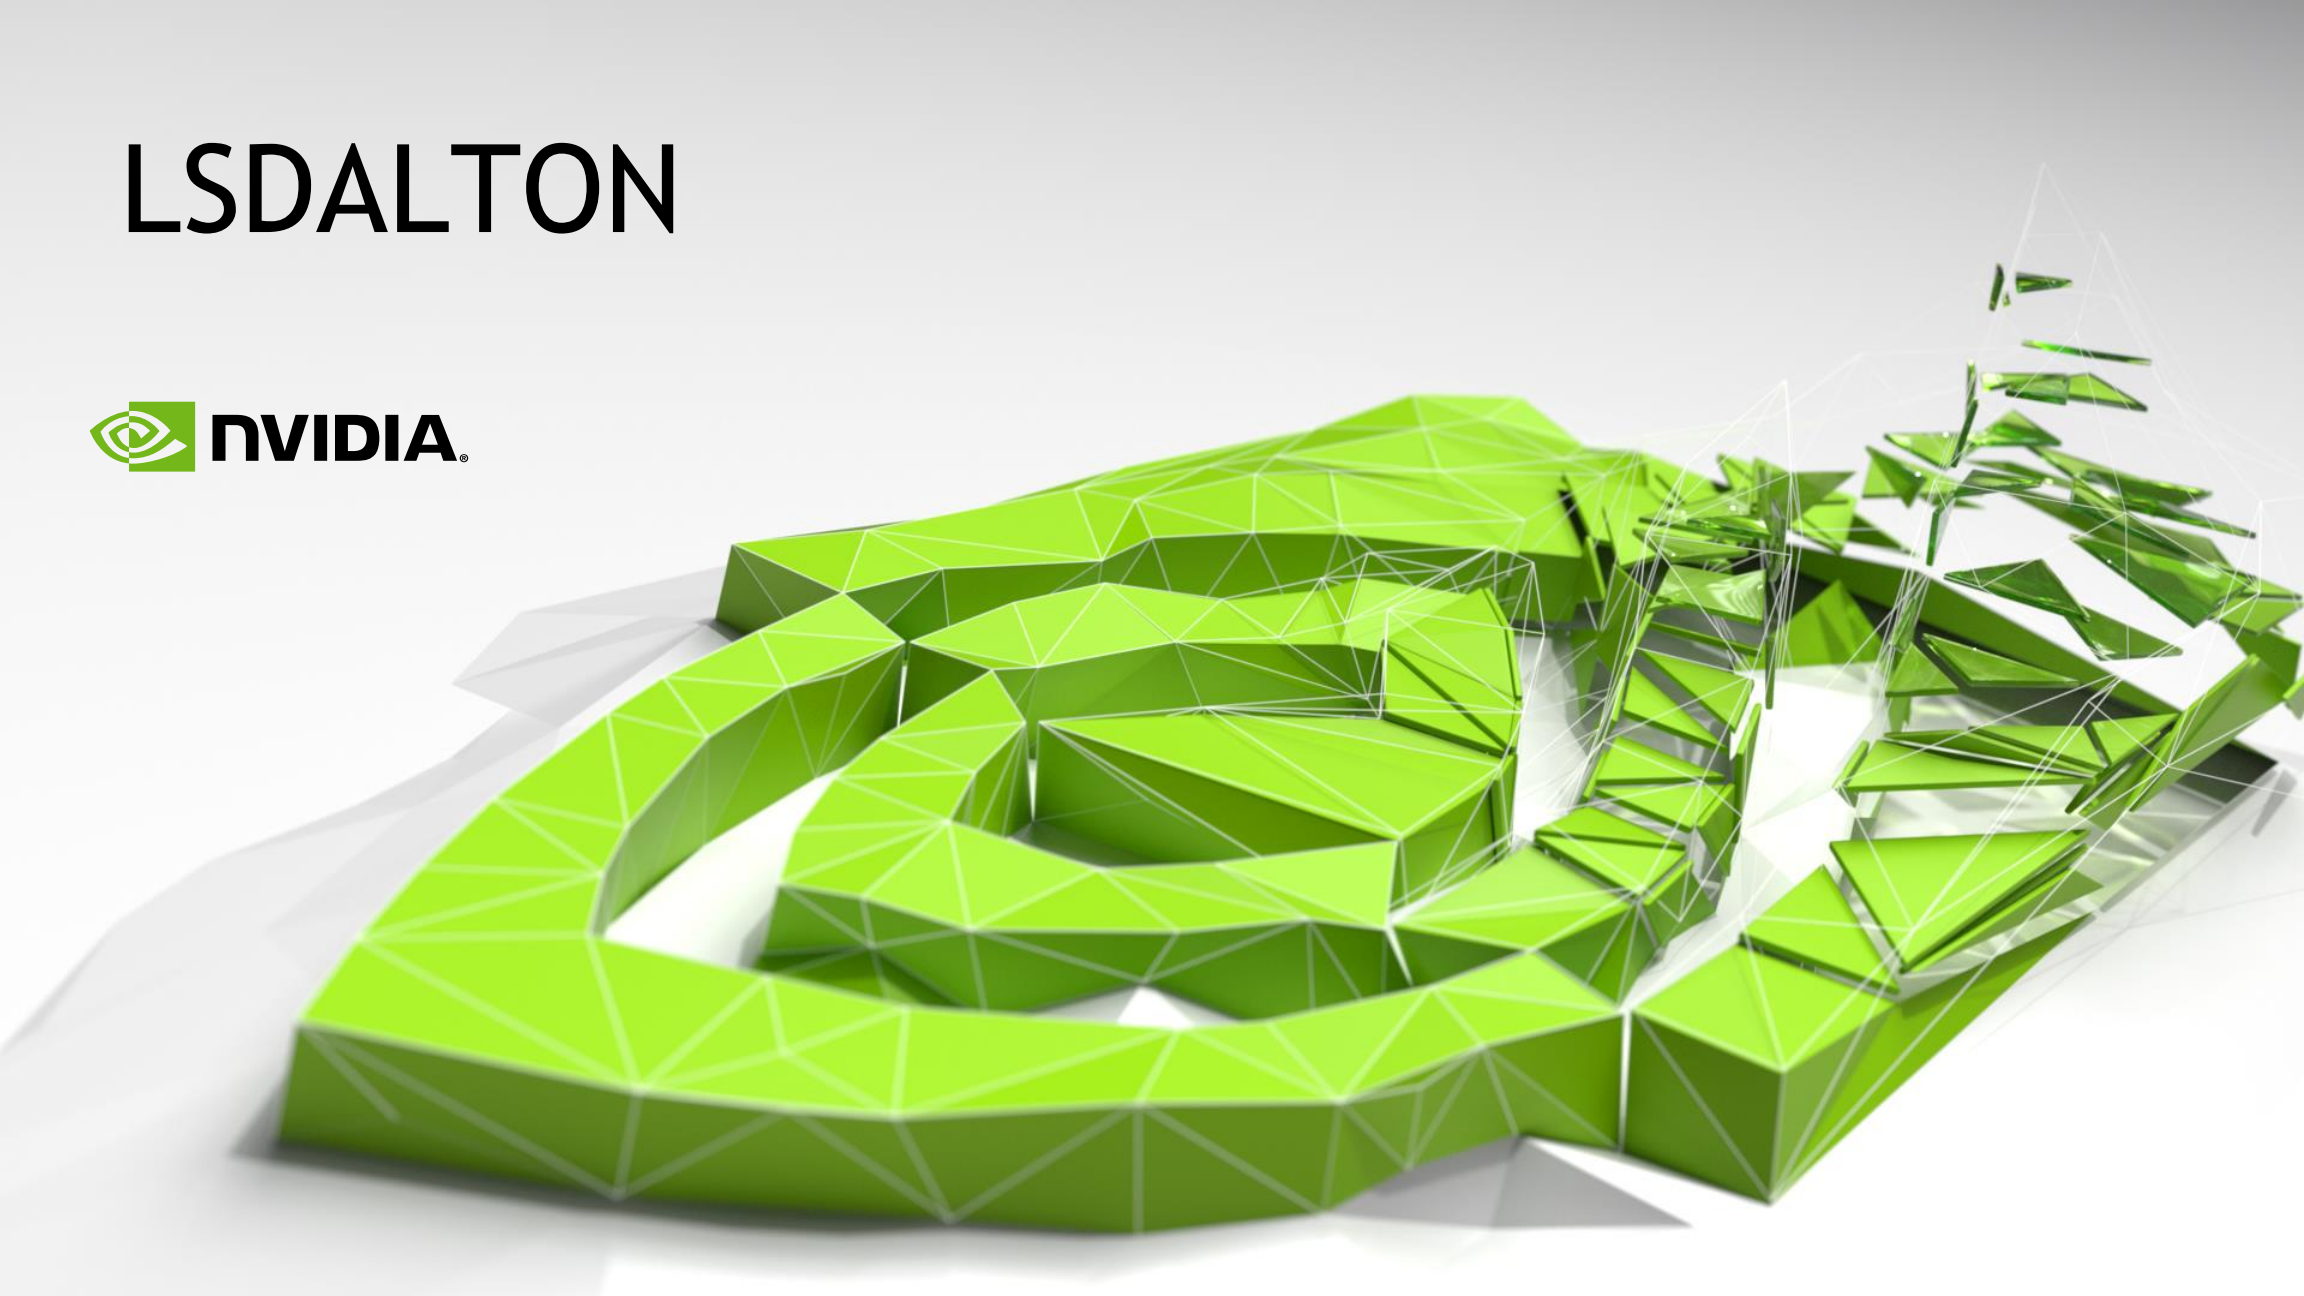

LSDALTON

Large-scale application for calculating high-accuracy molecular energies

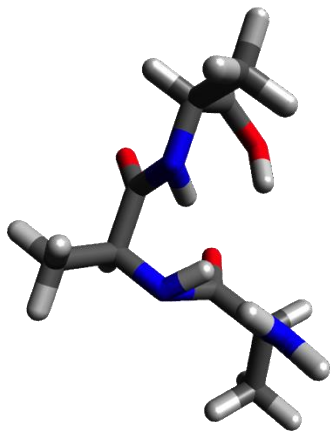

"*OpenACC makes GPU computing approachable for domain scientists. Initial OpenACC implementation required only minor effort, and more importantly, no modifications of our existing CPU implementation.*"

Janus Juul Eriksen, PhD Fellow
qLEAP Center for Theoretical Chemistry, Aarhus University

Minimal Effort

Lines of Code Modified	# of Weeks Required	# of Codes to Maintain
<100 Lines	1 Week	1 Source

Big Performance

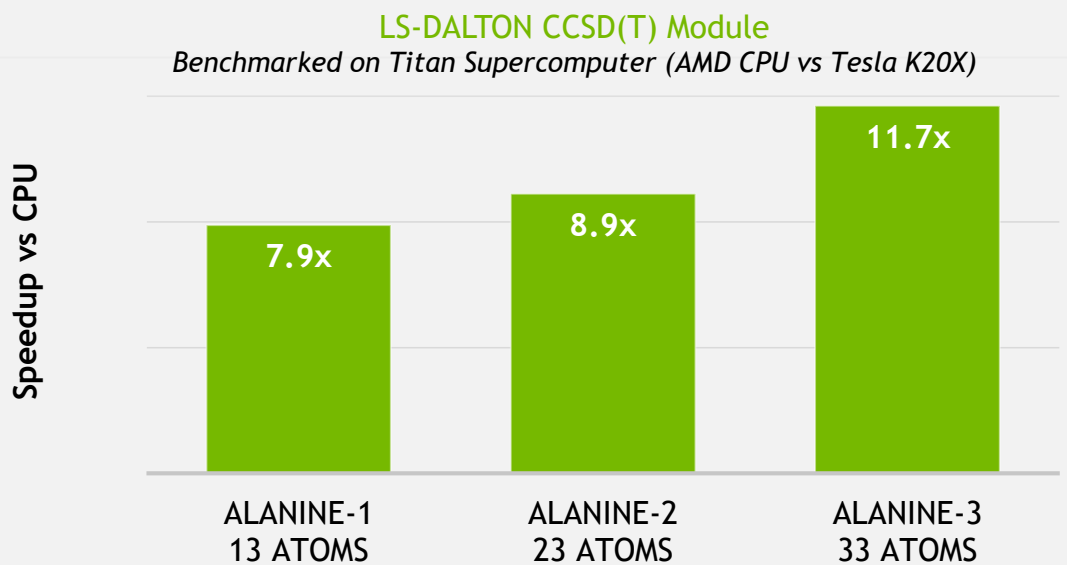

https://developer.nvidia.com/openacc/success-stories

NWChem

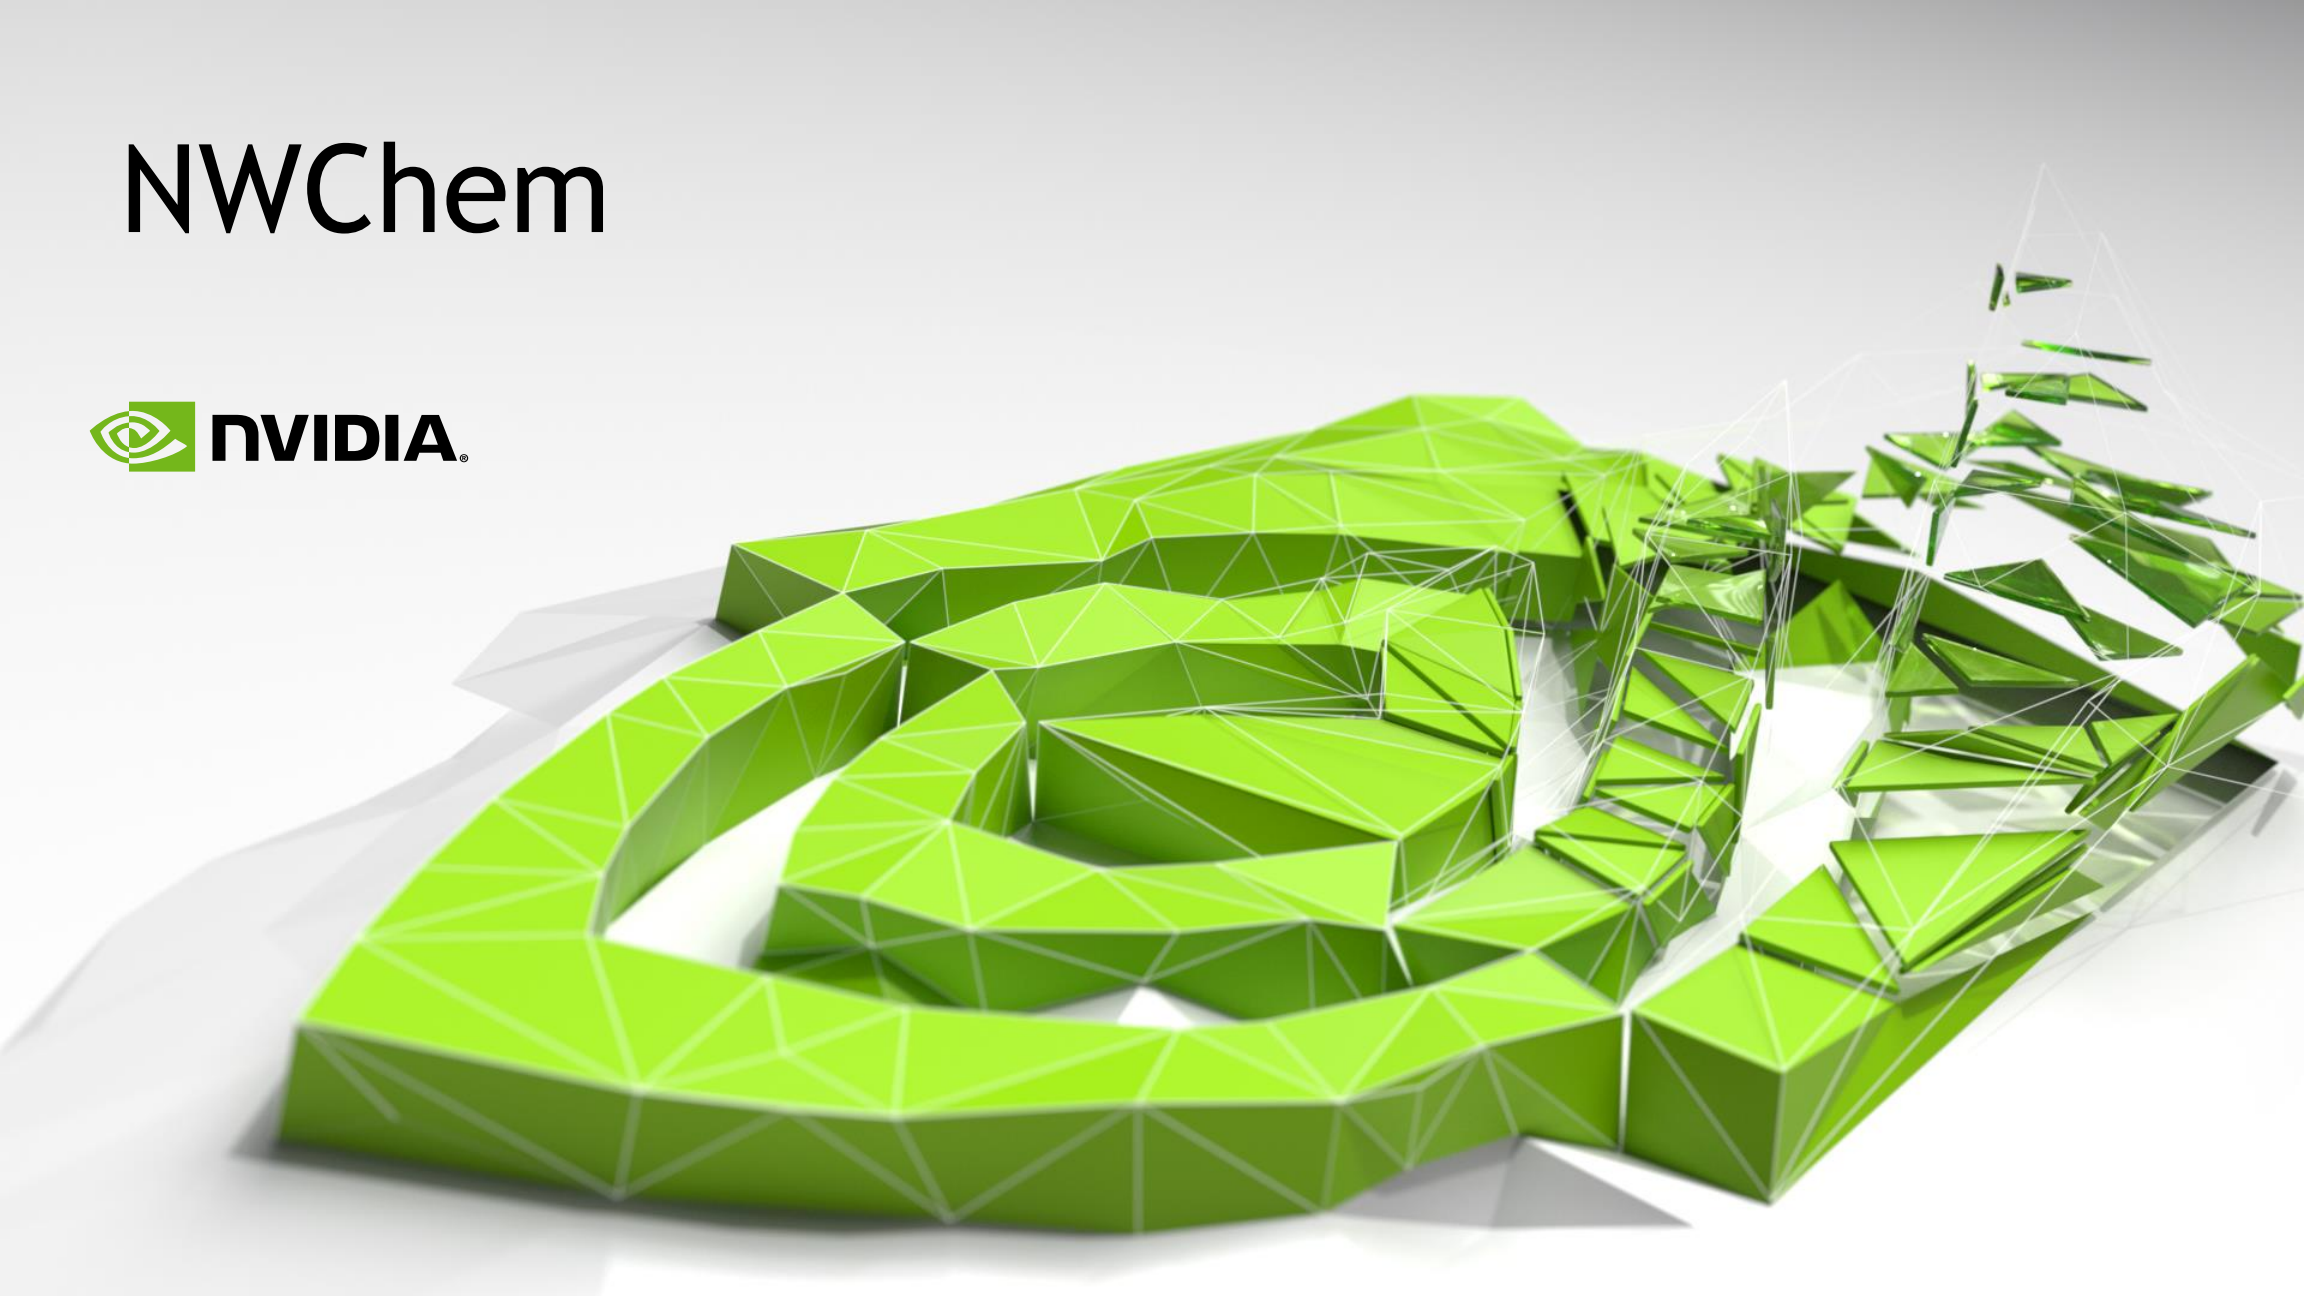

NWChem 6.3 Release with GPU Acceleration

- Addresses large complex and challenging molecular-scale scientific problems in the areas of catalysis, materials, geochemistry and biochemistry on highly scalable, parallel computing platforms to obtain the fastest time-to-solution
- Researchers can for the first time be able to perform large scale coupled cluster with perturbative triples calculations utilizing the NVIDIA GPU technology. A highly scalable multi-reference coupled cluster capability will also be available in NWChem 6.3.
- The software, released under the Educational Community License 2.0, can be downloaded from the NWChem website at www.nwchem-sw.org

NWChem - Speedup of the non-iterative calculation for various configurations/tile sizes

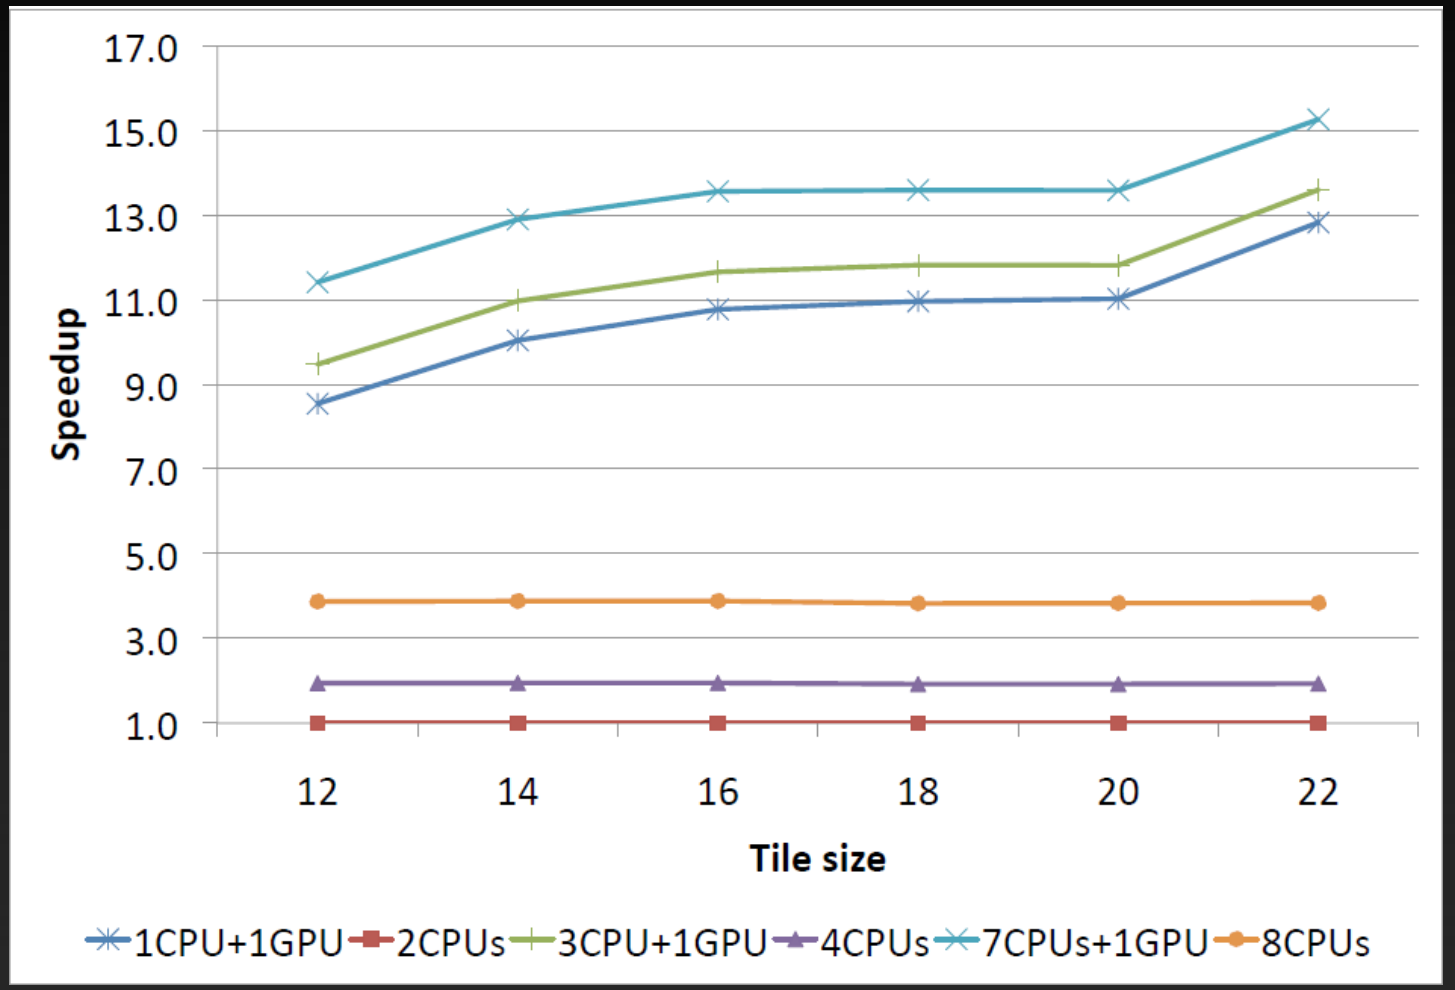

System: cluster consisting of dual-socket nodes constructed from:

- 8-core AMD Interlagos processors
- 64 GB of memory
- Tesla M2090 (Fermi) GPUs

The nodes are connected using a high-performance QDR Infiniband interconnect

Courtesy of Kowolski, K., Bhaskaran-Nair, at al @ PNNL, JCTC (submitted)

Kepler, Faster Performance (NWChem)

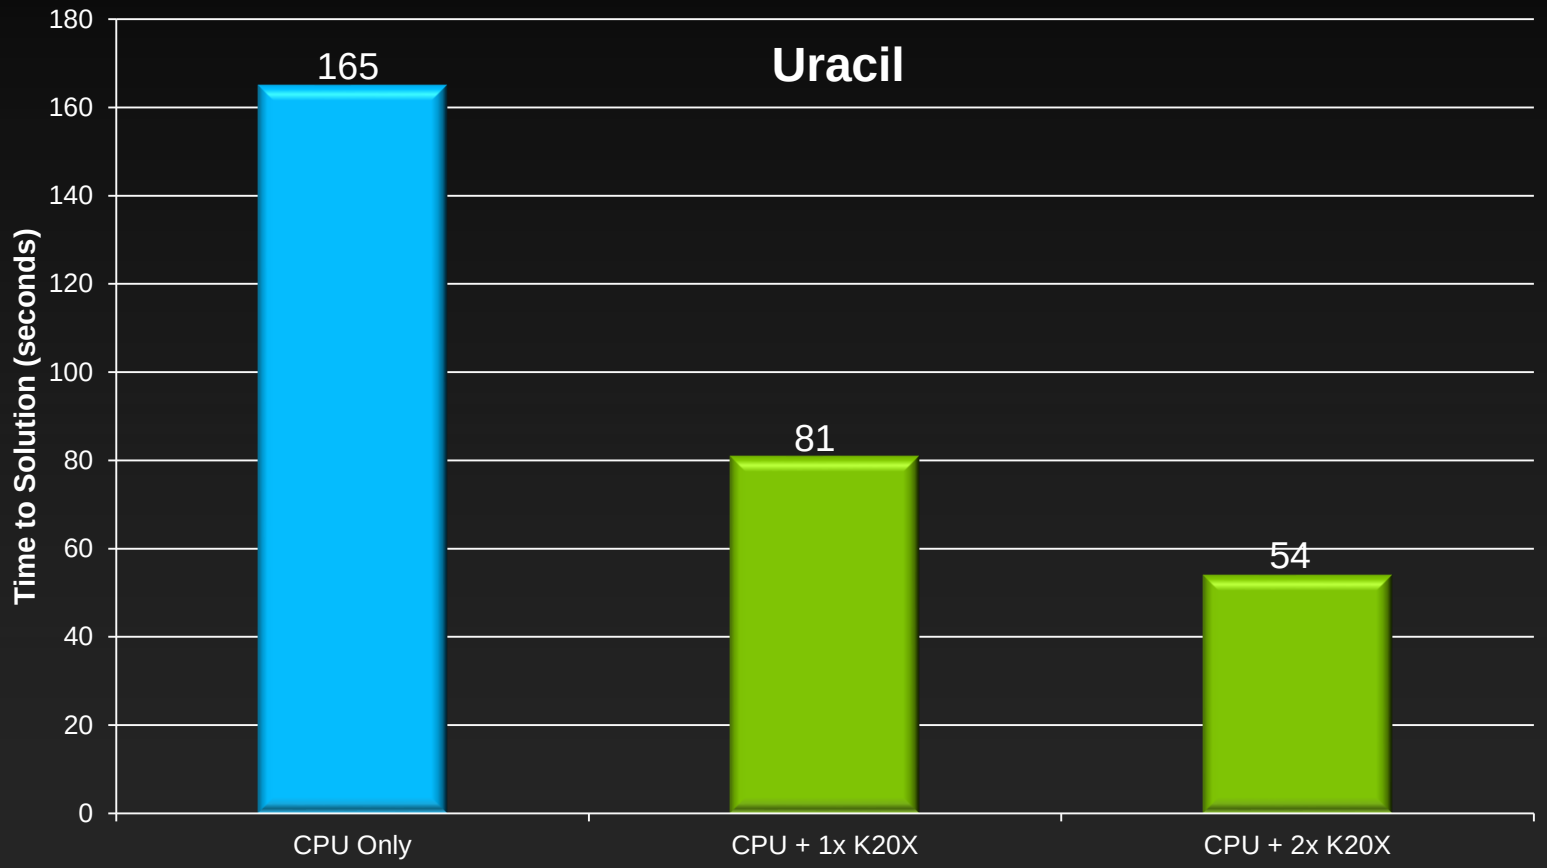

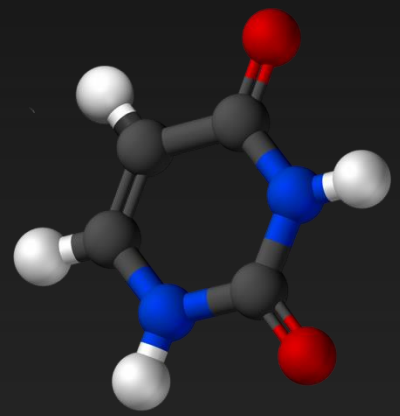

Uracil Molecule

Performance improves by 2x with one GPU and by 3.1x with 2 GPUs

Quantum Espresso 5.4.0

March 2018

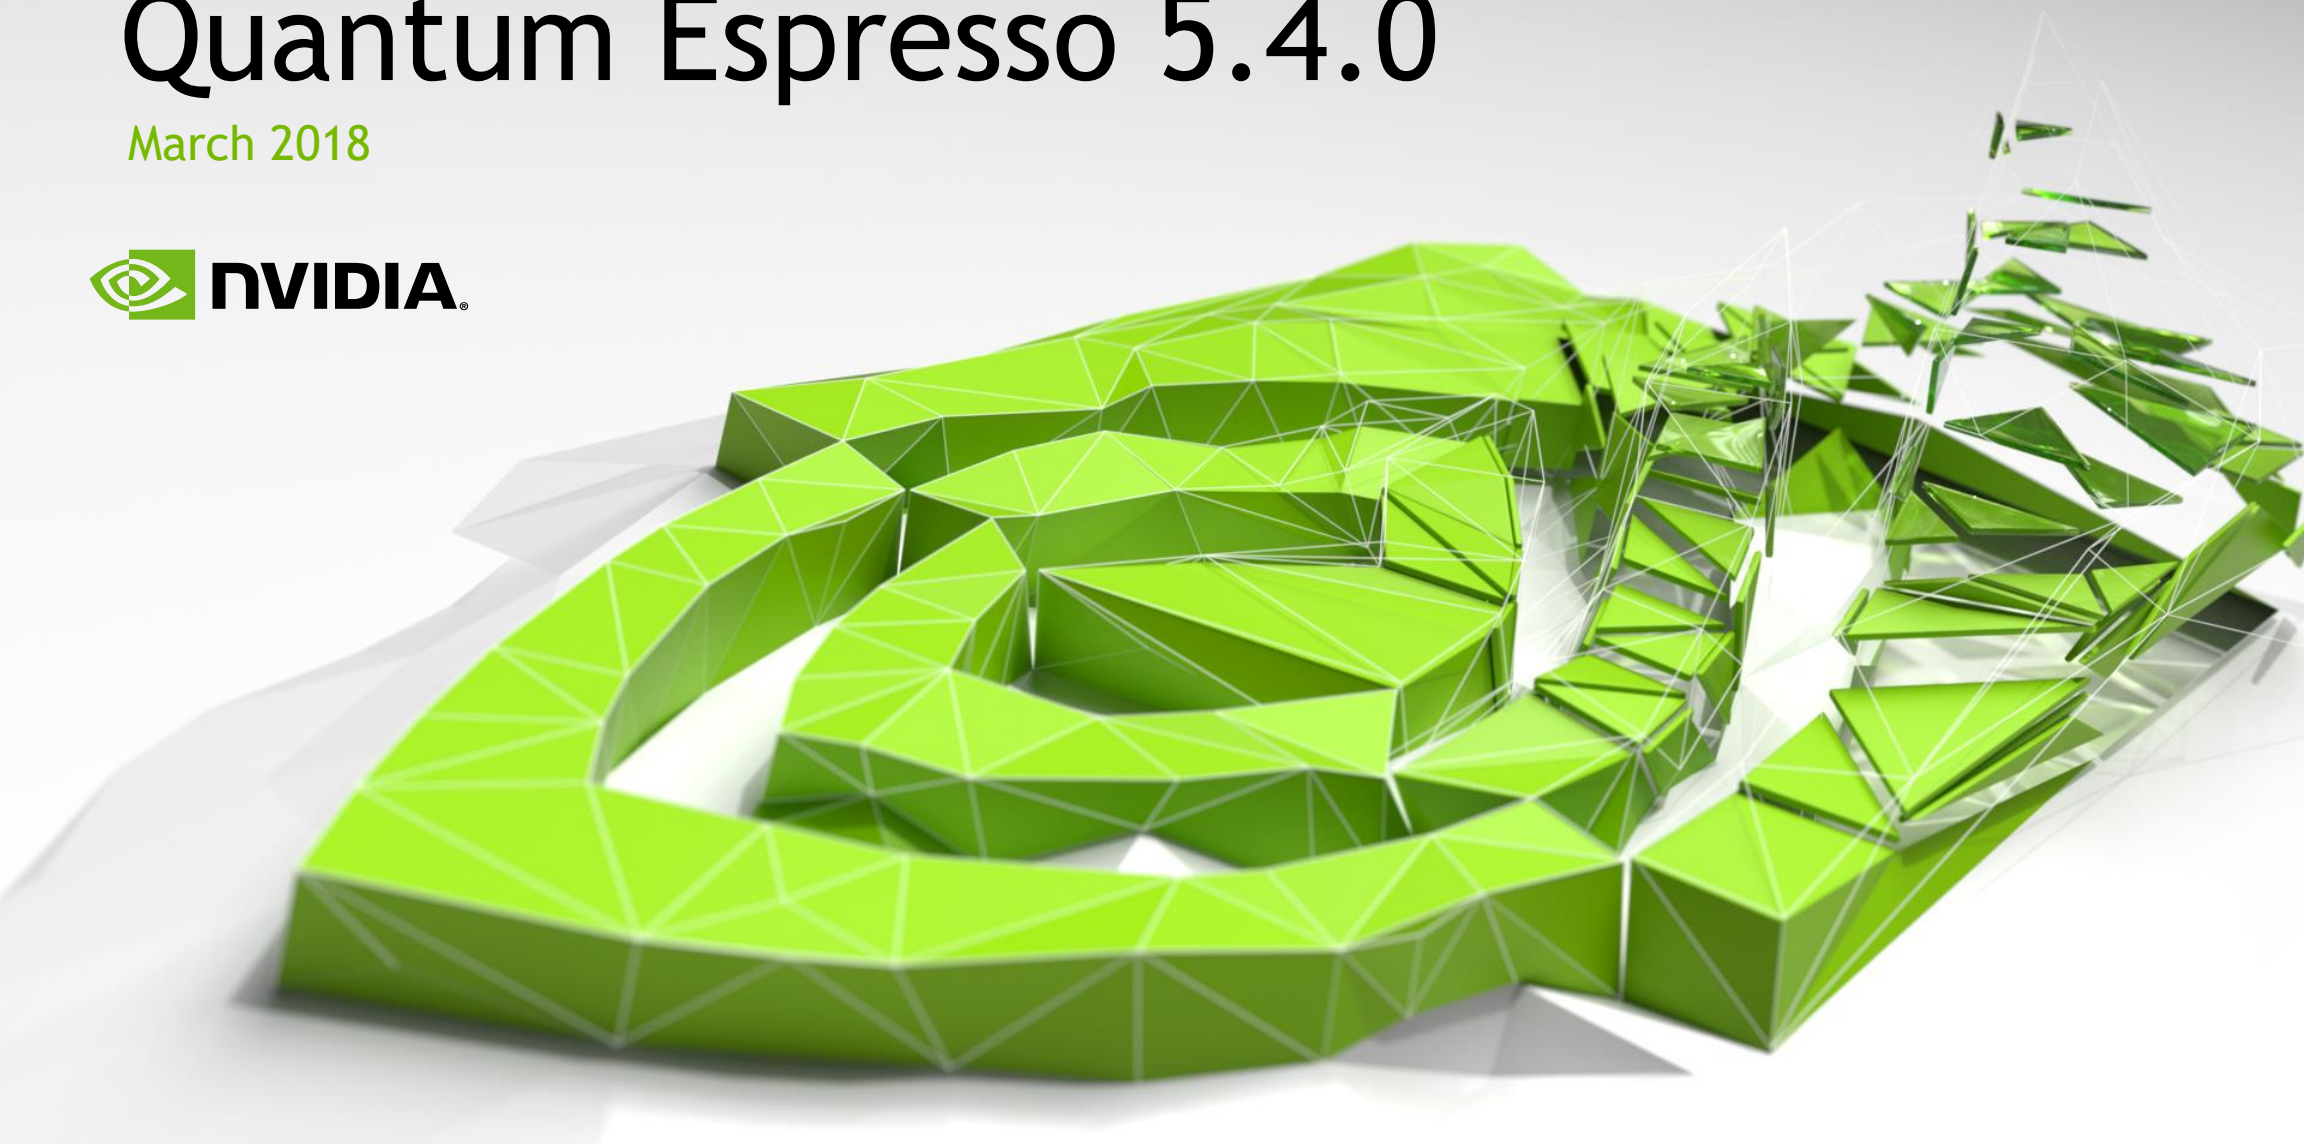

QUANTUM ESPRESSO

Quantum Chemistry Suite

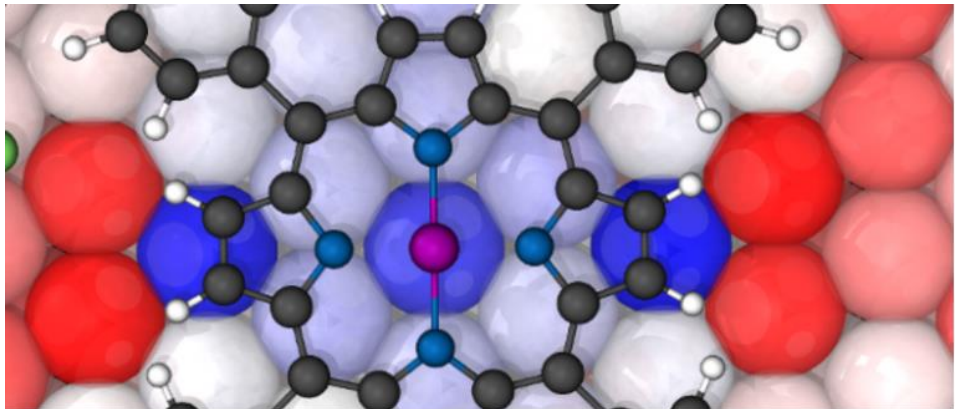

QUANTUMESPRESSO

www.quantum-espresso.org

Filippo Spiga
Head of Research Software Engineering
University of Cambridge

> CUDA Fortran gives us the full performance potential of the CUDA programming model and NVIDIA GPUs. !$CUF KERNELS directives give us productivity and source code maintainability. It's the best of both worlds.

PGI

QUANTUM ESPRESSO PERFORMANCE

Benchmark Si63GE-vc-relax

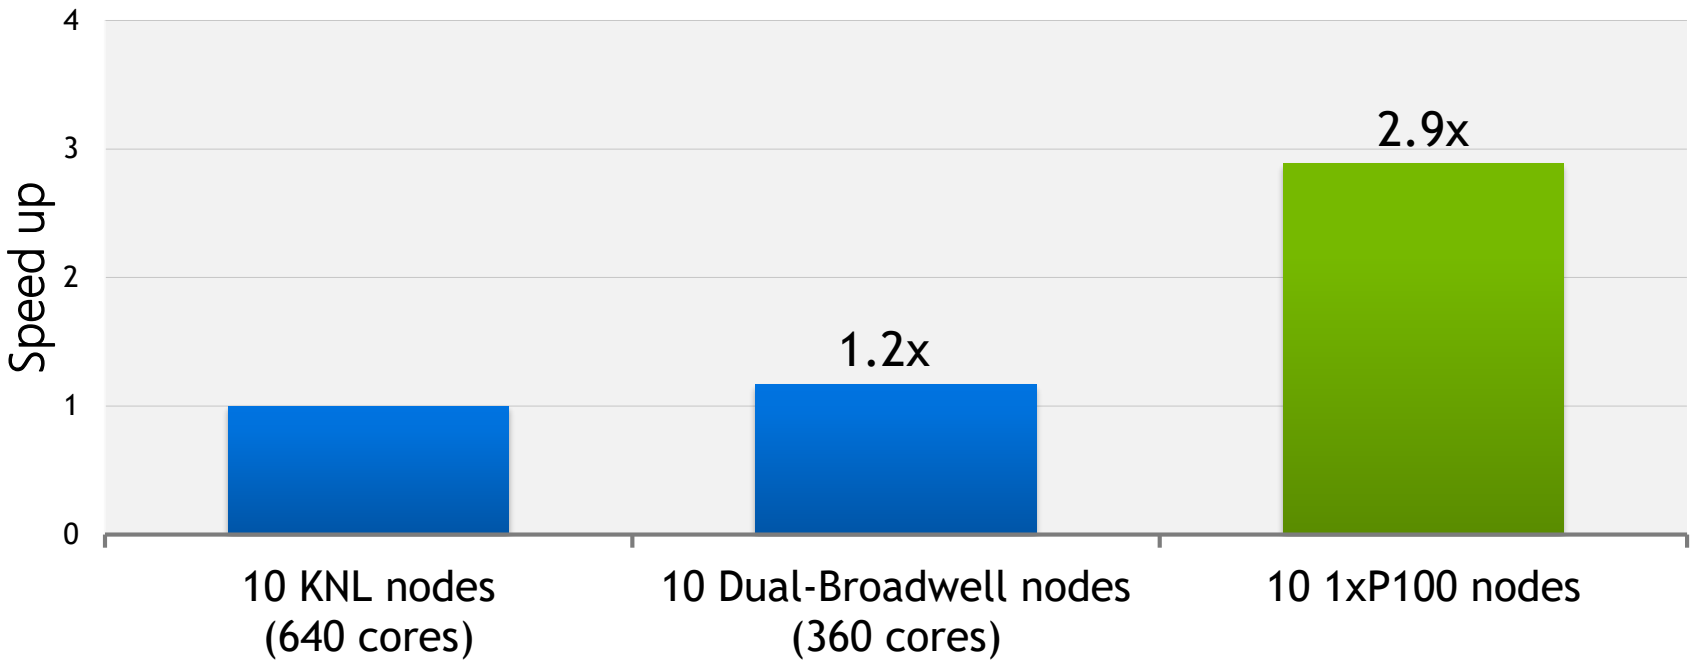

System information:
KNL: Cineca System A2 nodes (Xeon Phi 7250 @ 1.4GHz, 68 cores, 16GB MCDRAM and 96GB DDR4 per node)
Broadwell: CSCS Piz Dora Cray XC50 nodes (2 Xeon E5-2695 V4 @ 2.1GHz, 36 cores and 128GB DDR4 per node.
P100: CSCS Piz Dora Cray XC50 nodes (1 Xeon E5-2695 V4 @ 2.1GHz, 18 cores, 64/128GB DD4 and one Tesla P100 per node)

TeraChem 1.5K

Speaker, Date

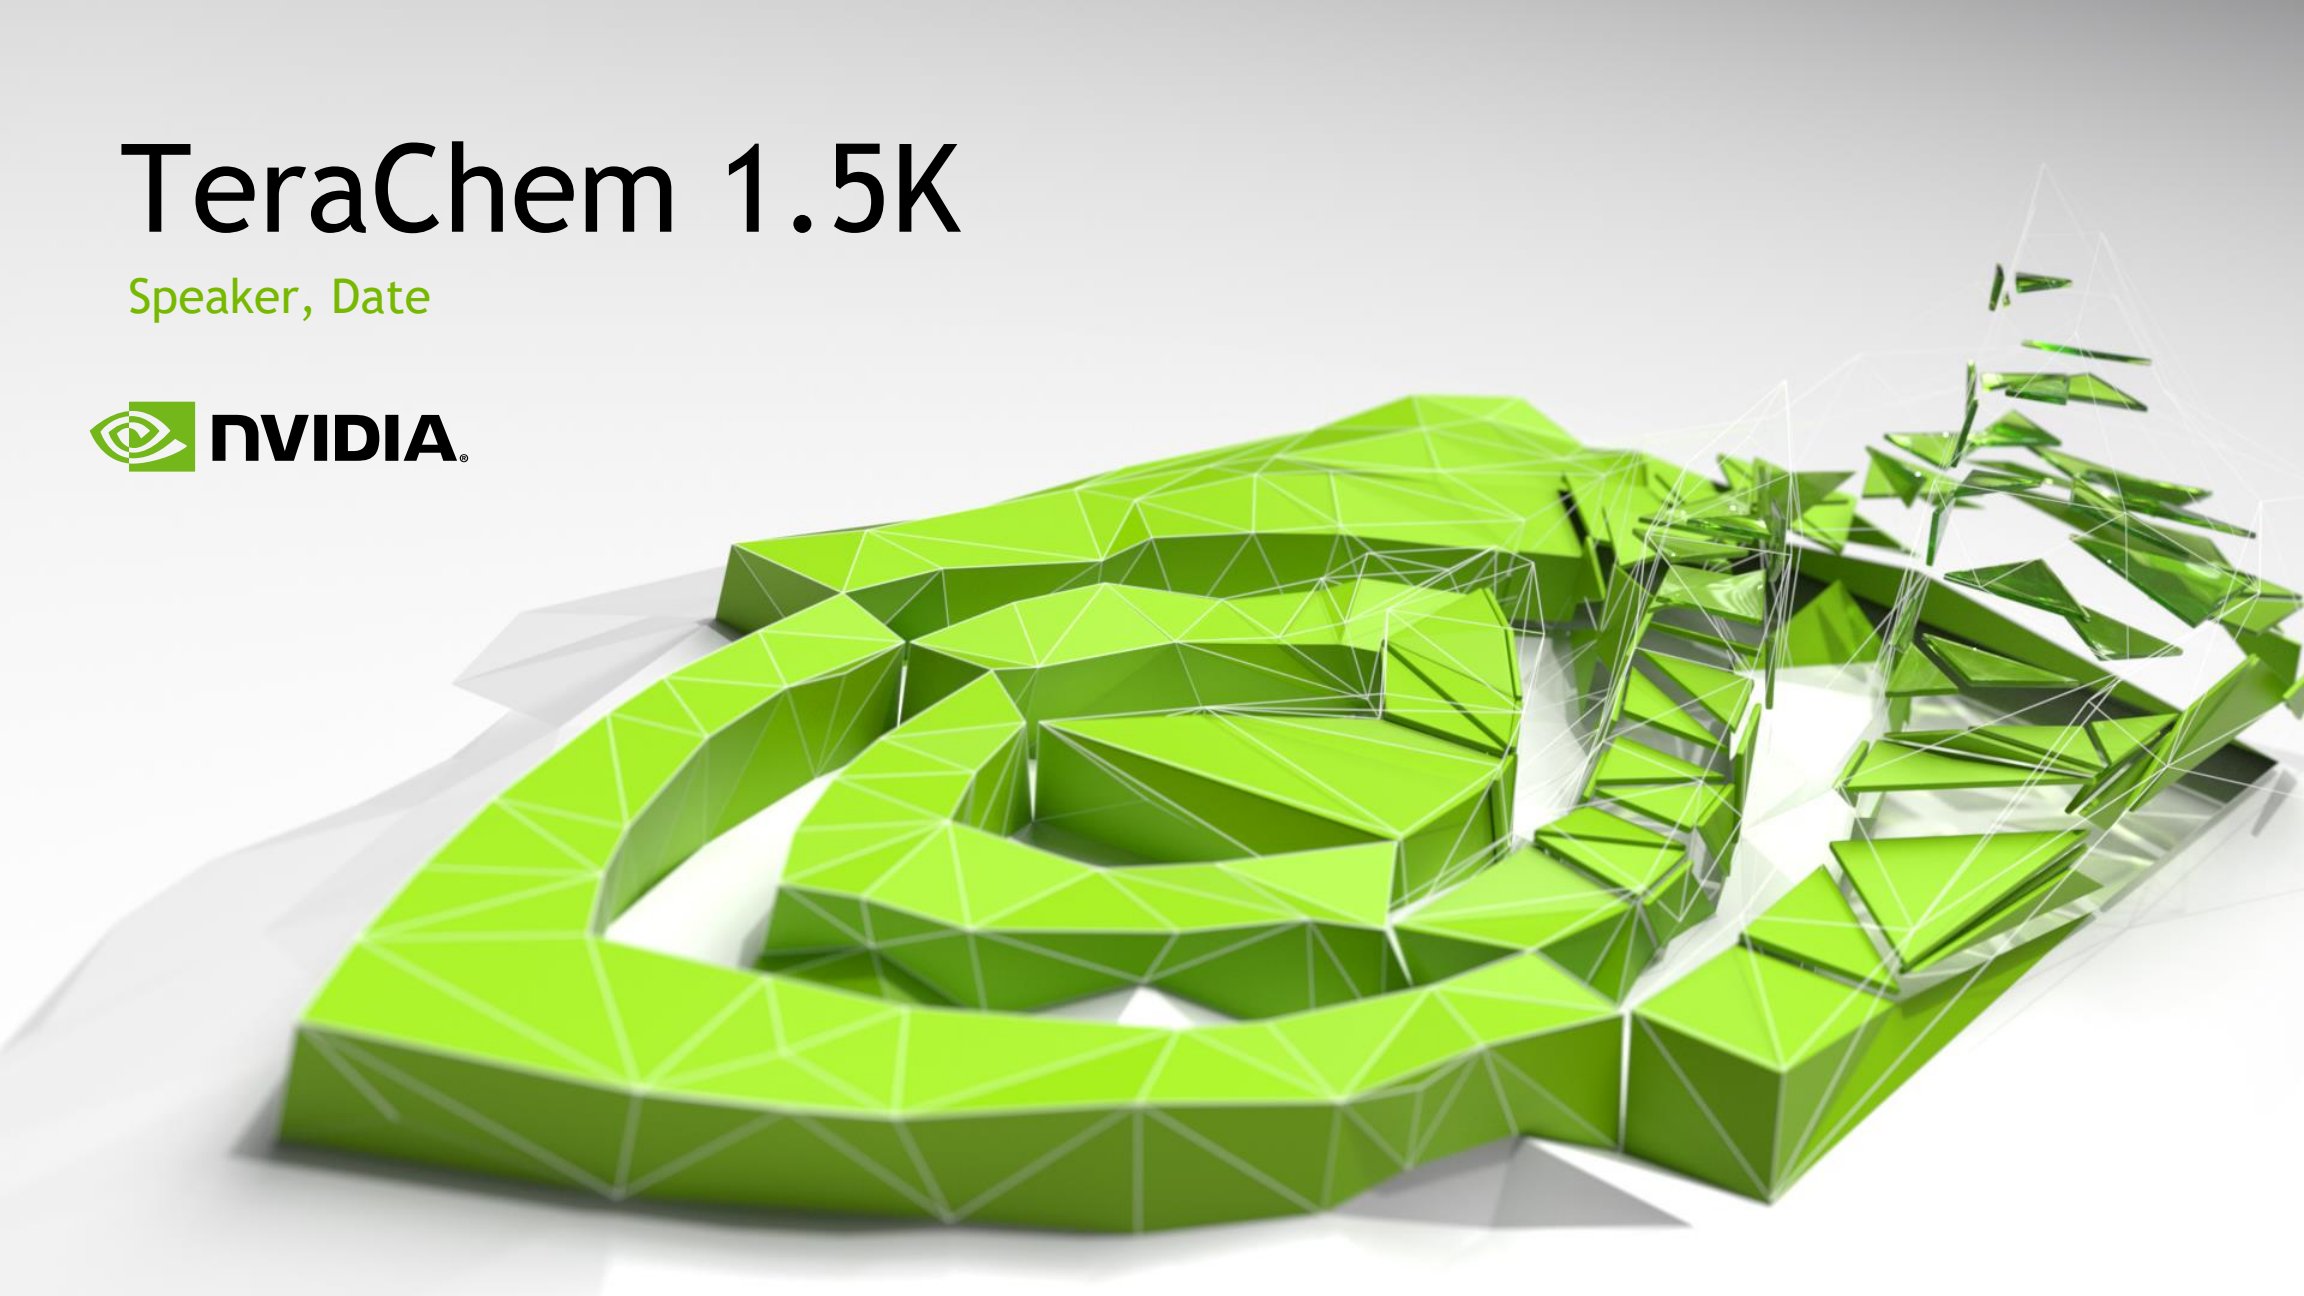

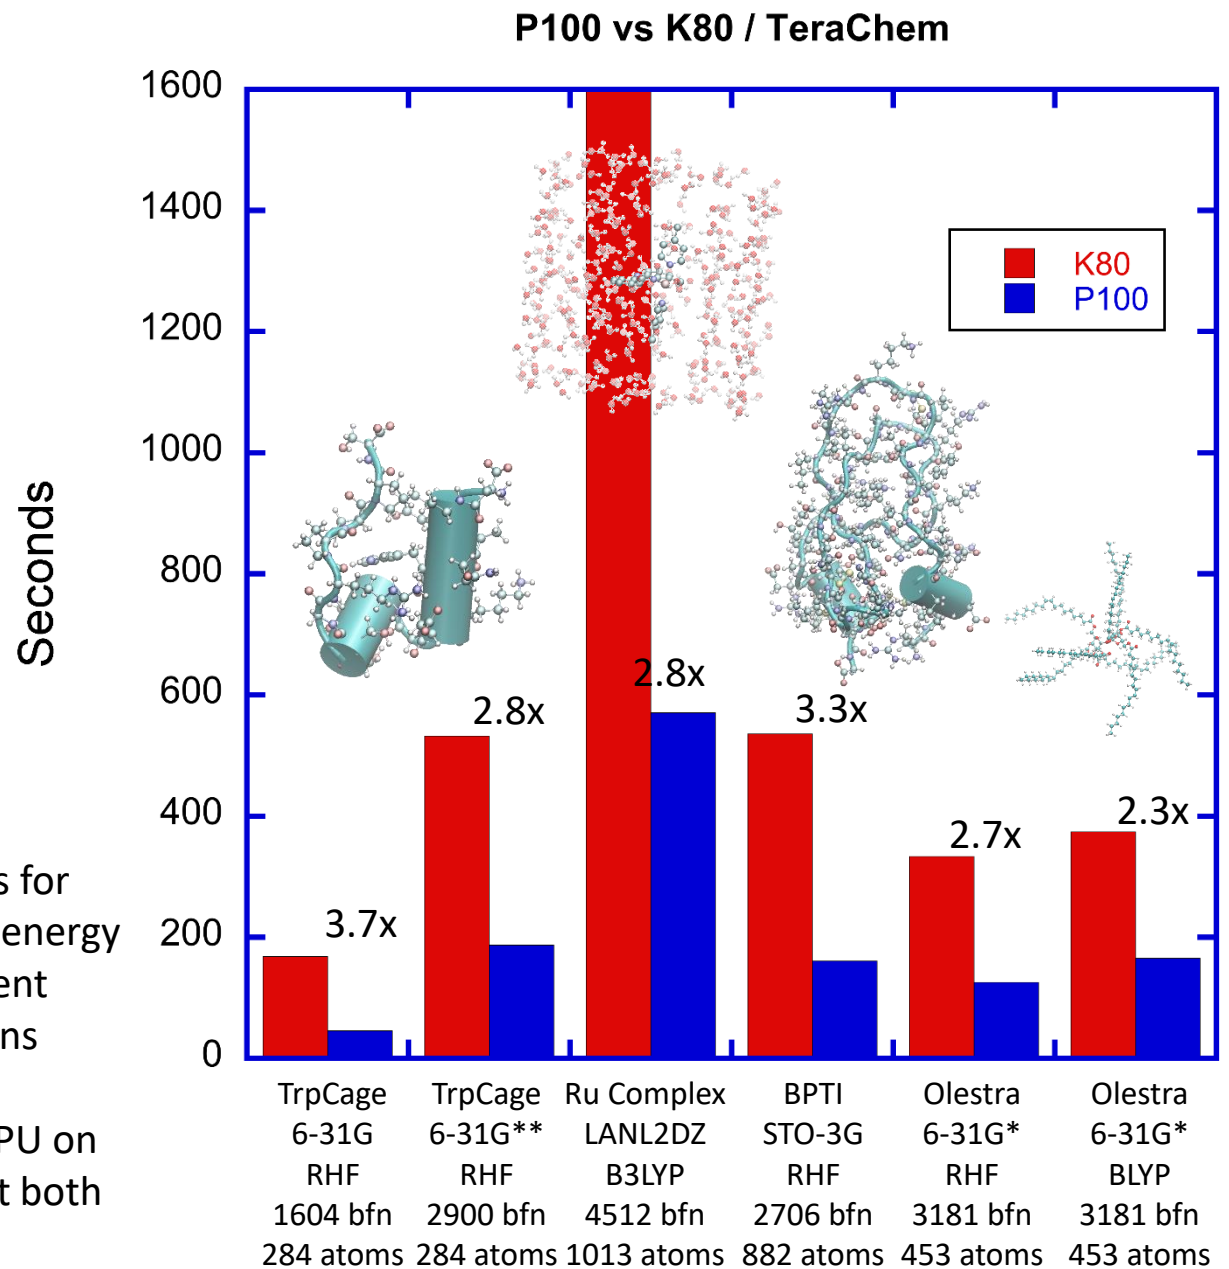

All timings for complete energy and gradient Calculations

K80 = 1 GPU on Board, not both

Slide courtesy of PetaChem LLC / Todd Martinez

TERACHEM 1.5K; TRIPCAGE ON TESLA K40S

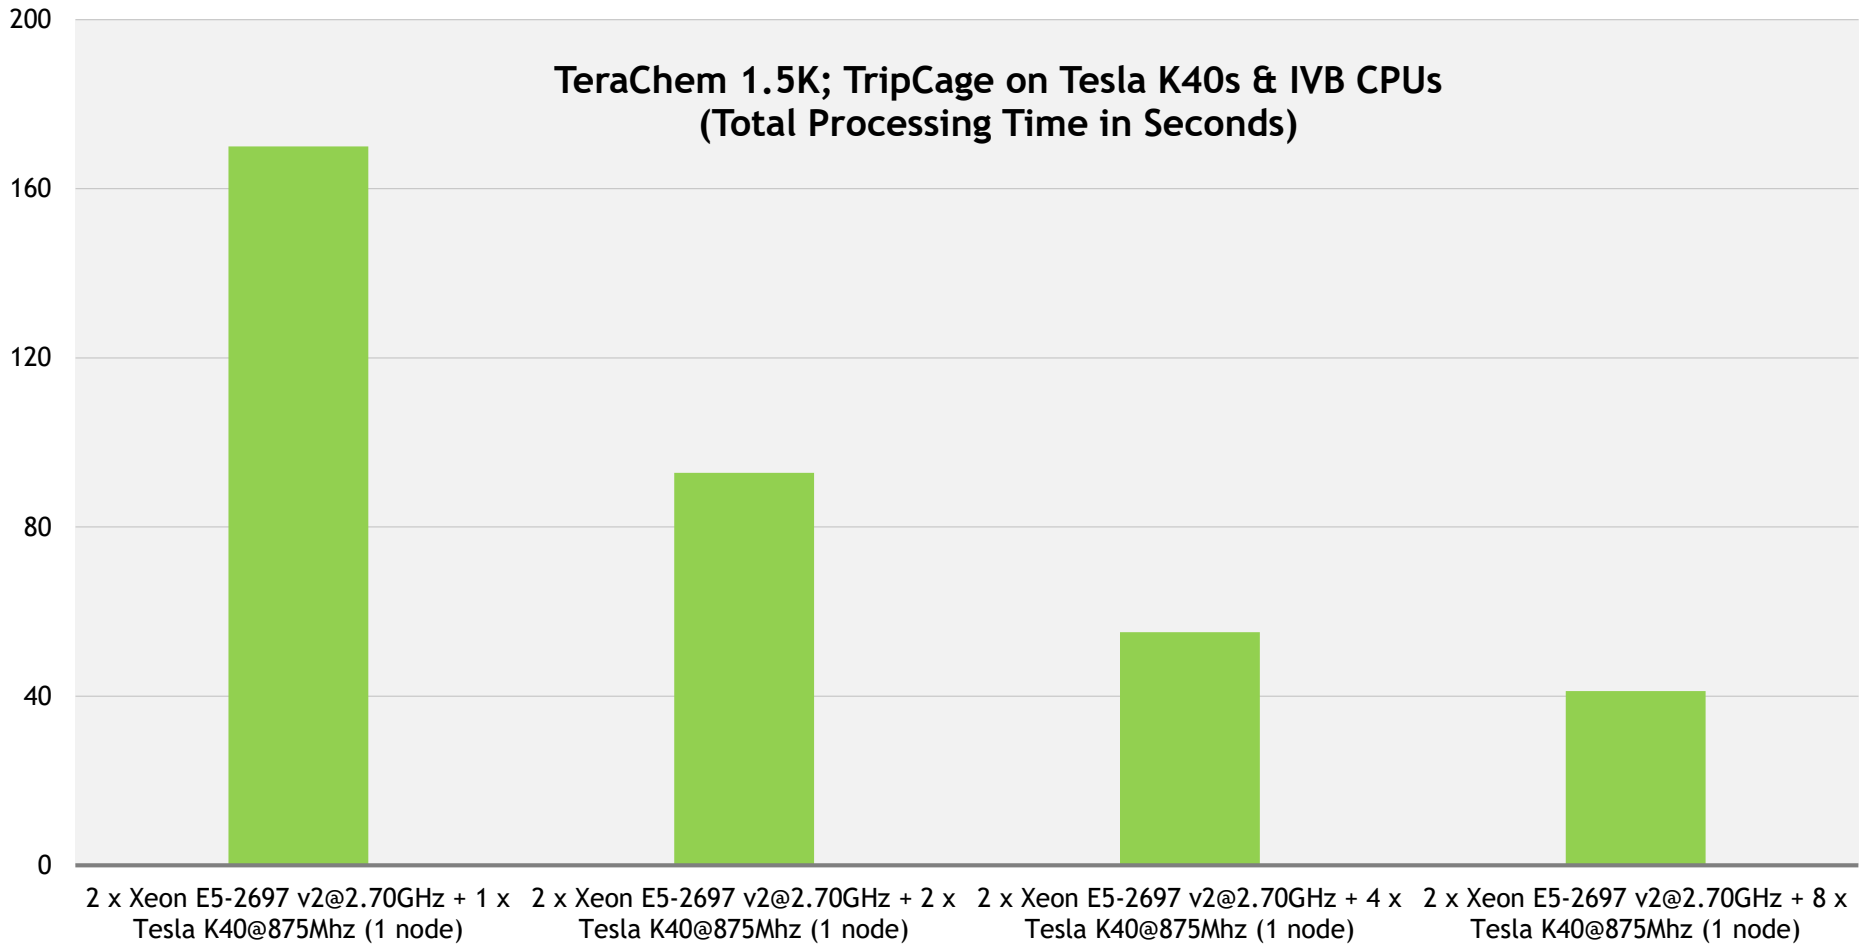

TERACHEM 1.5K; TRIPCAGE ON TESLA K40S & HASWELL CPUS

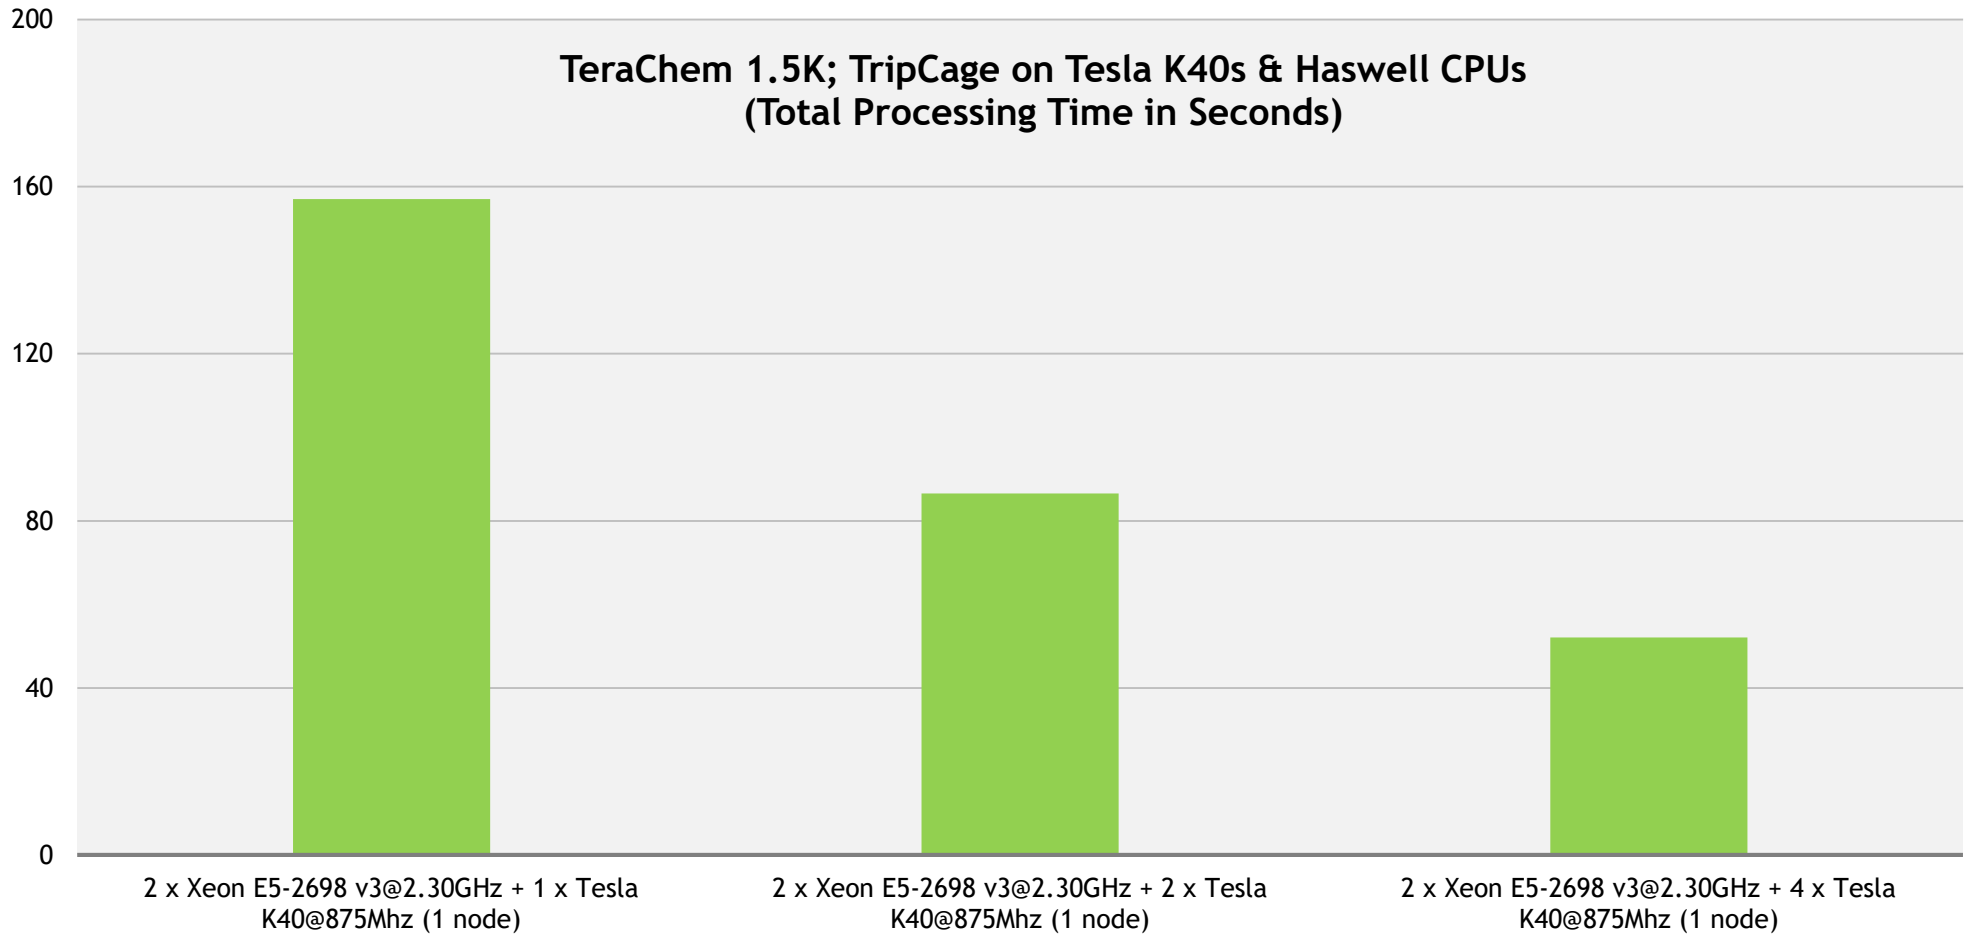

TERACHEM 1.5K; TRIPCAGE ON TESLA K80S & IVB CPUS

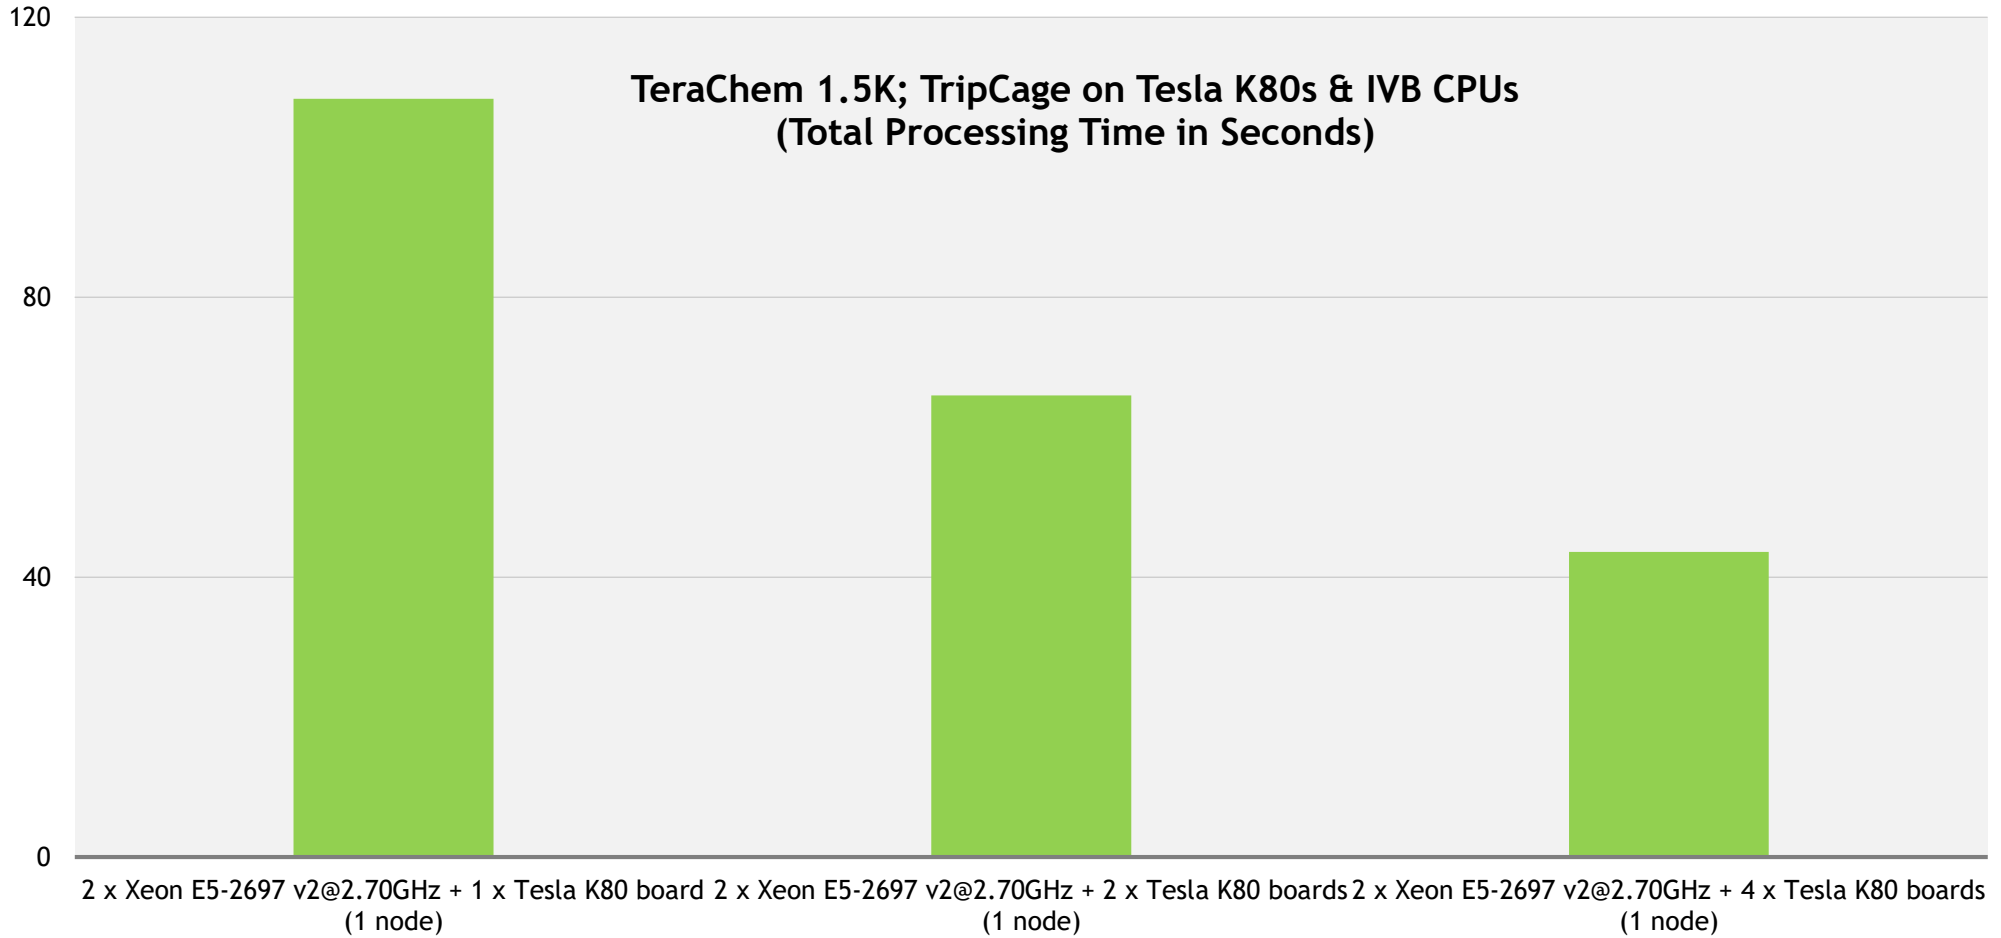

TERACHEM 1.5K; TRIPCAGE ON TESLA K80S & HASWELL CPUS

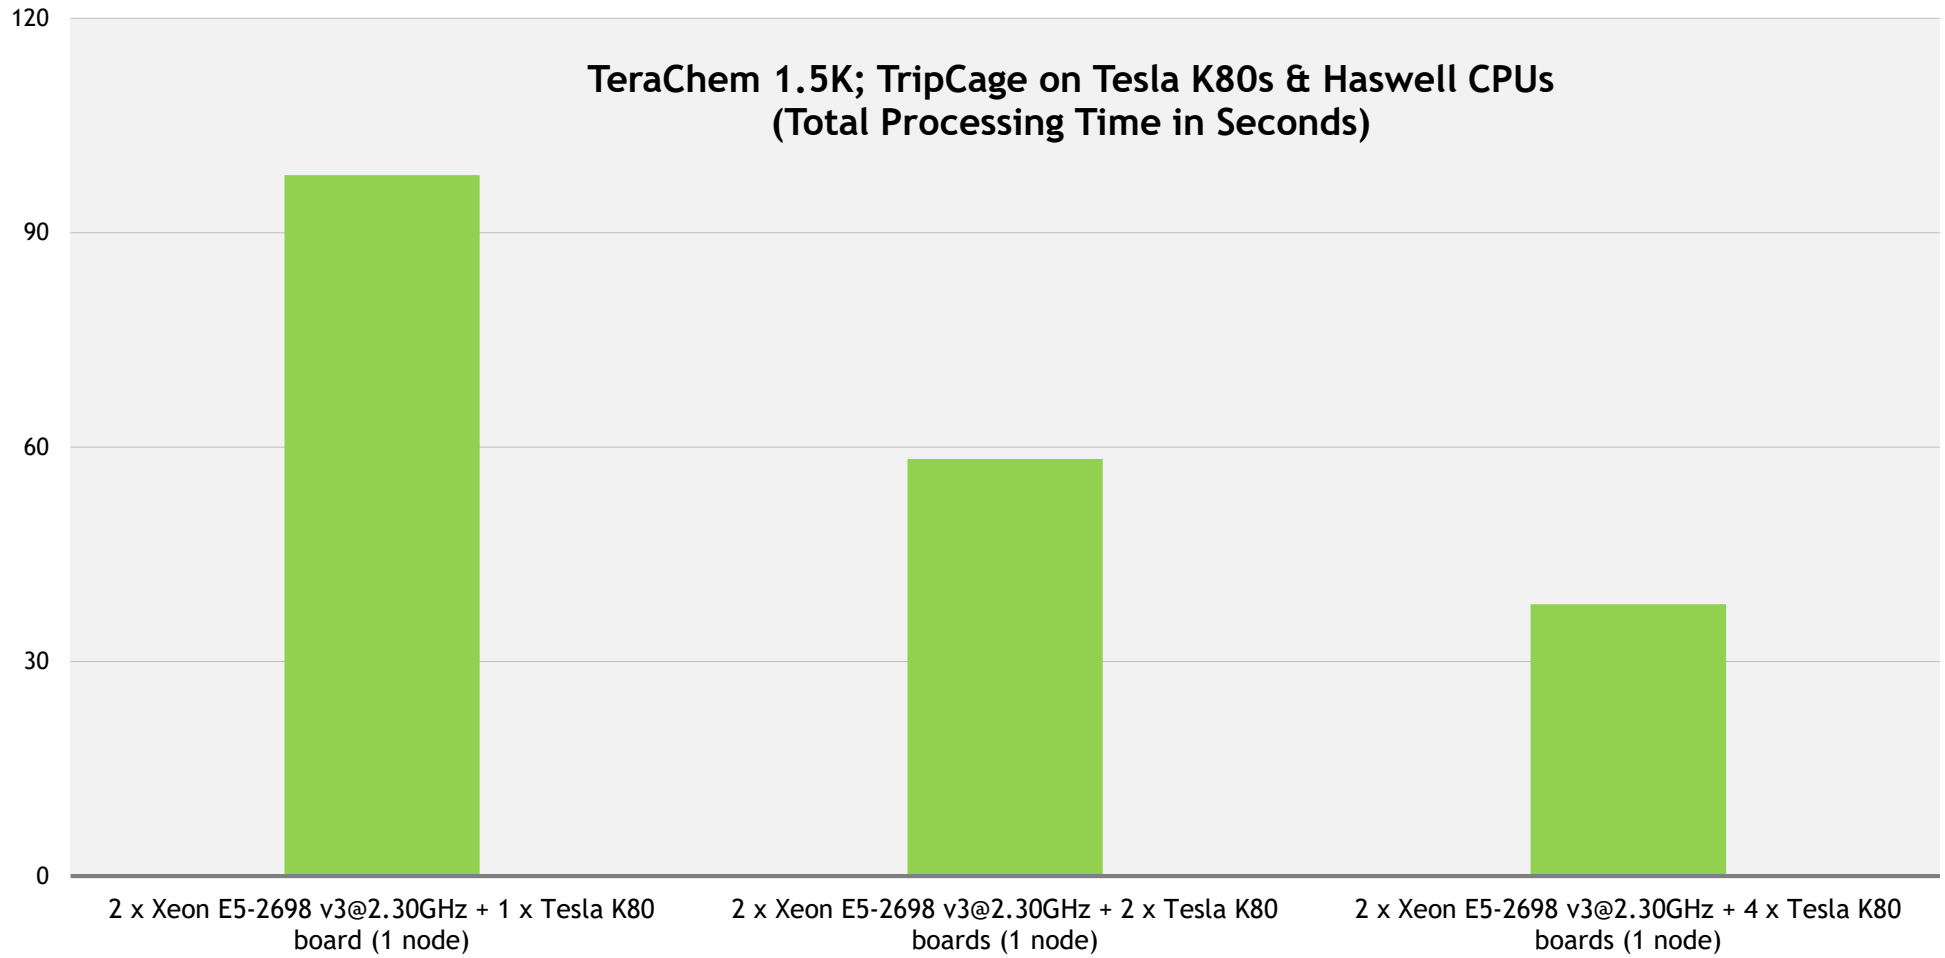

TERACHEM 1.5K; BPTI ON TESLA K40S & IVB CPUS

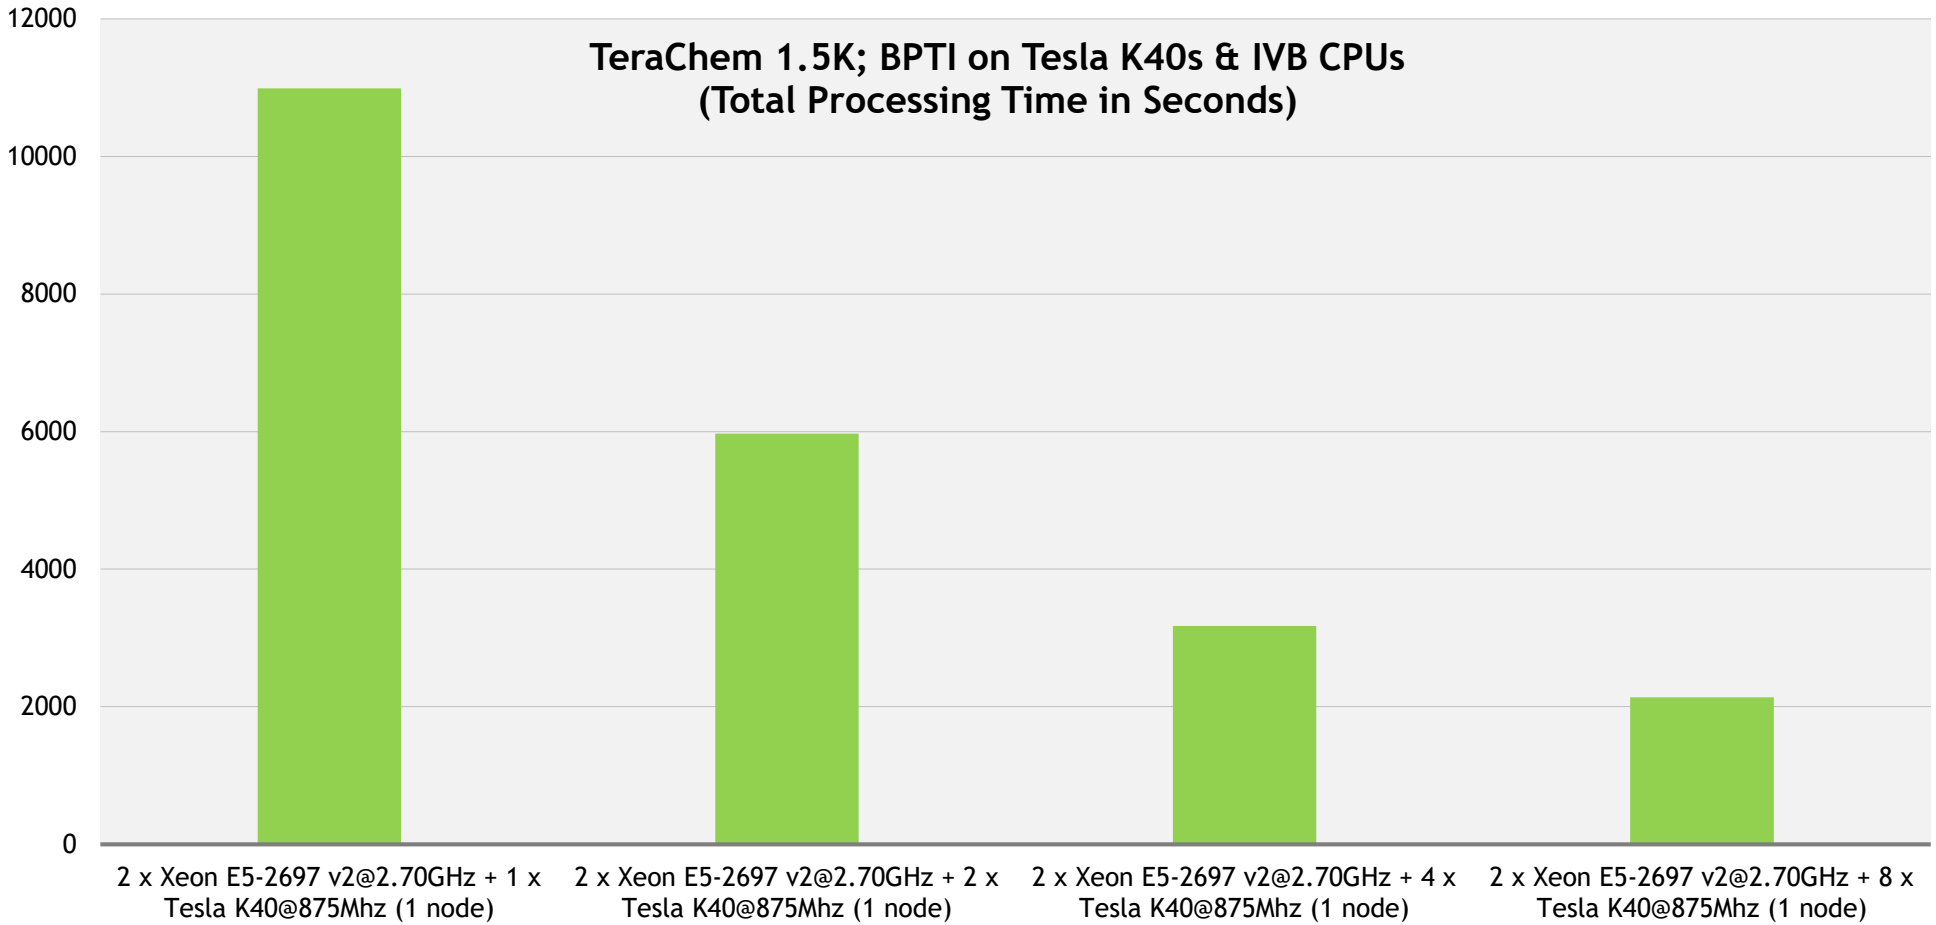

TERACHEM 1.5K; BPTI ON TESLA K80S & IVB CPUS

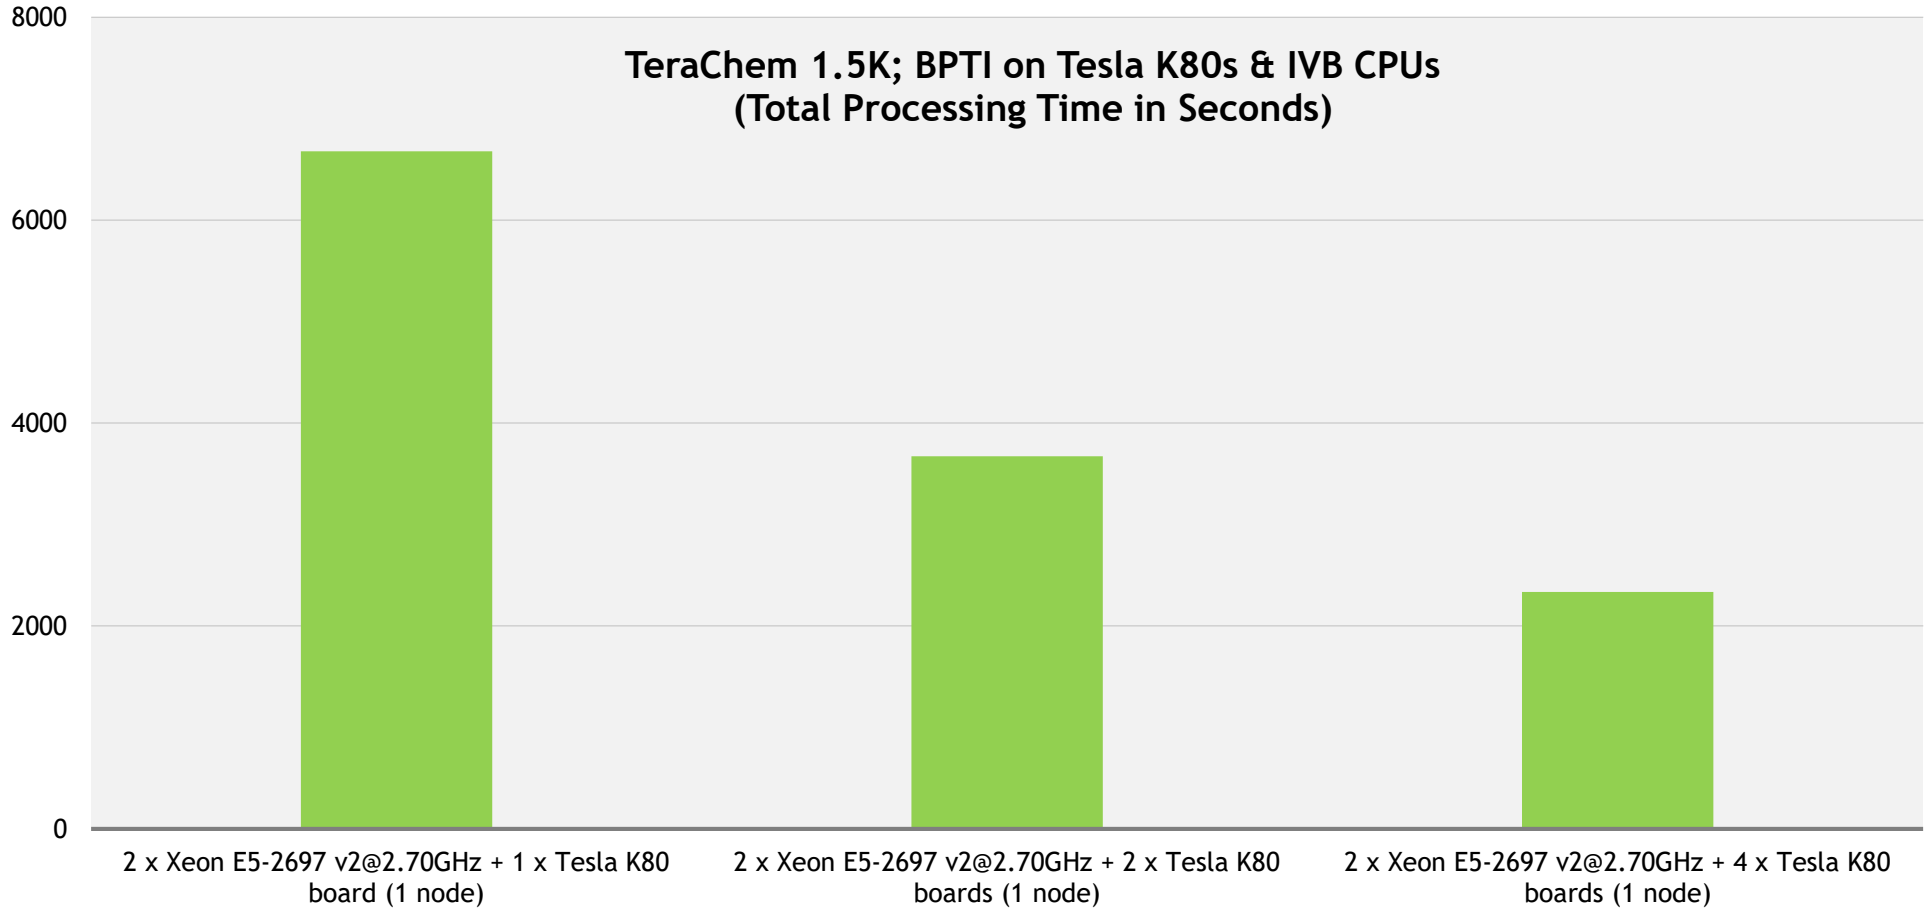

TERACHEM 1.5K; BPTI ON TESLA K40S & IVB CPUS

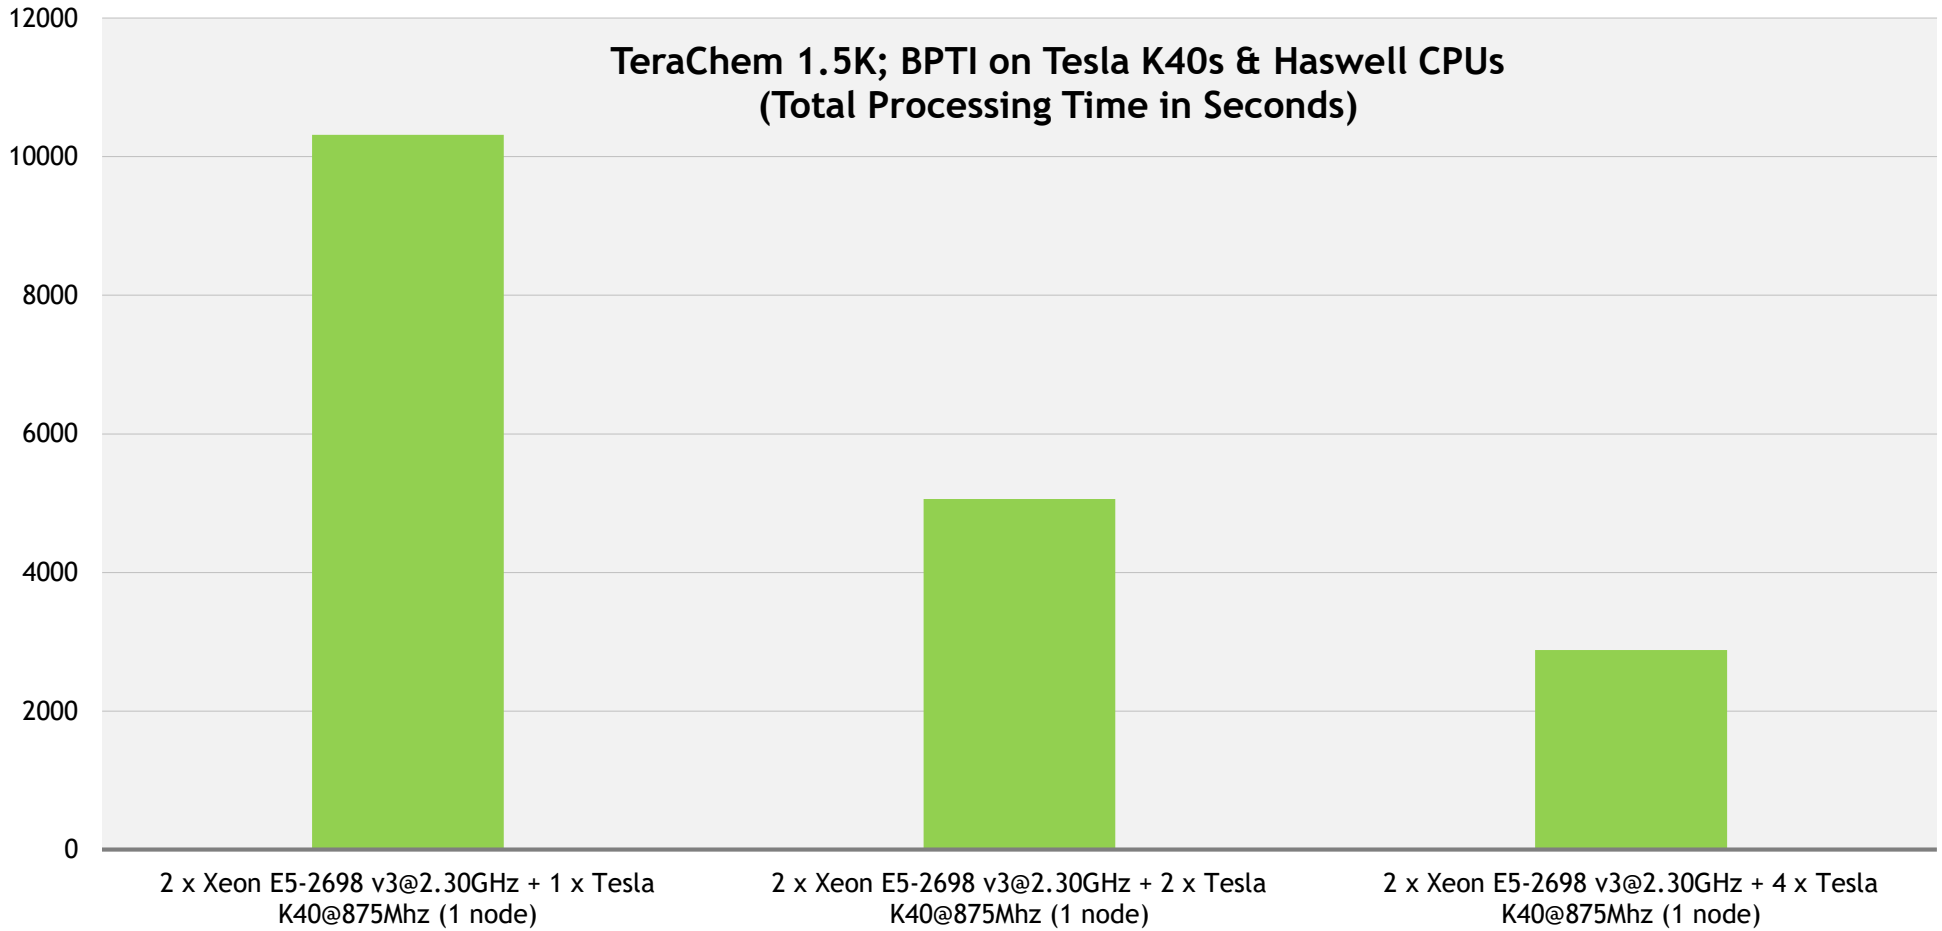

TERACHEM 1.5K; BPTI ON TESLA K80S & HASWELL CPUS

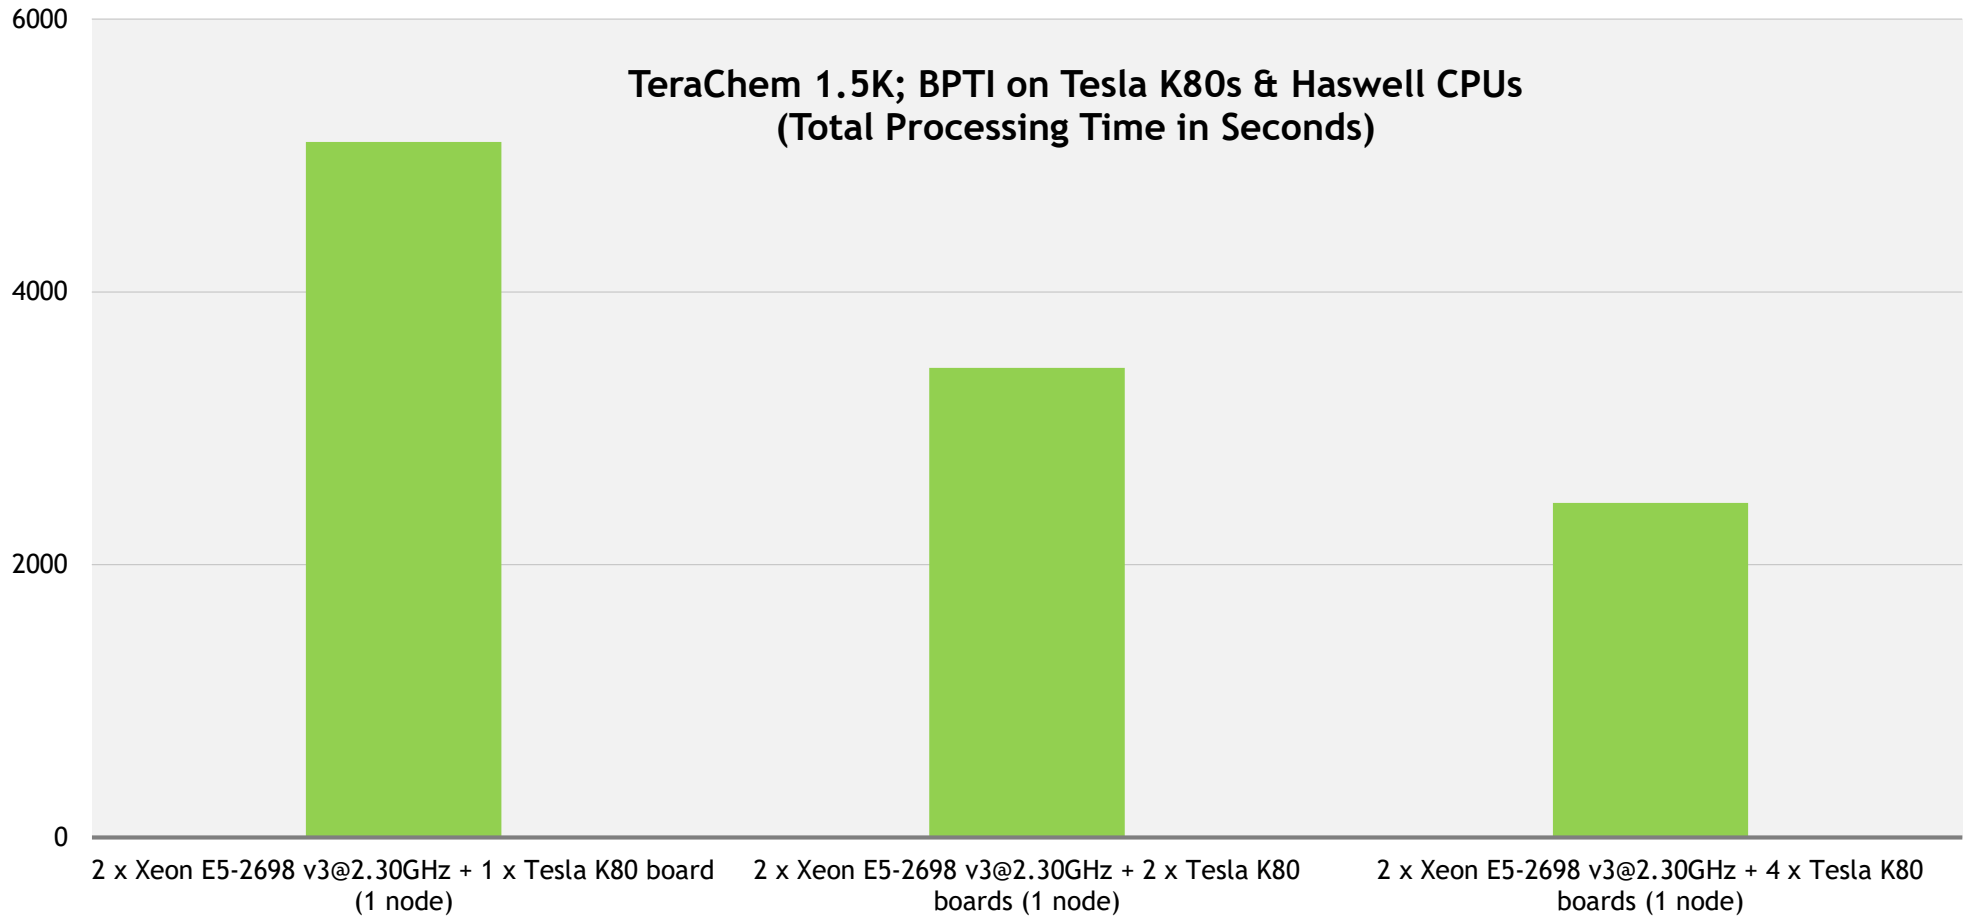

Beyond VASP 5.4.4 Development

March 2018

VASP

The Vienna Ab Initio Simulation Package

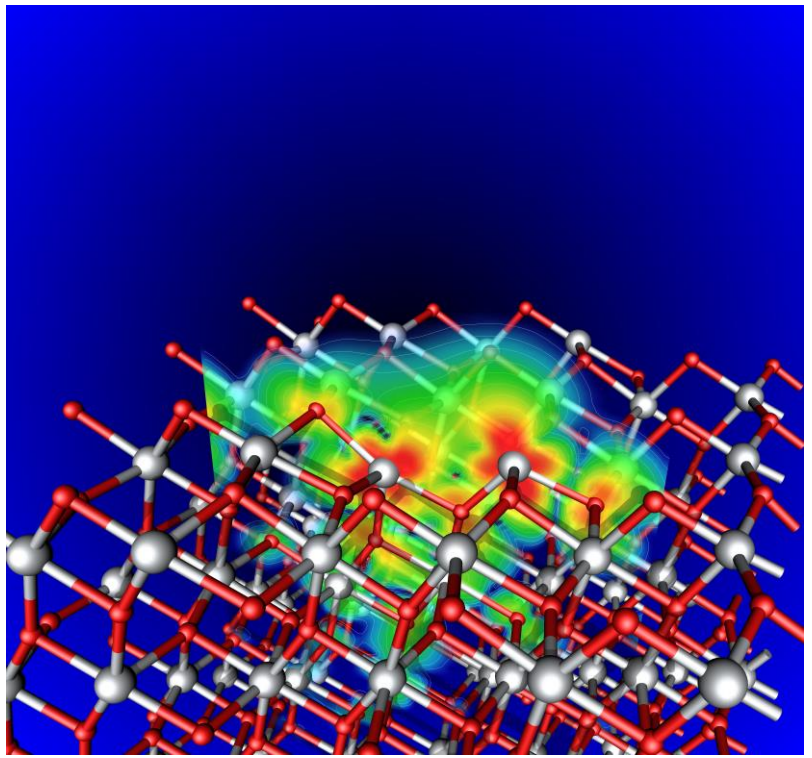

Prof. Georg Kresse
Computational Materials Physics
University of Vienna

“For VASP, OpenACC is ***the*** way forward for GPU acceleration. Performance is similar to CUDA, and OpenACC dramatically decreases GPU development and maintenance efforts. We’re excited to collaborate with NVIDIA and PGI as an early adopter of Unified Memory.”

PGI

B.aP107 on V100s PCIe

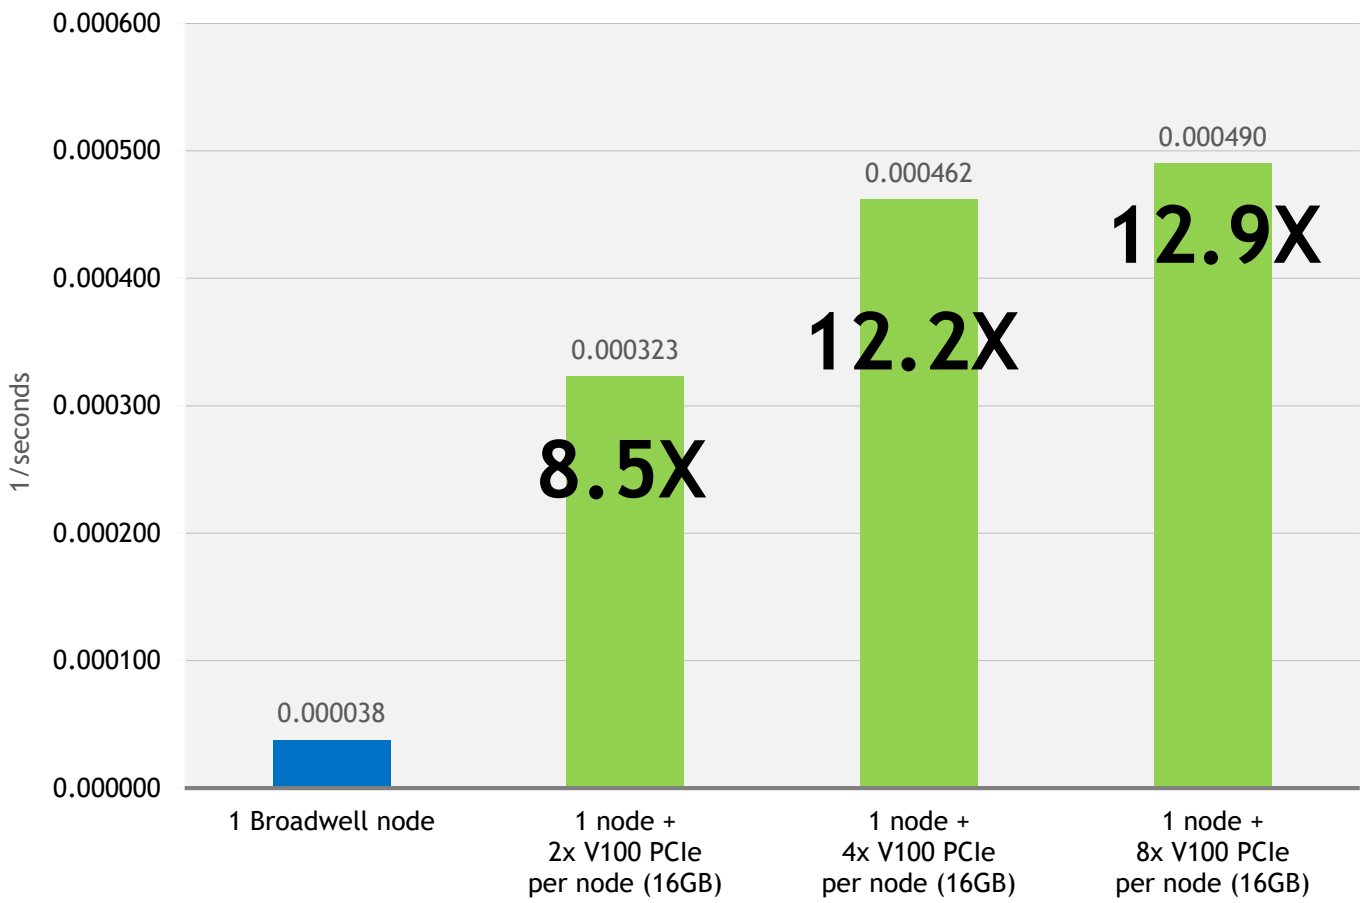

(Untuned on Volta)
Running VASP version 5.4.4

The blue node contains Dual Intel Xeon E5-2690 v4@2.6GHz [3.5GHz Turbo] (Broadwell) CPUs

The green nodes contain Dual Intel Xeon E5-2690 v4@2.6GHz [3.5GHz Turbo] (Broadwell) CPUs + Tesla V100 PCIe (16GB) GPUs

107 Boron atoms (symmetry broken 107-atom B' variant)
216 bands
110592 plane waves
Hybrid functional calculation (exact exchange) with blocked Davidson. No KPoint parallelization.
Hybrid Functional with blocked Davidson (ALGO=Normal)
LHFCALC=.True. (Exact Exchange)

Silica IFPEN on V100s PCIe

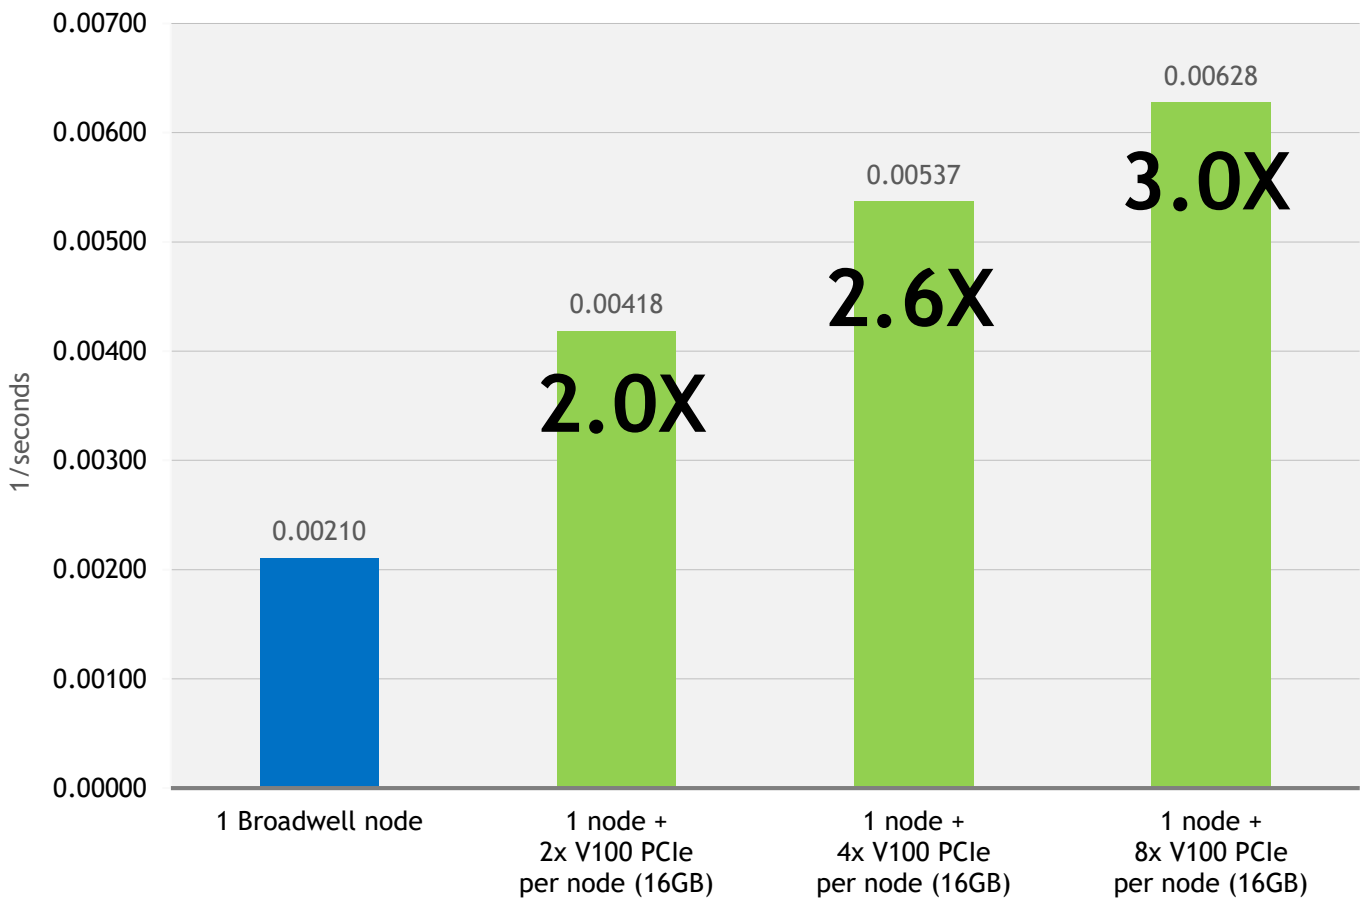

(Untuned on Volta)
Running VASP version 5.4.4

The blue node contains Dual Intel Xeon E5-2690 v4@2.6GHz [3.5GHz Turbo] (Broadwell) CPUs

The green nodes contain Dual Intel Xeon E5-2690 v4@2.6GHz [3.5GHz Turbo] (Broadwell) CPUs + Tesla V100 PCIe (16GB) GPUs

240 ions, cristobalite (high) bulk
720 bands
? plane waves
ALGO = Very Fast (RMM-DIIS)

VASP Dev Version OpenACC Performance

silica IFPEN, RMM-DIIS on P100

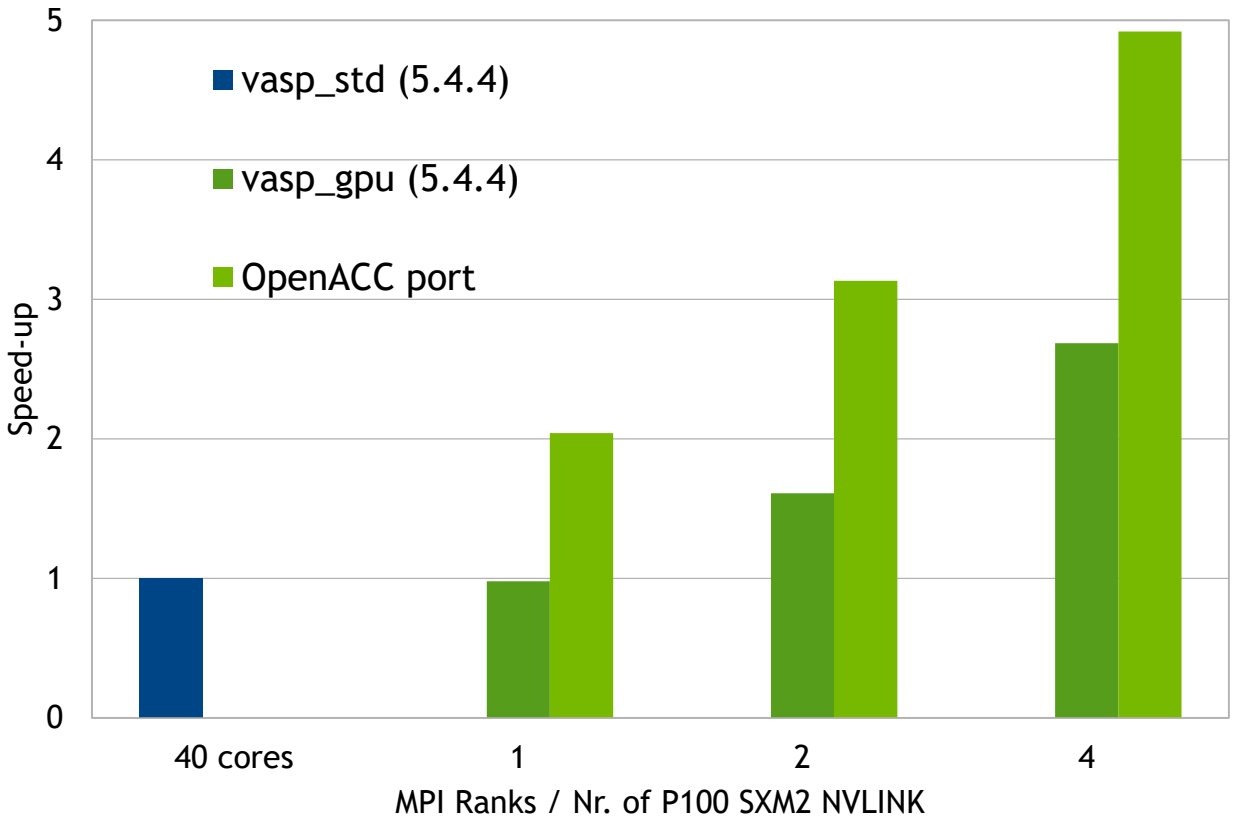

Reasons for performance gain:

- OpenACC port covers more VASP routines
- OpenACC port planned top down, with complete analysis of the call tree
- OpenACC port leverages improvements in latest VASP Fortran source base

PGI

VASP Dev Version OpenACC Performance

VASP with OpenACC is an ongoing project. The developers in Vienna are integrating some OpenACC into their code base and there are plans to continue porting in that direction. Given that there is ongoing work, chances are the next major release will contain some of the results of this work. But from NV's side, we can neither comment on any release schedule for VASP in general, nor on any feature set for upcoming releases.

VASP 5.4.4

February 2018

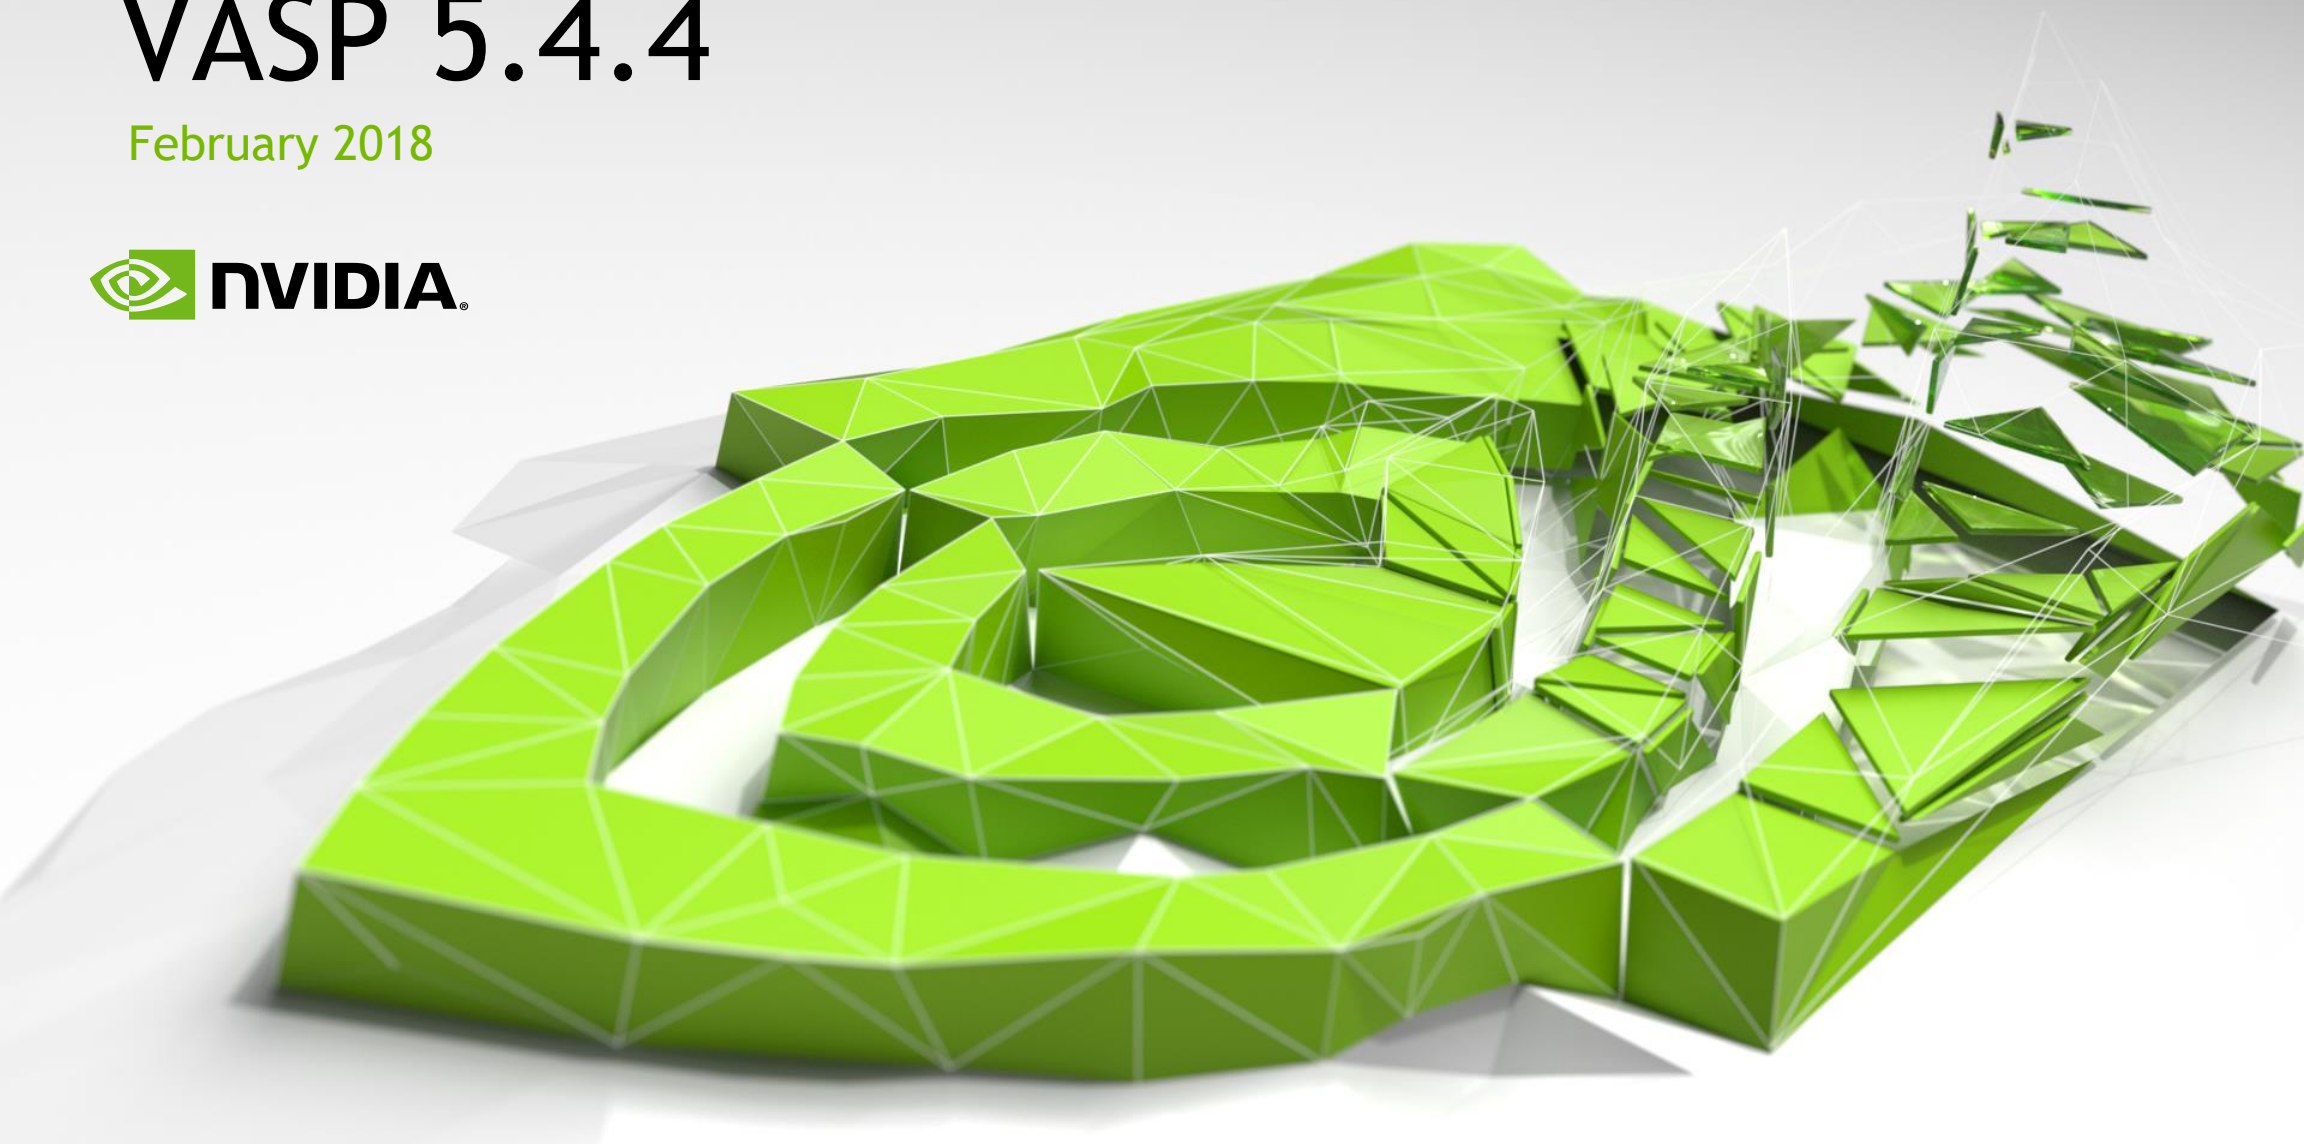

VASP Silica IFPEN
on Gold 6140 vs V100 PCIe (32GB)

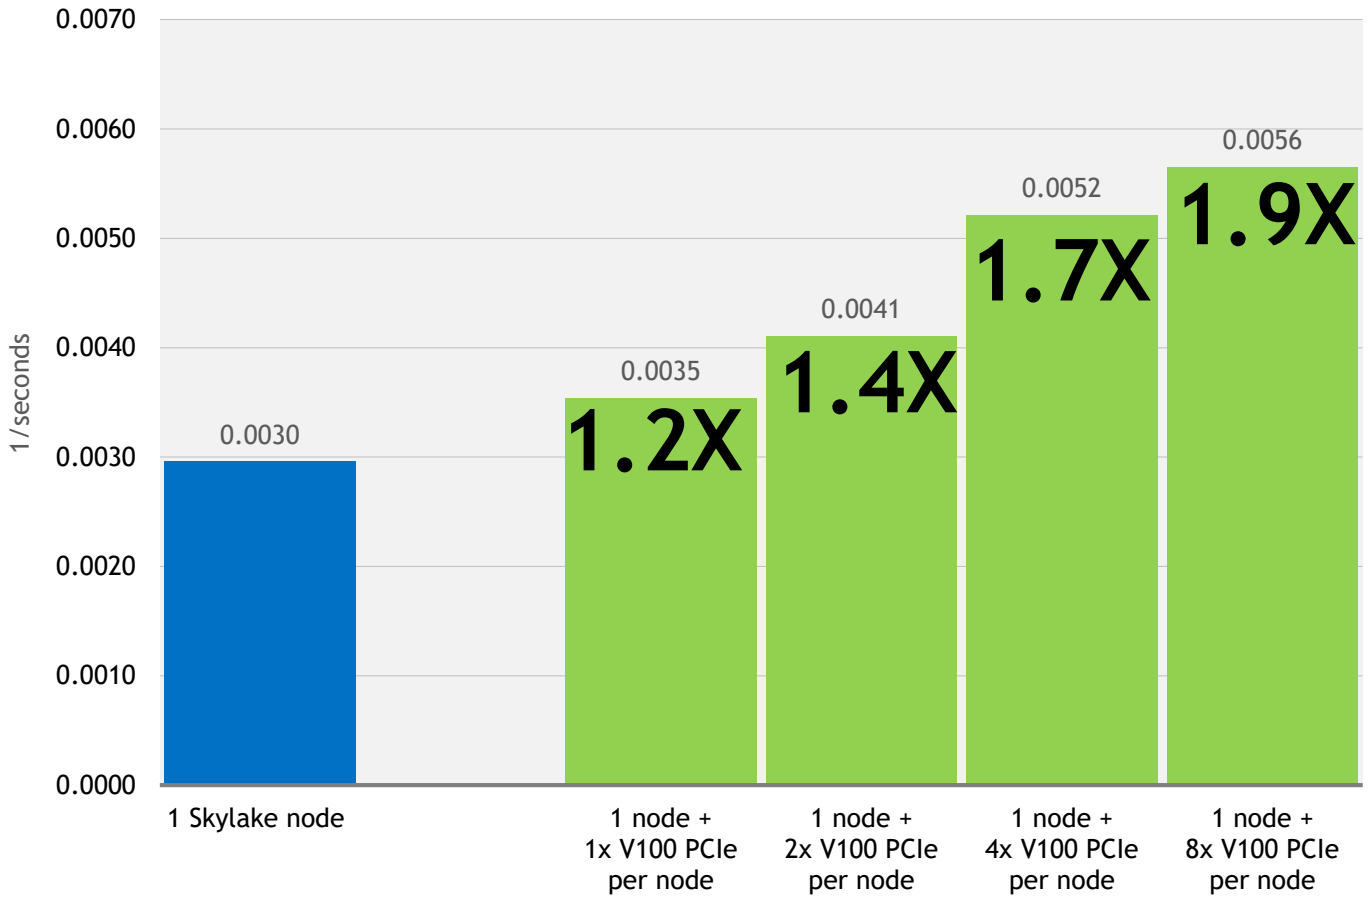

(Untuned on Volta)
Running VASP version 5.4.4

The blue node contains Dual Intel Xeon Gold 6140@2.30GHz [2.301GHz Turbo] (Skylake) CPUs

The green nodes contain Dual Intel Xeon E5-2690 v4@2.6GHz [3.5GHz Turbo] (Broadwell) CPUs + Tesla V100 PCIe (32GB) GPUs

240 ions, cristobalite (high) bulk
720 bands
? plane waves
ALGO = Very Fast (RMM-DIIS)

VASP Silica IFPEN
on Gold 6140 vs V100 SXM2 (32GB)

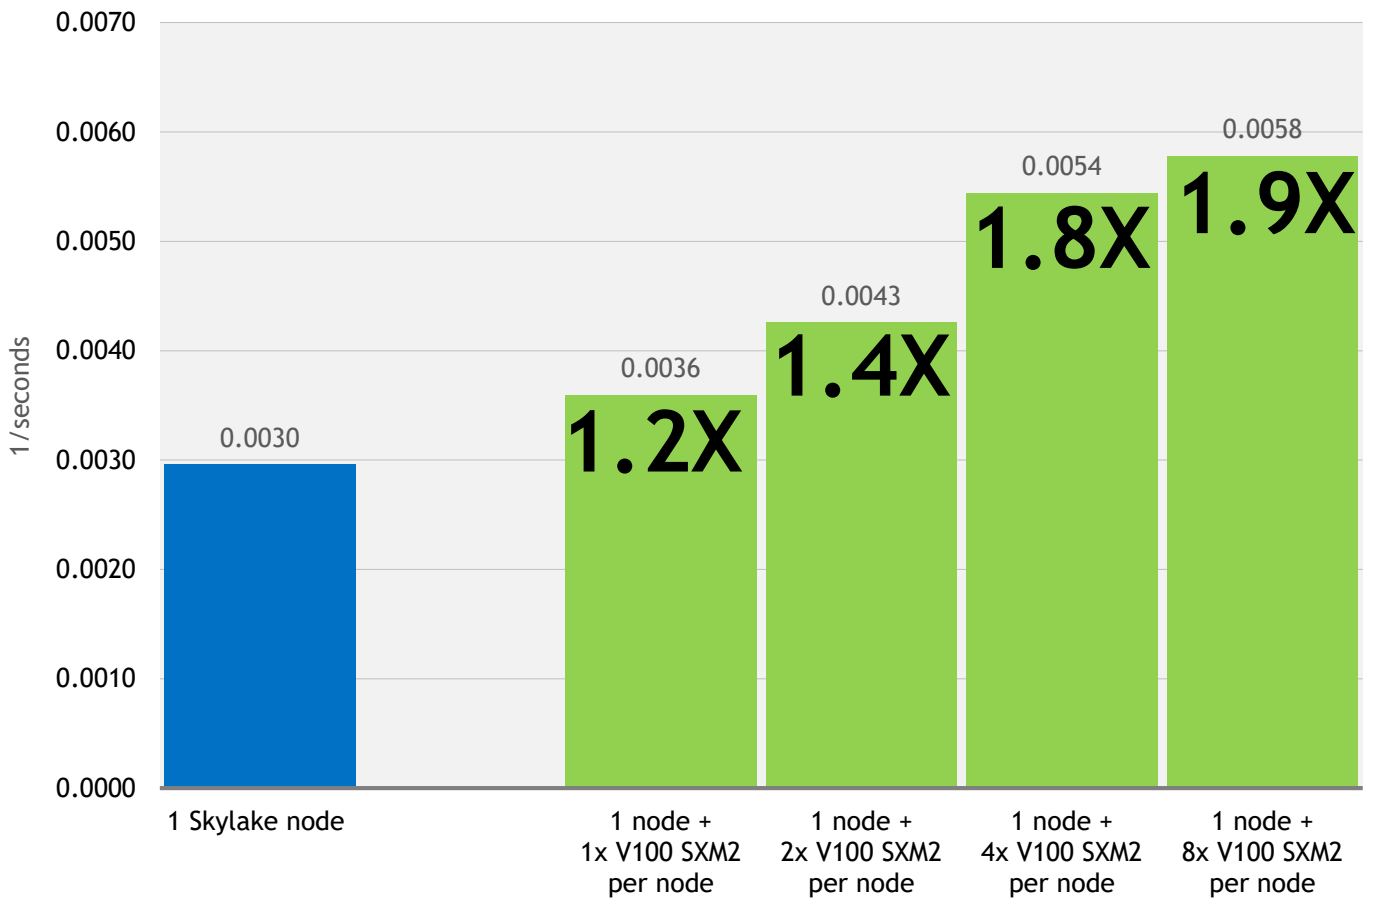

(Untuned on Volta)
Running VASP version 5.4.4

The blue node contains Dual Intel Xeon Gold 6140@2.30GHz [2.301GHz Turbo] (Skylake) CPUs

The green nodes contain Dual Intel Xeon E5-2698 v4@2.2GHz [3.6GHz Turbo] (Broadwell) CPUs + Tesla V100 SXM2 (32GB) GPUs

240 ions, cristobalite (high) bulk
720 bands
? plane waves
ALGO = Very Fast (RMM-DIIS)

VASP Silica IFPEN
on Platinum 8180 vs V100 PCIe (32GB)

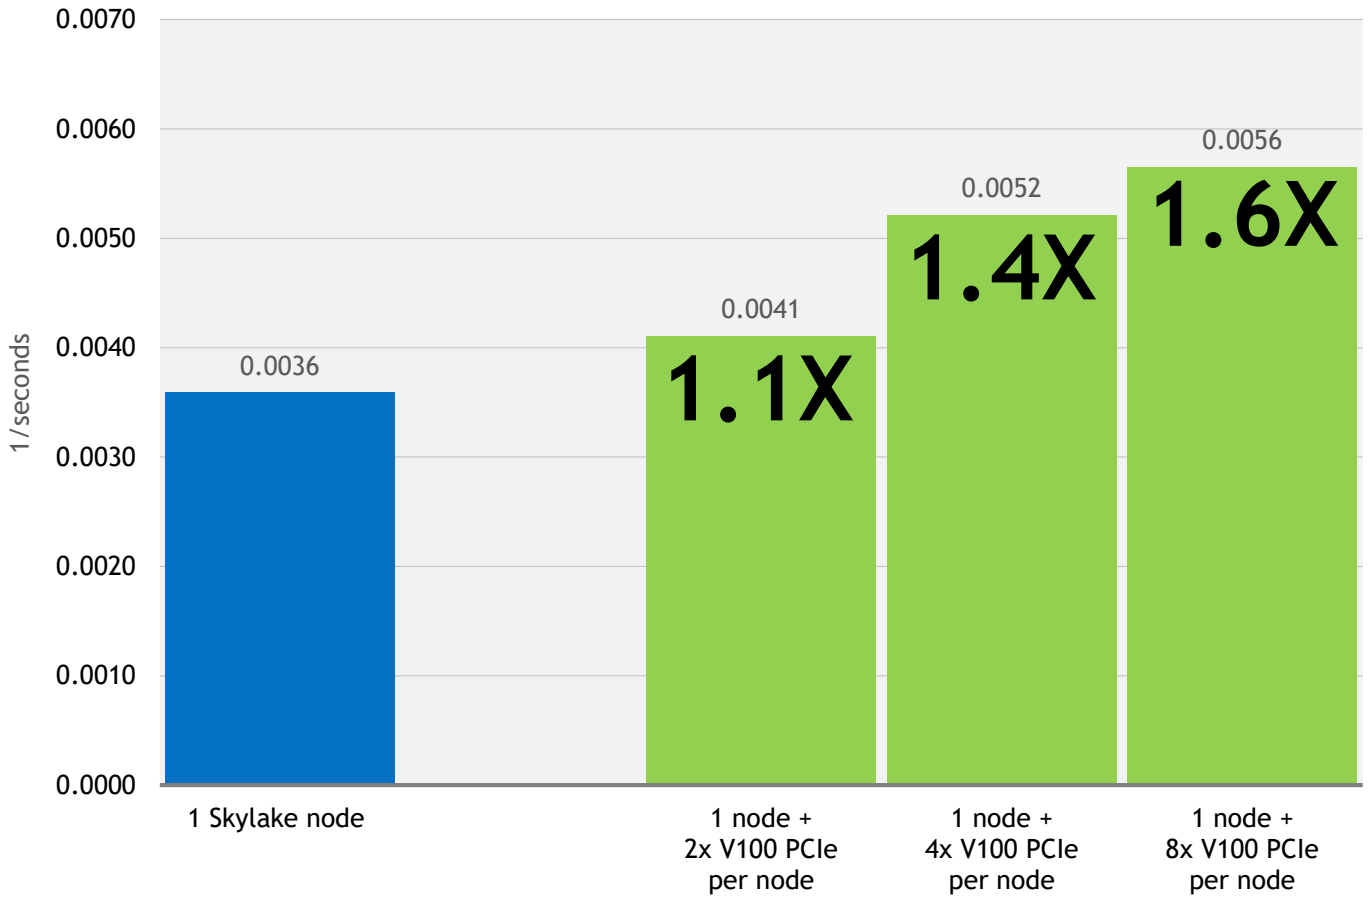

(Untuned on Volta)
Running VASP version 5.4.4

The blue node contains Dual Intel Xeon Platinum 8180@2.50GHz [2.501GHz Turbo] (Skylake) CPUs

The green nodes contain Dual Intel Xeon E5-2690 v4@2.6GHz [3.5GHz Turbo] (Broadwell) CPUs + Tesla V100 PCIe (32GB) GPUs

240 ions, cristobalite (high) bulk
720 bands
? plane waves
ALGO = Very Fast (RMM-DIIS)

VASP Silica IFPEN
on Platinum 8180 vs V100 SXM2 (32GB)

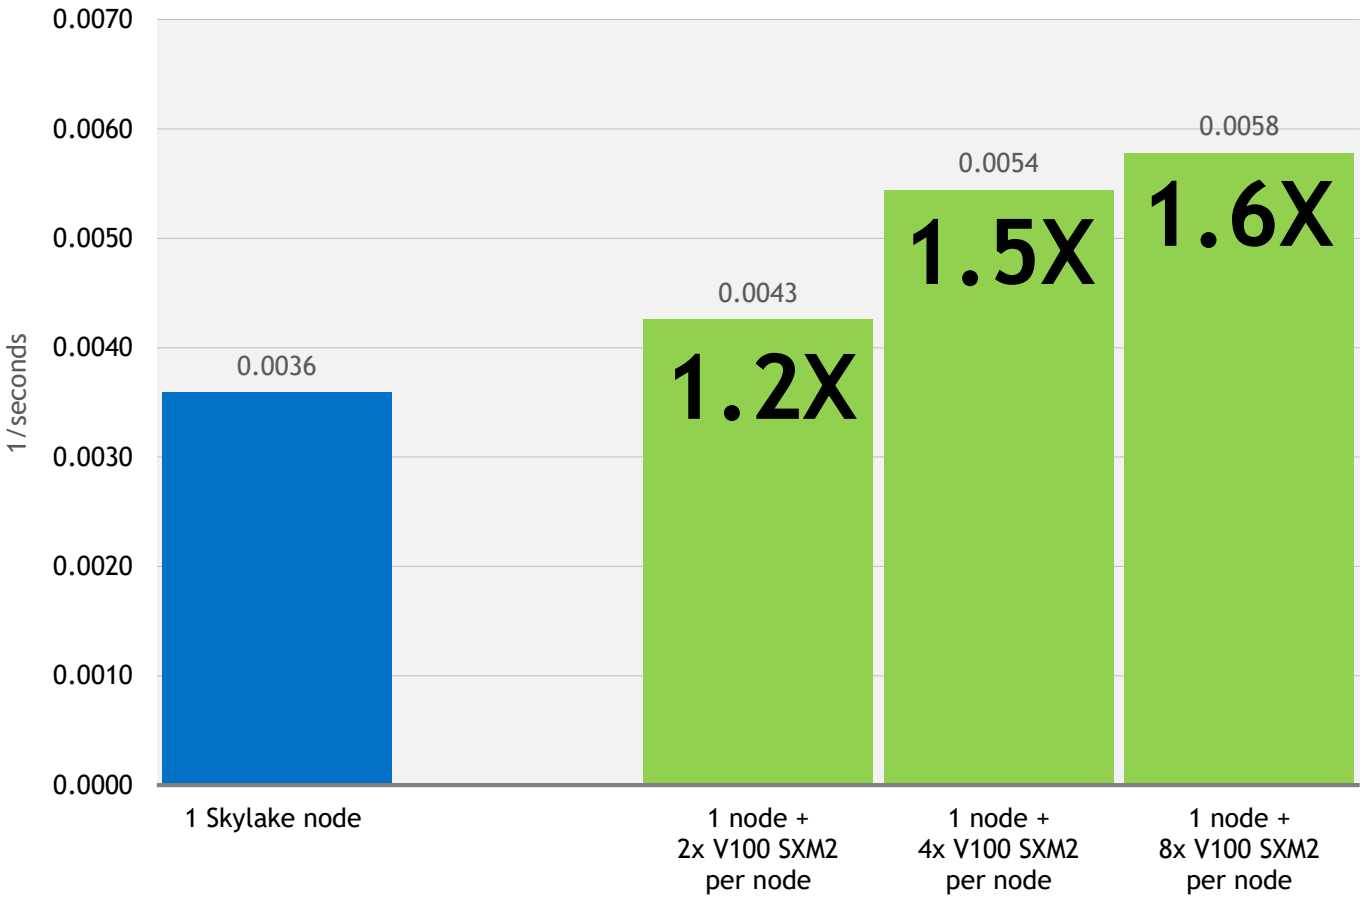

(Untuned on Volta)
Running VASP version 5.4.4

The blue node contains Dual Intel Xeon Platinum 8180@2.50GHz [2.501GHz Turbo] (Skylake) CPUs

The green nodes contain Dual Intel Xeon E5-2698 v4@2.2GHz [3.6GHz Turbo] (Broadwell) CPUs + Tesla V100 SXM2 (32GB) GPUs

240 ions, cristobalite (high) bulk
720 bands
? plane waves
ALGO = Very Fast (RMM-DIIS)

VASP Si-Huge
on Gold 6140 vs V100 PCIe (32GB)

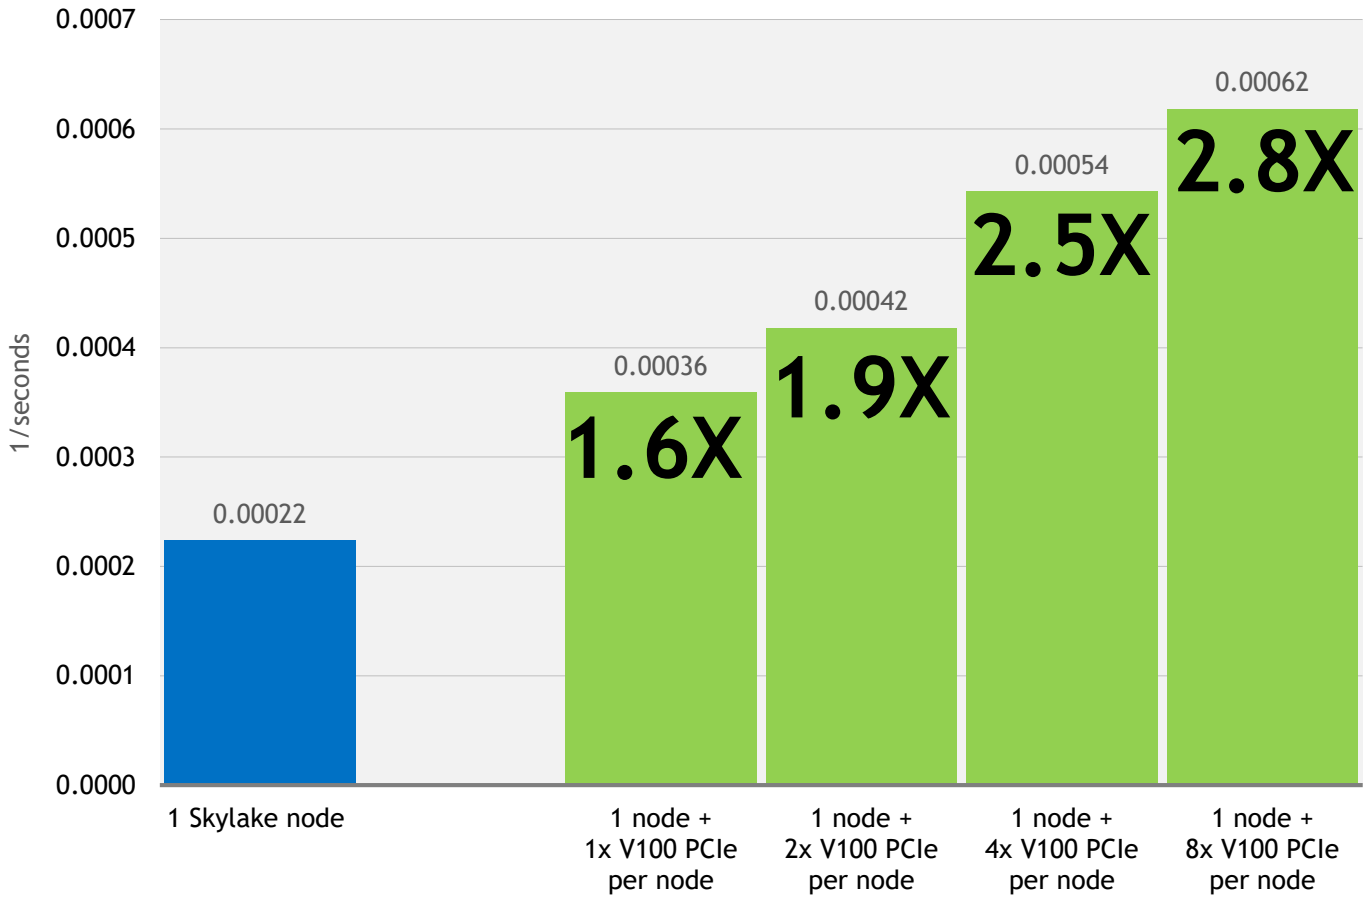

(Untuned on Volta)
Running VASP version 5.4.4

The blue node contains Dual Intel Xeon Gold 6140@2.30GHz [2.301GHz Turbo] (Skylake) CPUs

The green nodes contain Dual Intel Xeon E5-2690 v4@2.6GHz [3.5GHz Turbo] (Broadwell) CPUs + Tesla V100 PCIe (32GB) GPUs

512 Si atoms
1282 bands
864000 Plane Waves
Algo = Normal (blocked Davidson)

VASP Si-Huge
on Gold 6140 vs V100 SXM2 (32GB)

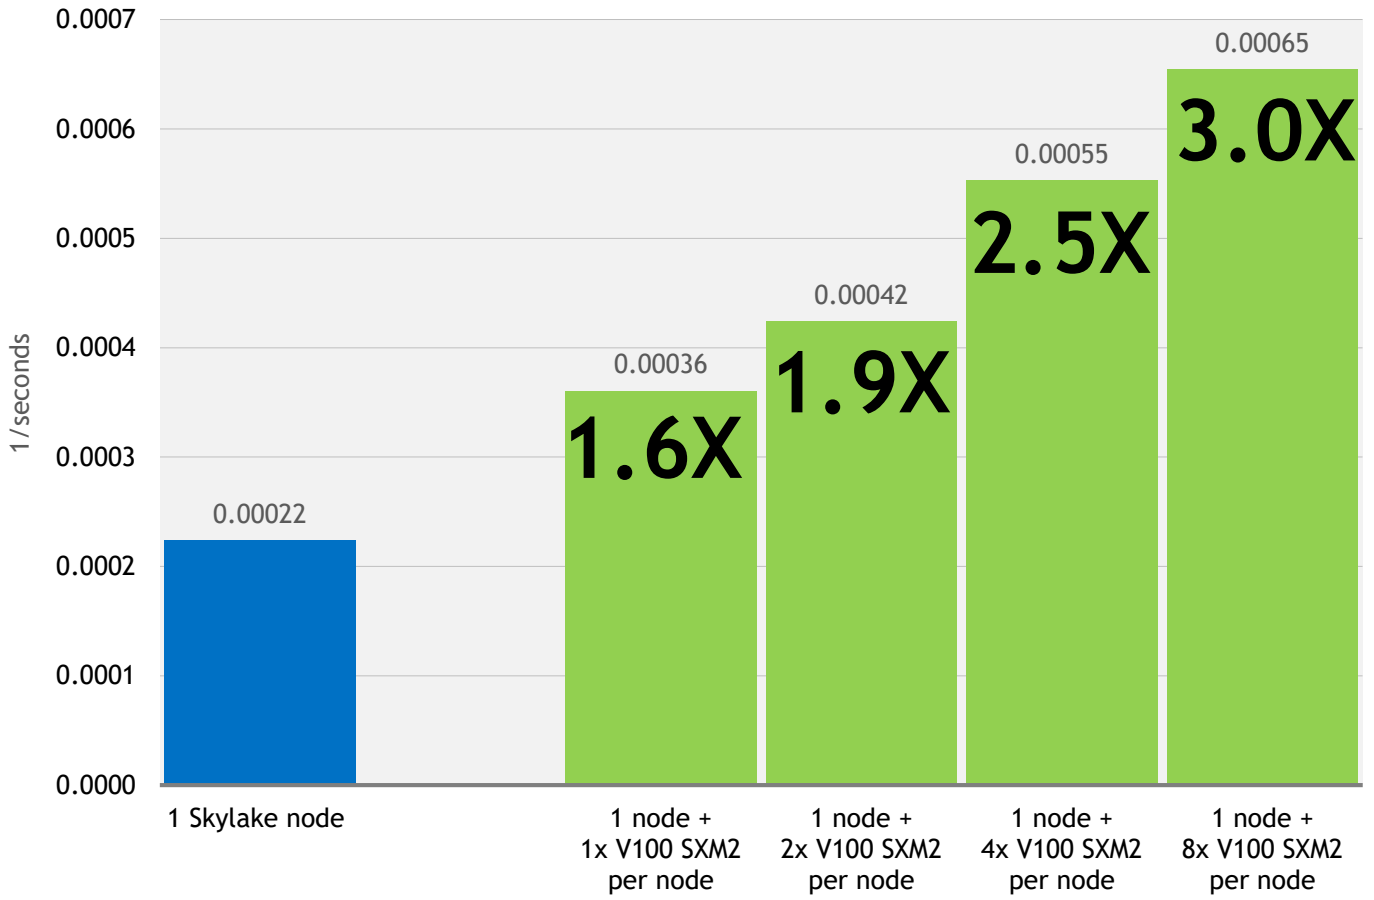

(Untuned on Volta)
Running VASP version 5.4.4

The blue node contains Dual Intel Xeon Gold 6140@2.30GHz [2.301GHz Turbo] (Skylake) CPUs

The green nodes contain Dual Intel Xeon E5-2698 v4@2.2GHz [3.6GHz Turbo] (Broadwell) CPUs + Tesla V100 SXM2 (32GB) GPUs

512 Si atoms
1282 bands
864000 Plane Waves
Algo = Normal (blocked Davidson)

VASP Si-Huge
on Platinum 8180 vs V100 PCIe (32GB)

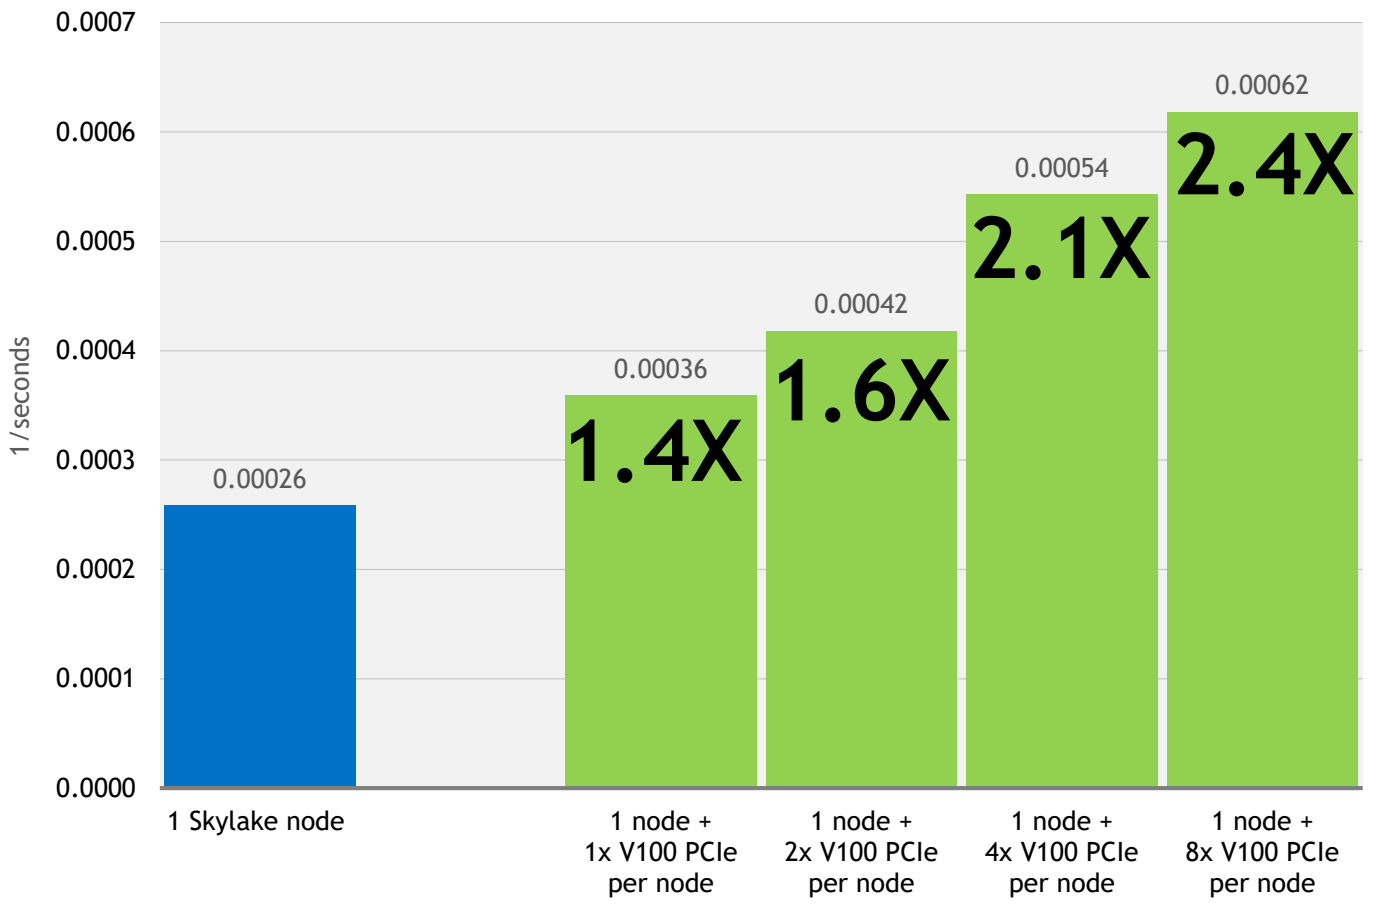

(Untuned on Volta)
Running VASP version 5.4.4

The blue node contains Dual Intel Xeon Platinum 8180@2.50GHz [2.501GHz Turbo] (Skylake) CPUs

The green nodes contain Dual Intel Xeon E5-2690 v4@2.6GHz [3.5GHz Turbo] (Broadwell) CPUs + Tesla V100 PCIe (32GB) GPUs

512 Si atoms
1282 bands
864000 Plane Waves
Algo = Normal (blocked Davidson)

VASP Si-Huge
on Platinum 8180 vs V100 SXM2 (32GB)

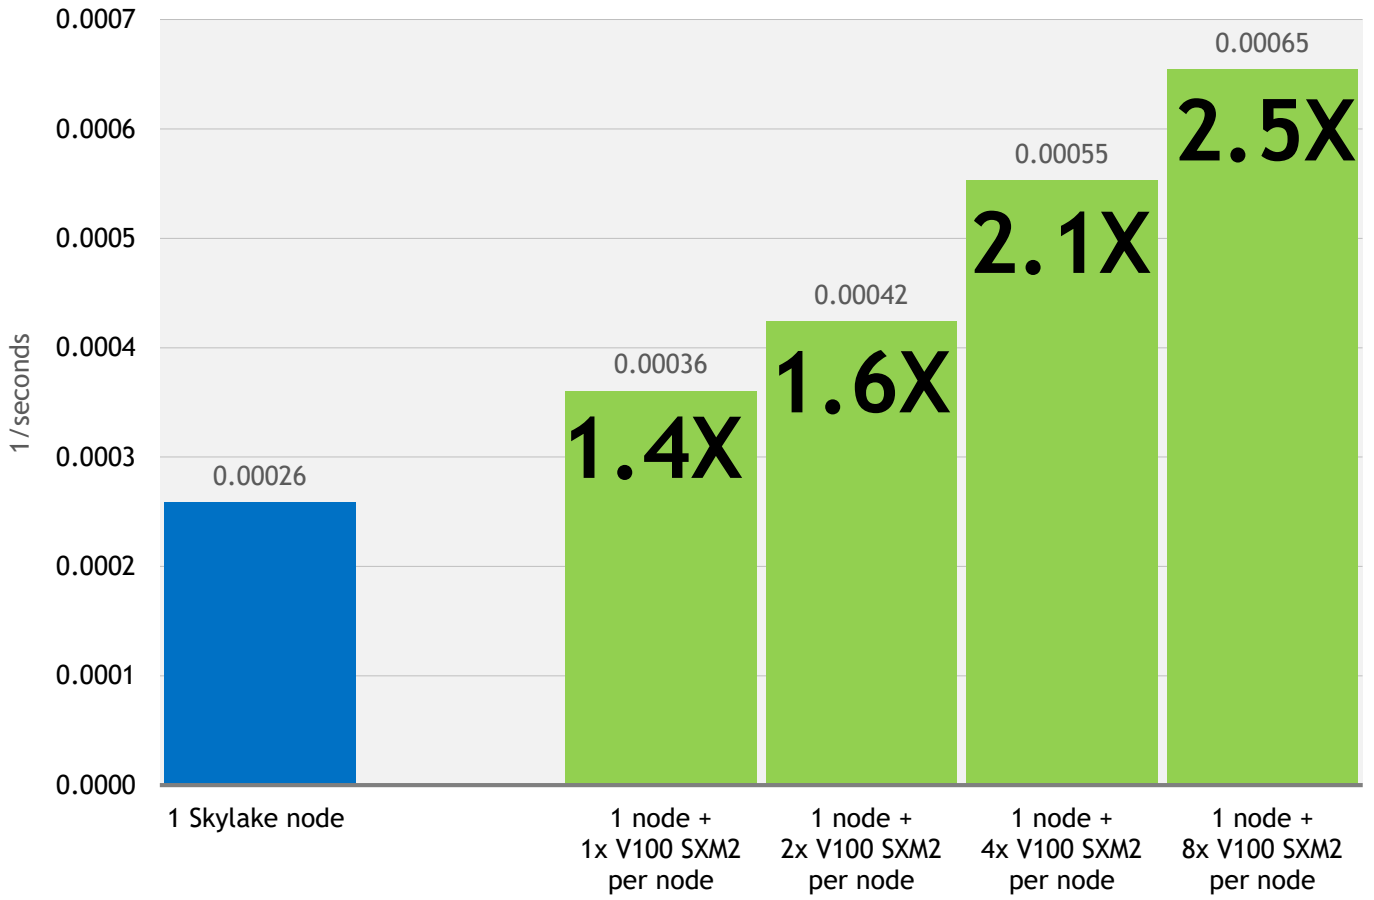

(Untuned on Volta)
Running VASP version 5.4.4

The blue node contains Dual Intel Xeon Platinum 8180@2.50GHz [2.501GHz Turbo] (Skylake) CPUs

The green nodes contain Dual Intel Xeon E5-2698 v4@2.2GHz [3.6GHz Turbo] (Broadwell) CPUs + Tesla V100 SXM2 (32GB) GPUs

512 Si atoms
1282 bands
864000 Plane Waves
Algo = Normal (blocked Davidson)

VASP B.hR105
on Gold 6140 vs V100 PCIe (32GB)

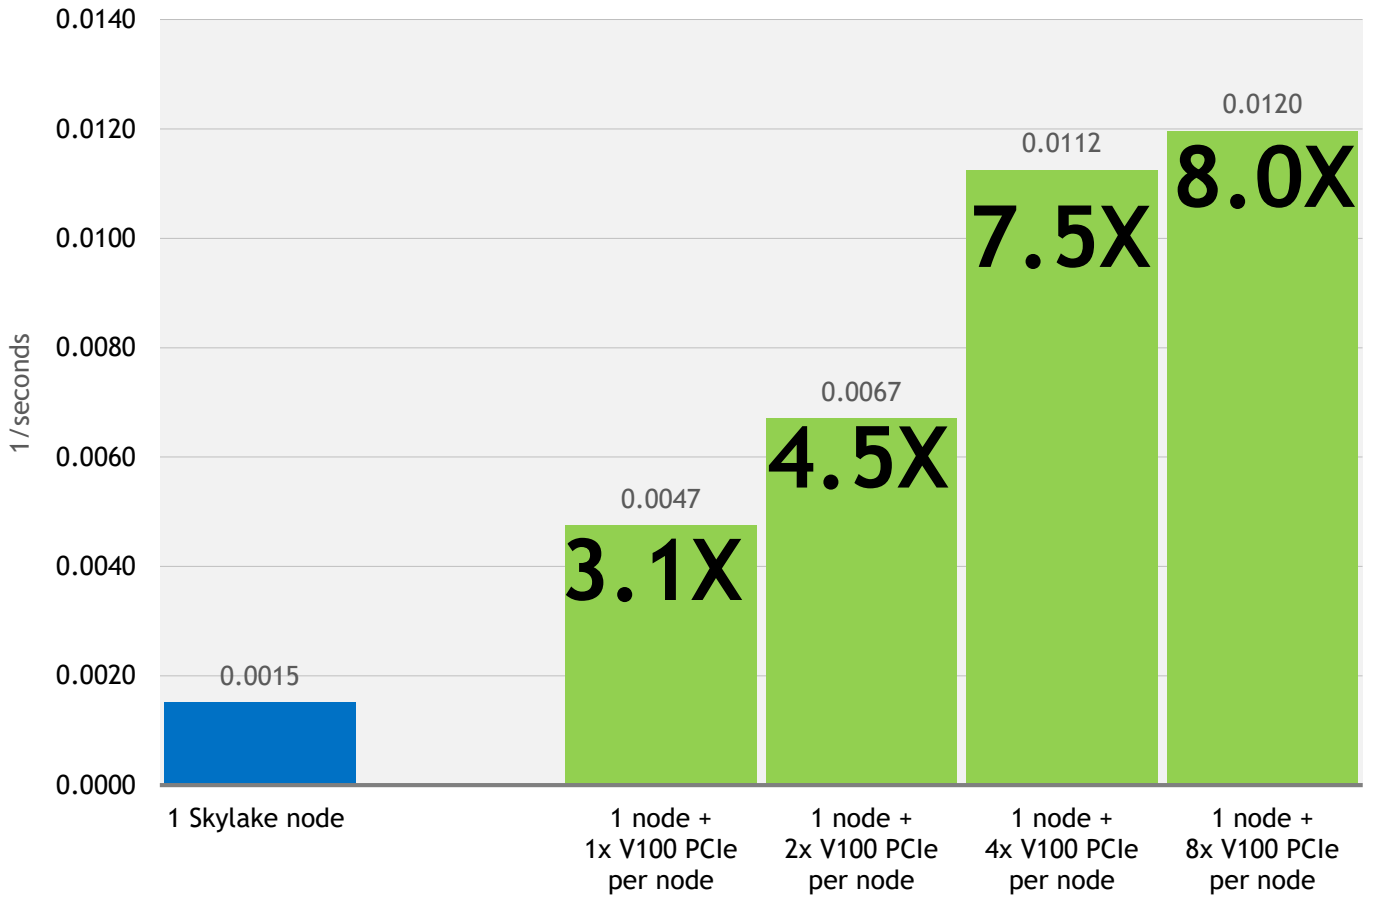

(Untuned on Volta)
Running VASP version 5.4.4

The blue node contains Dual Intel Xeon Gold 6140@2.30GHz [2.301GHz Turbo] (Skylake) CPUs

The green nodes contain Dual Intel Xeon E5-2690 v4@2.6GHz [3.5GHz Turbo] (Broadwell) CPUs + Tesla V100 PCIe (32GB) GPUs

105 Boron atoms (B-rhombohedral structure)
216 bands
110592 plane waves
Hybrid Functional with blocked Davicson (ALGO=Normal)
LHFCALC=.True. (Exact Exchange)

VASP B.hR105
on Gold 6140 vs V100 SXM2 (32GB)

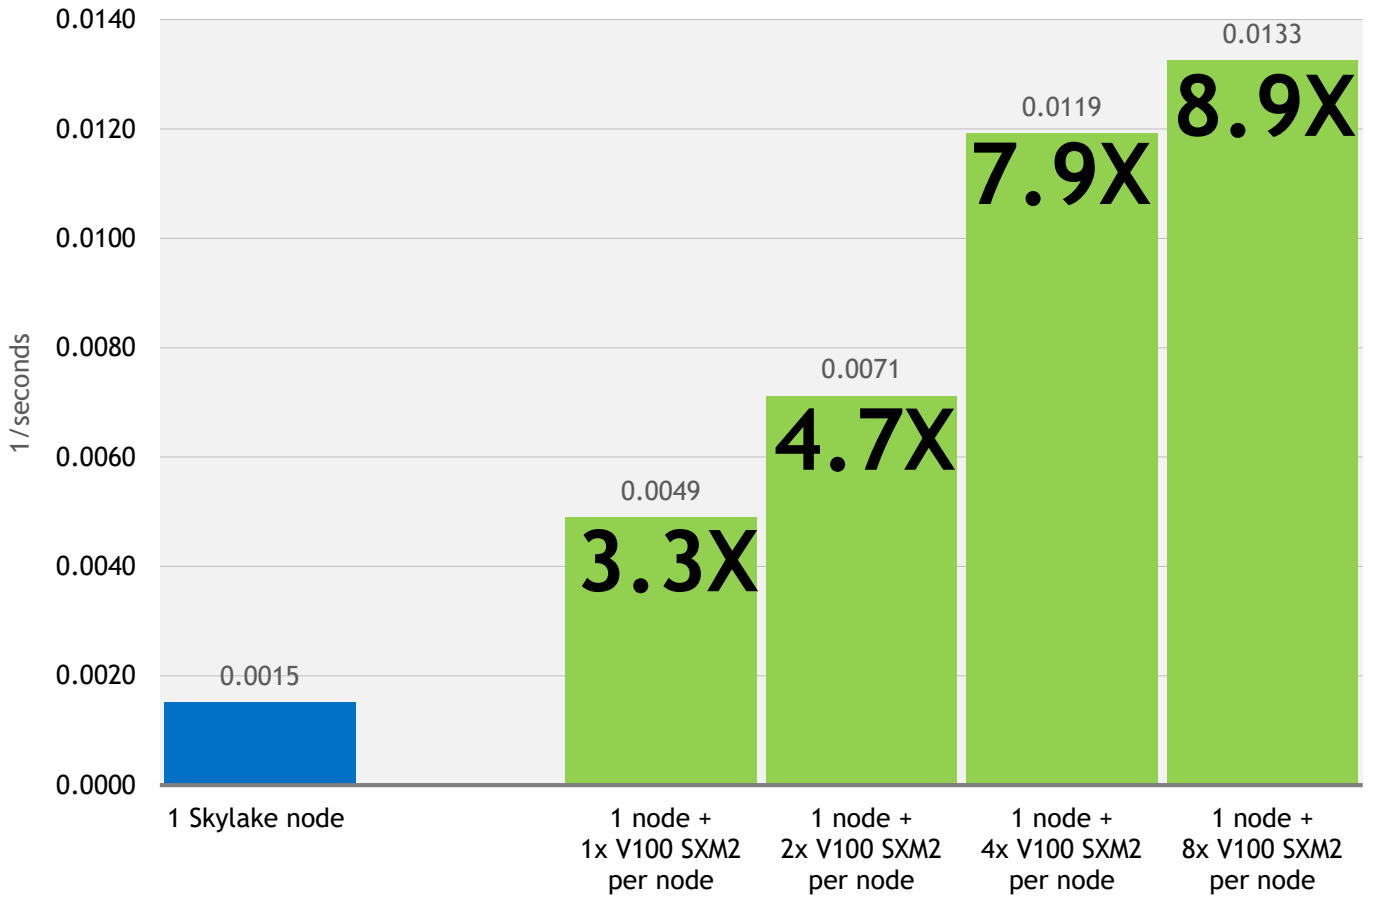

(Untuned on Volta)
Running VASP version 5.4.4

The blue node contains Dual Intel Xeon Gold 6140@2.30GHz [2.301GHz Turbo] (Skylake) CPUs

The green nodes contain Dual Intel Xeon E5-2698 v4@2.2GHz [3.6GHz Turbo] (Broadwell) CPUs + Tesla V100 SXM2 (32GB) GPUs

105 Boron atoms (B-rhombohedral structure)
216 bands
110592 plane waves
Hybrid Functional with blocked Davicson (ALGO=Normal)
LHFCALC=.True. (Exact Exchange)

VASP B.hR105
on Platinum 8180 vs V100 PCIe (32GB)

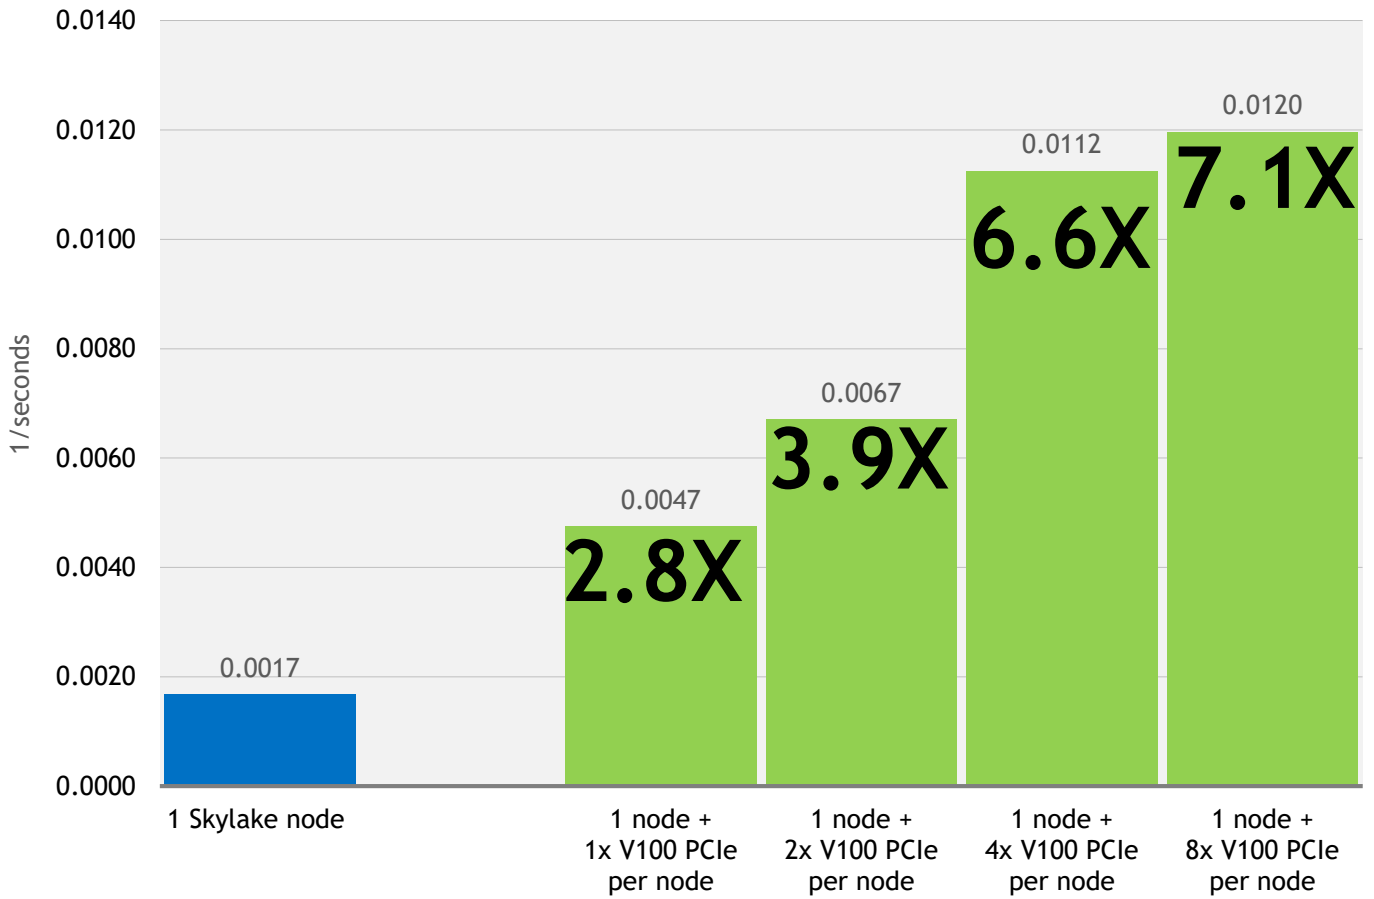

(Untuned on Volta)
Running VASP version 5.4.4

The blue node contains Dual Intel Xeon Platinum 8180@2.50GHz [2.501GHz Turbo] (Skylake) CPUs

The green nodes contain Dual Intel Xeon E5-2690 v4@2.6GHz [3.5GHz Turbo] (Broadwell) CPUs + Tesla V100 PCIe (32GB) GPUs

105 Boron atoms (B-rhombohedral structure)
216 bands
110592 plane waves
Hybrid Functional with blocked Davicson (ALGO=Normal)
LHFCALC=.True. (Exact Exchange)

VASP B.hR105
on Platinum 8180 vs V100 SXM2 (32GB)

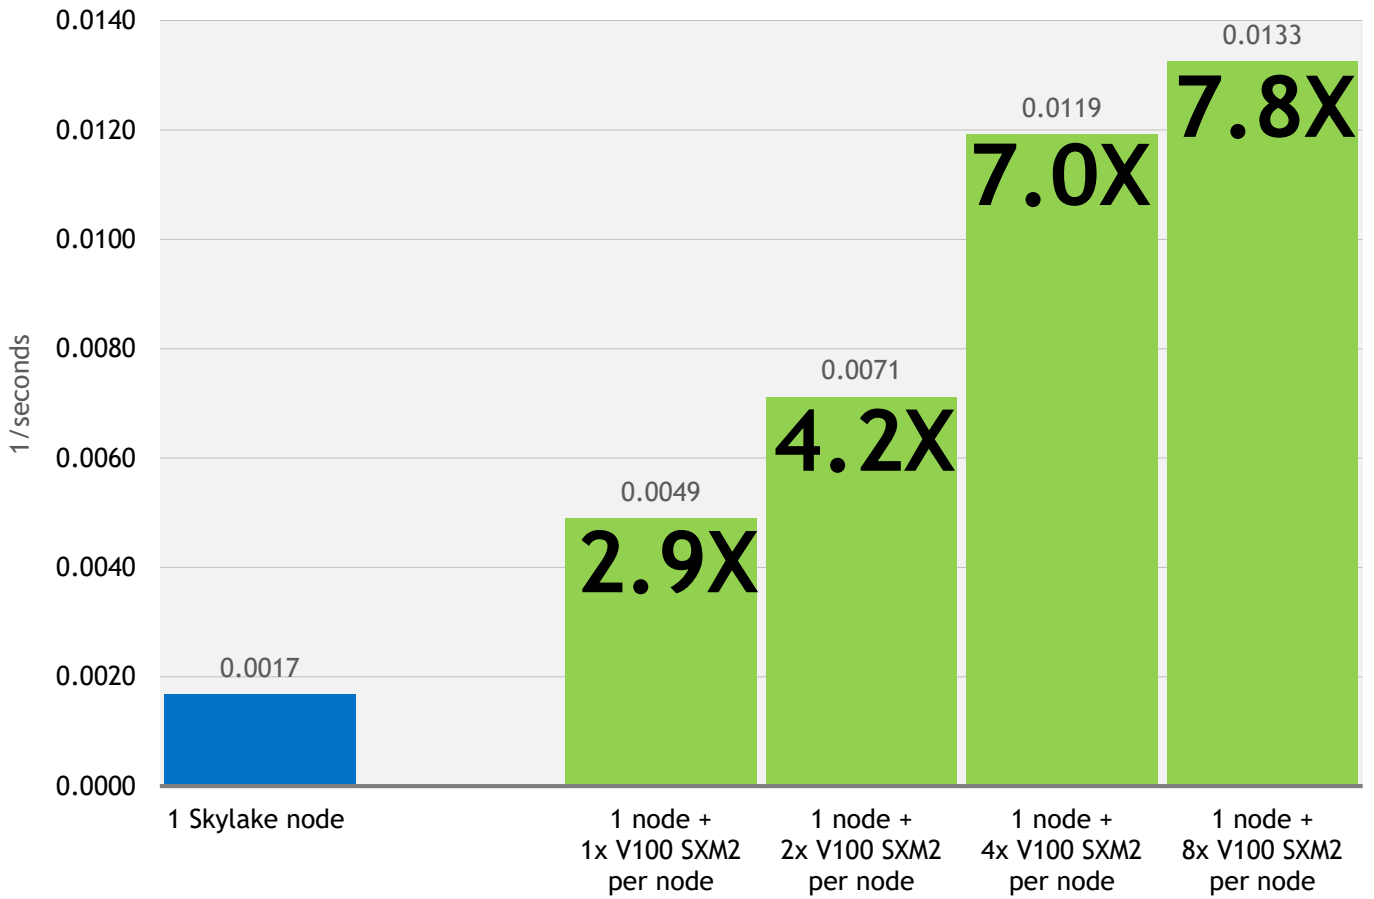

(Untuned on Volta)
Running VASP version 5.4.4

The blue node contains Dual Intel Xeon Platinum 8180@2.50GHz [2.501GHz Turbo] (Skylake) CPUs

The green nodes contain Dual Intel Xeon E5-2698 v4@2.2GHz [3.6GHz Turbo] (Broadwell) CPUs + Tesla V100 SXM2 (32GB) GPUs

105 Boron atoms (B-rhombohedral structure)
216 bands
110592 plane waves
Hybrid Functional with blocked Davicson (ALGO=Normal)
LHFCALC=.True. (Exact Exchange)

VASP 5.4.4

October 2017

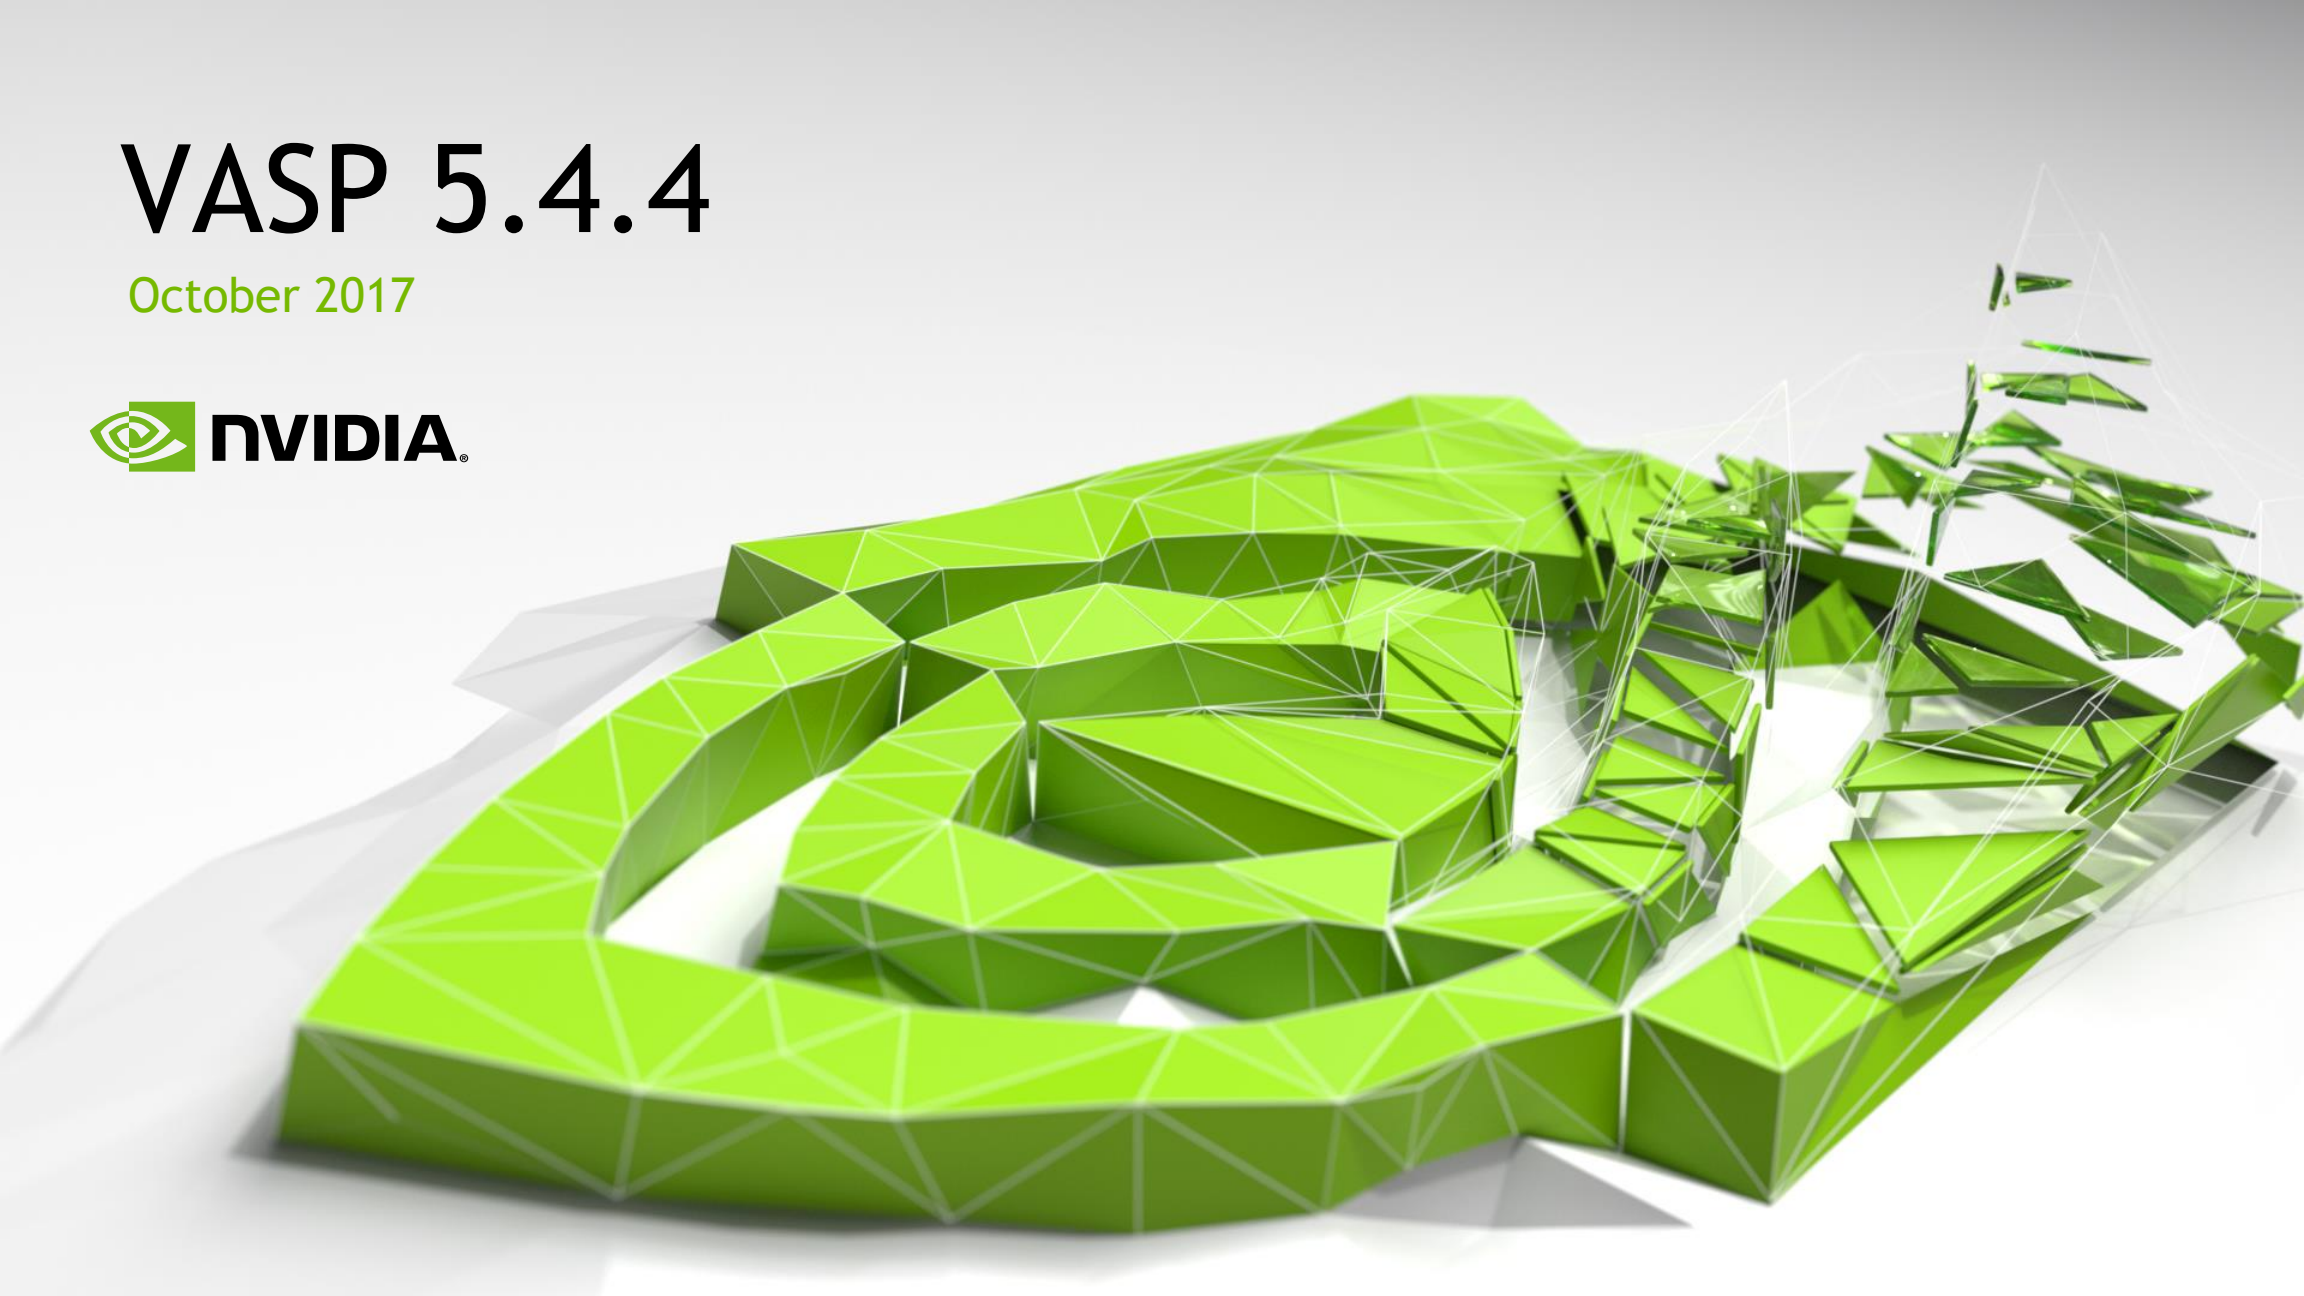

Silica IFPEN on V100s PCIe

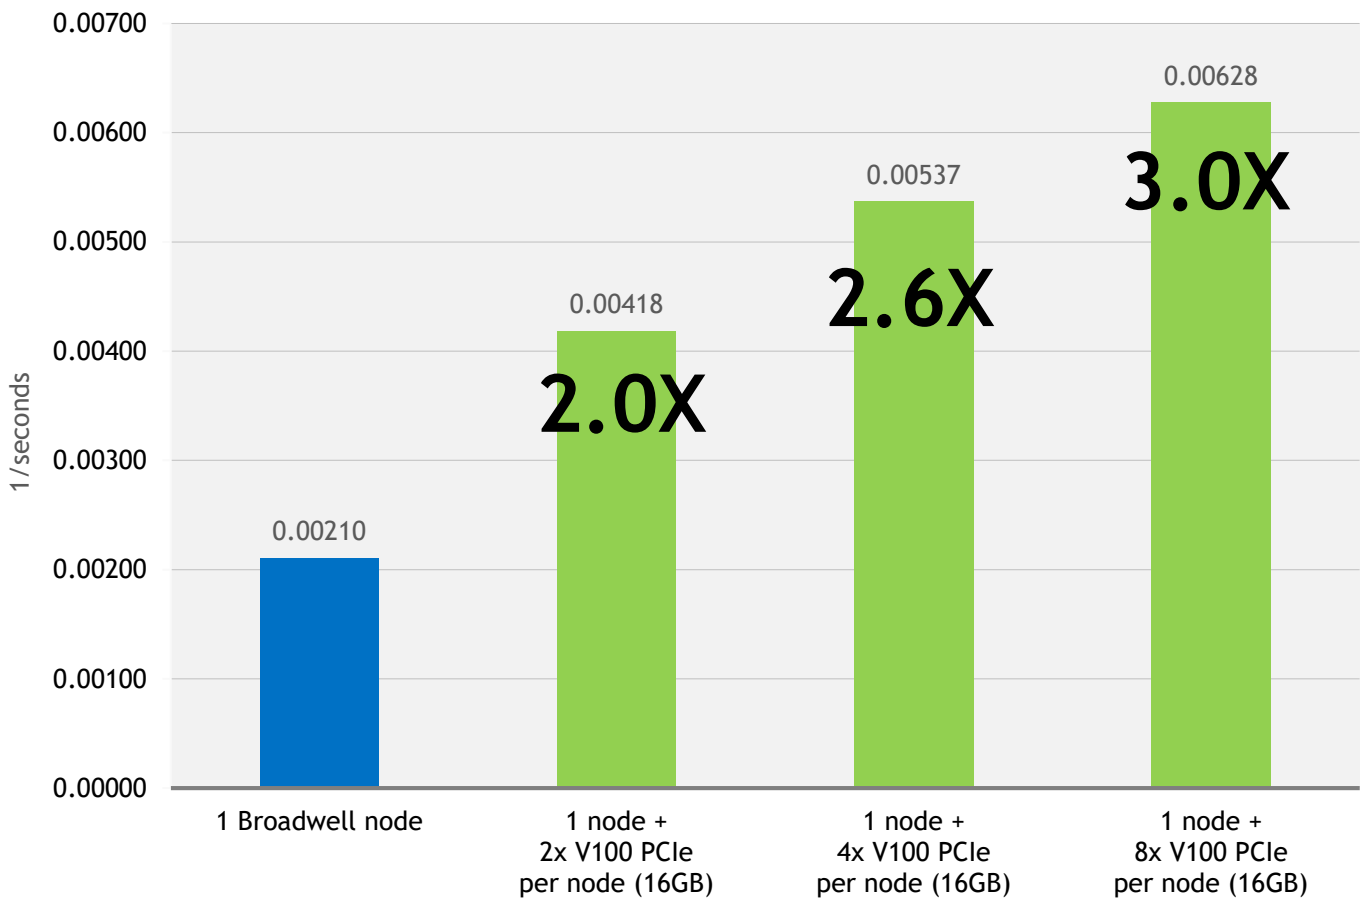

(Untuned on Volta)
Running VASP version 5.4.4

The blue node contains Dual Intel Xeon E5-2690 v4@2.6GHz [3.5GHz Turbo] (Broadwell) CPUs

The green nodes contain Dual Intel Xeon E5-2690 v4@2.6GHz [3.5GHz Turbo] (Broadwell) CPUs + Tesla V100 PCIe (16GB) GPUs

240 ions, cristobalite (high) bulk
720 bands
? plane waves
ALGO = Very Fast (RMM-DIIS)

Silica IFPEN on V100s SXM2

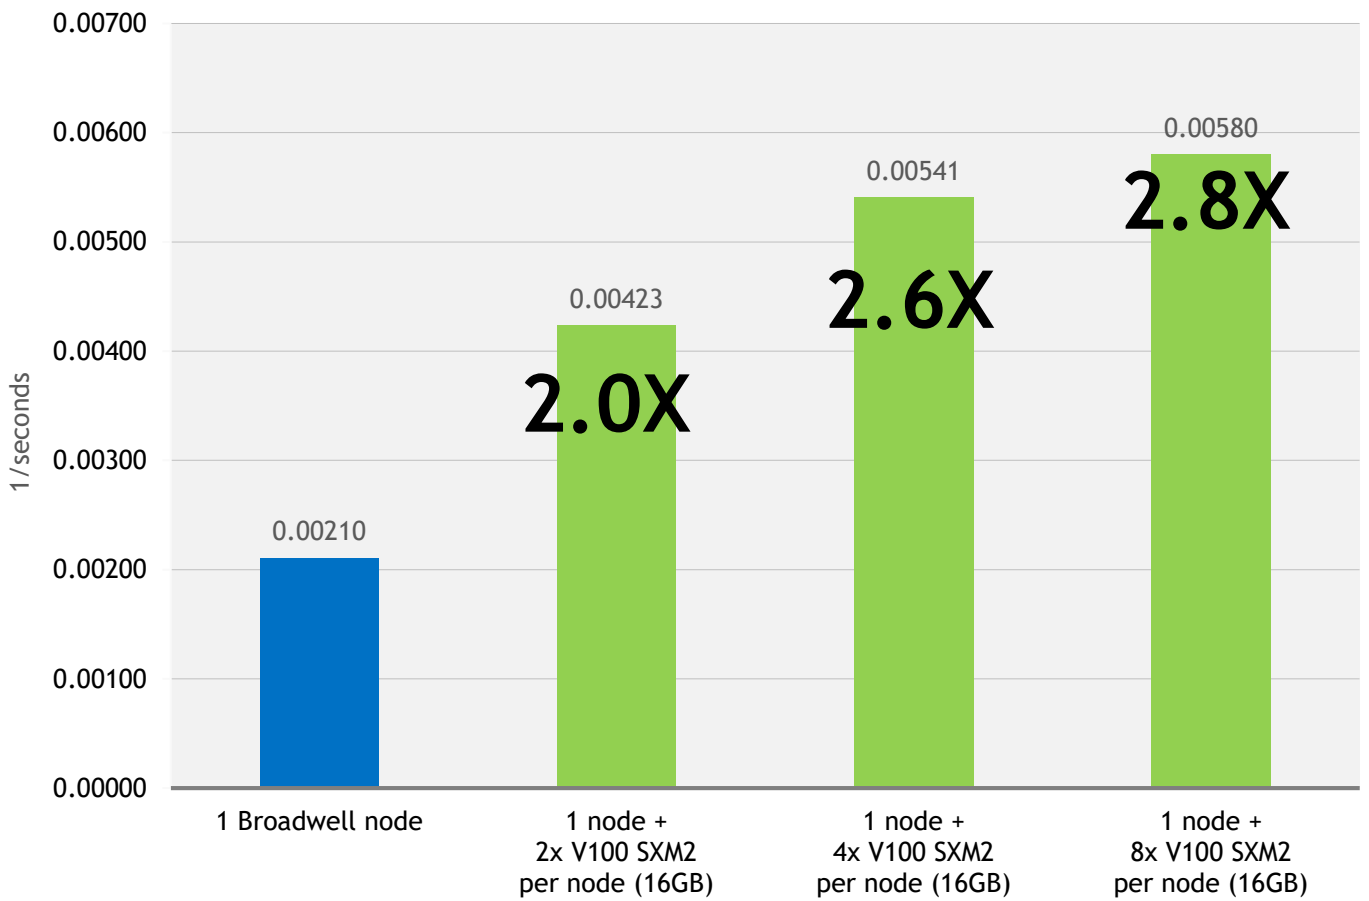

(Untuned on Volta)
Running VASP version 5.4.4

The blue node contains Dual Intel Xeon E5-2698 v4@2.2GHz [3.6GHz Turbo] (Broadwell) CPUs

The green nodes contain Dual Intel Xeon E5-2698 v4@2.2GHz [3.6GHz Turbo] (Broadwell) CPUs + Tesla V100 SXM2 (16GB) GPUs

240 ions, cristobalite (high) bulk
720 bands
? plane waves
ALGO = Very Fast (RMM-DIIS)

Si-Huge on V100s PCIe

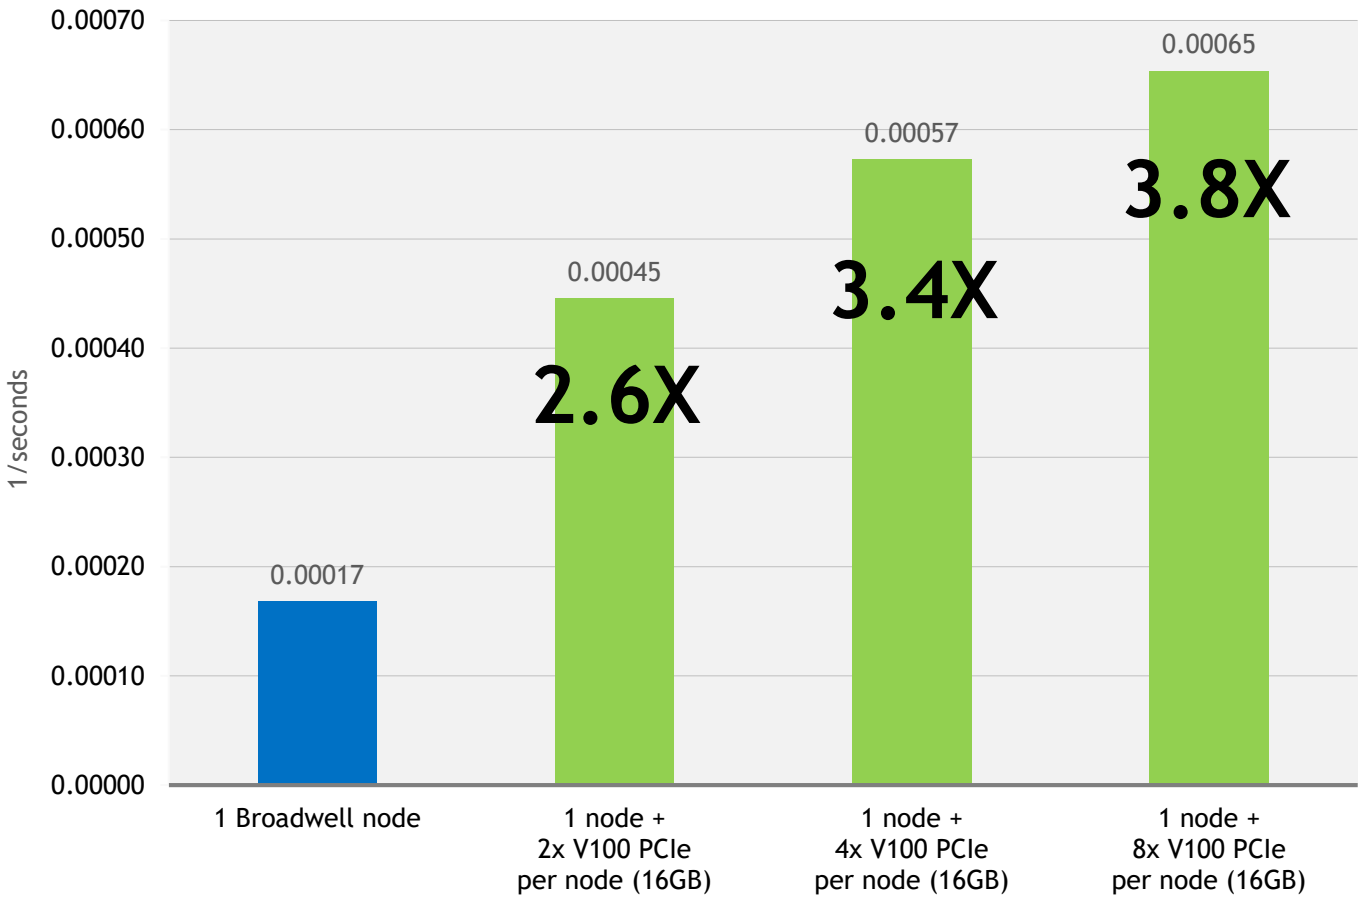

(Untuned on Volta)
Running VASP version 5.4.4

The blue node contains Dual Intel Xeon E5-2690 v4@2.6GHz [3.5GHz Turbo] (Broadwell) CPUs

The green nodes contain Dual Intel Xeon E5-2690 v4@2.6GHz [3.5GHz Turbo] (Broadwell) CPUs + Tesla V100 PCIe (16GB) GPUs

512 Si atoms
1282 bands
864000 Plane Waves
Algo = Normal (blocked Davidson)

Si-Huge on V100s SXM2

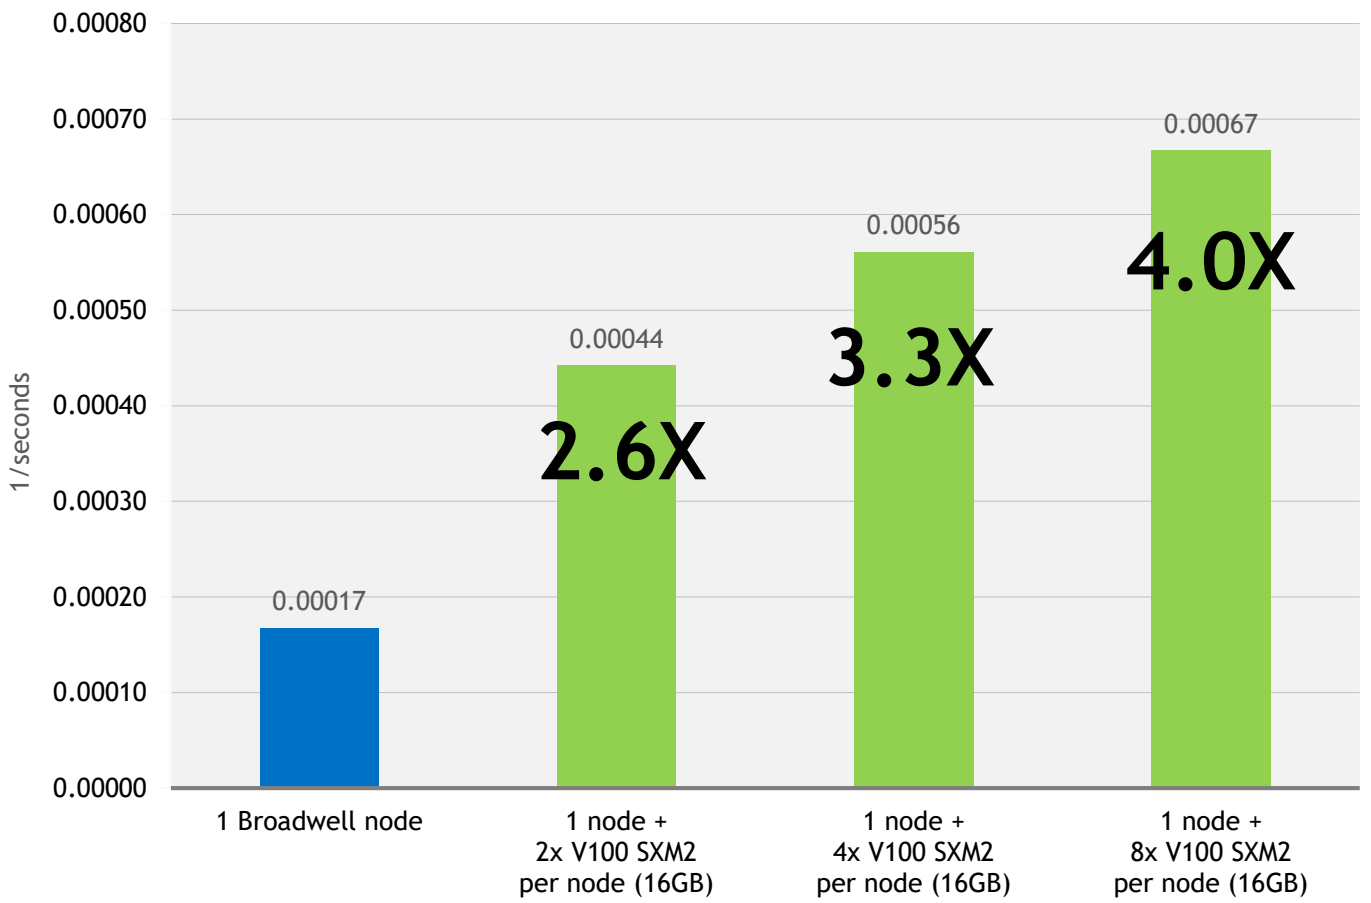

(Untuned on Volta)
Running VASP version 5.4.4

The blue node contains Dual Intel Xeon E5-2698 v4@2.2GHz [3.6GHz Turbo] (Broadwell) CPUs

The green nodes contain Dual Intel Xeon E5-2698 v4@2.2GHz [3.6GHz Turbo] (Broadwell) CPUs + Tesla V100 SXM2 (16GB) GPUs

512 Si atoms
1282 bands
864000 Plane Waves
Algo = Normal (blocked Davidson)

SupportedSystems on V100s PCIe

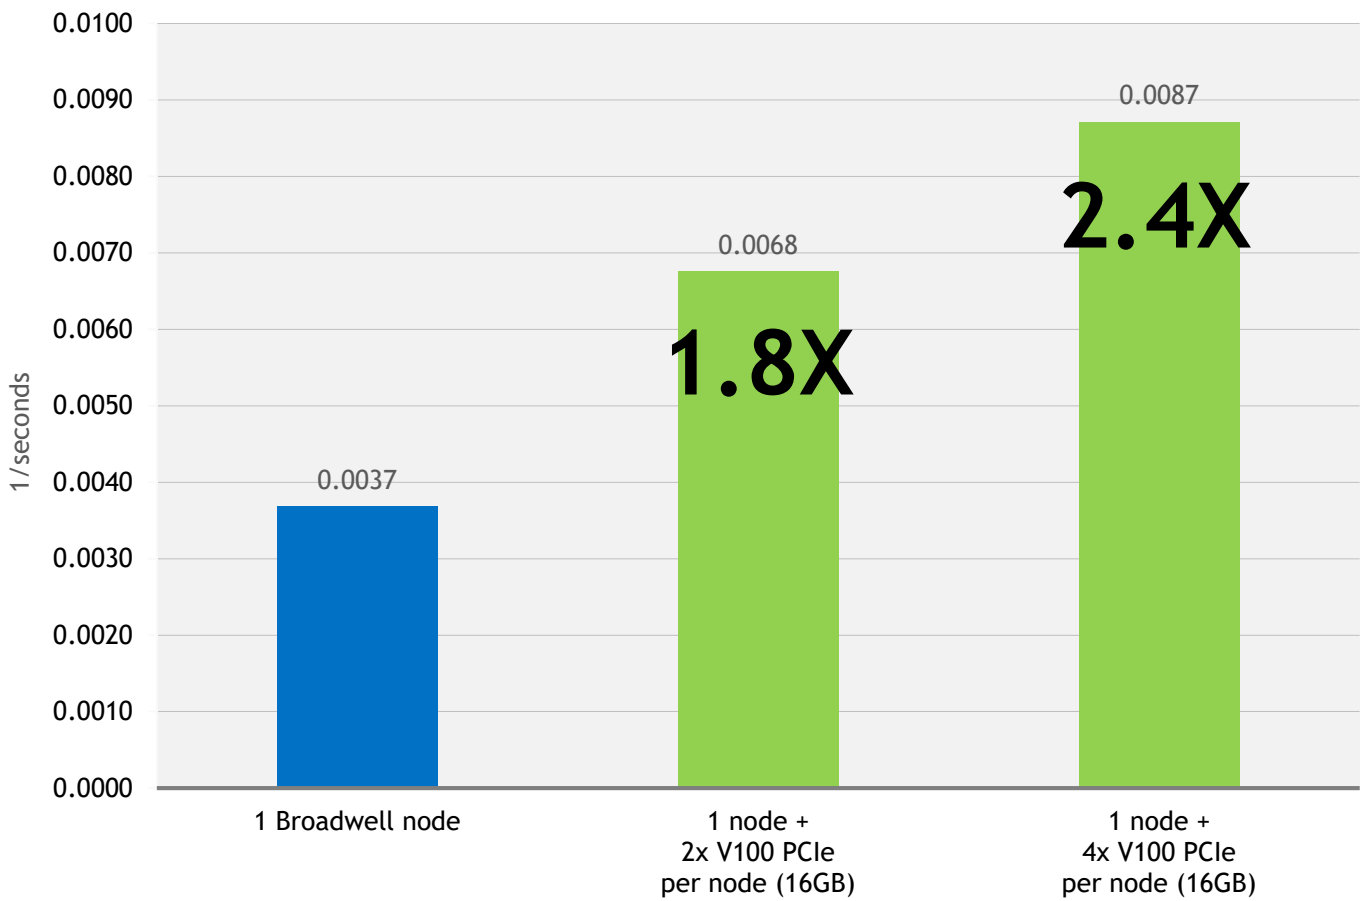

(Untuned on Volta)
Running VASP version 5.4.4

The blue node contains Dual Intel Xeon E5-2690 v4@2.6GHz [3.5GHz Turbo] (Broadwell) CPUs

The green nodes contain Dual Intel Xeon E5-2690 v4@2.6GHz [3.5GHz Turbo] (Broadwell) CPUs + Tesla V100 PCIe (16GB) GPUs

267 ions
788 bands
762048 plane waves
ALGO = Fast (Davidson + RMM-DIIS)

SupportedSystems on V100s SXM2

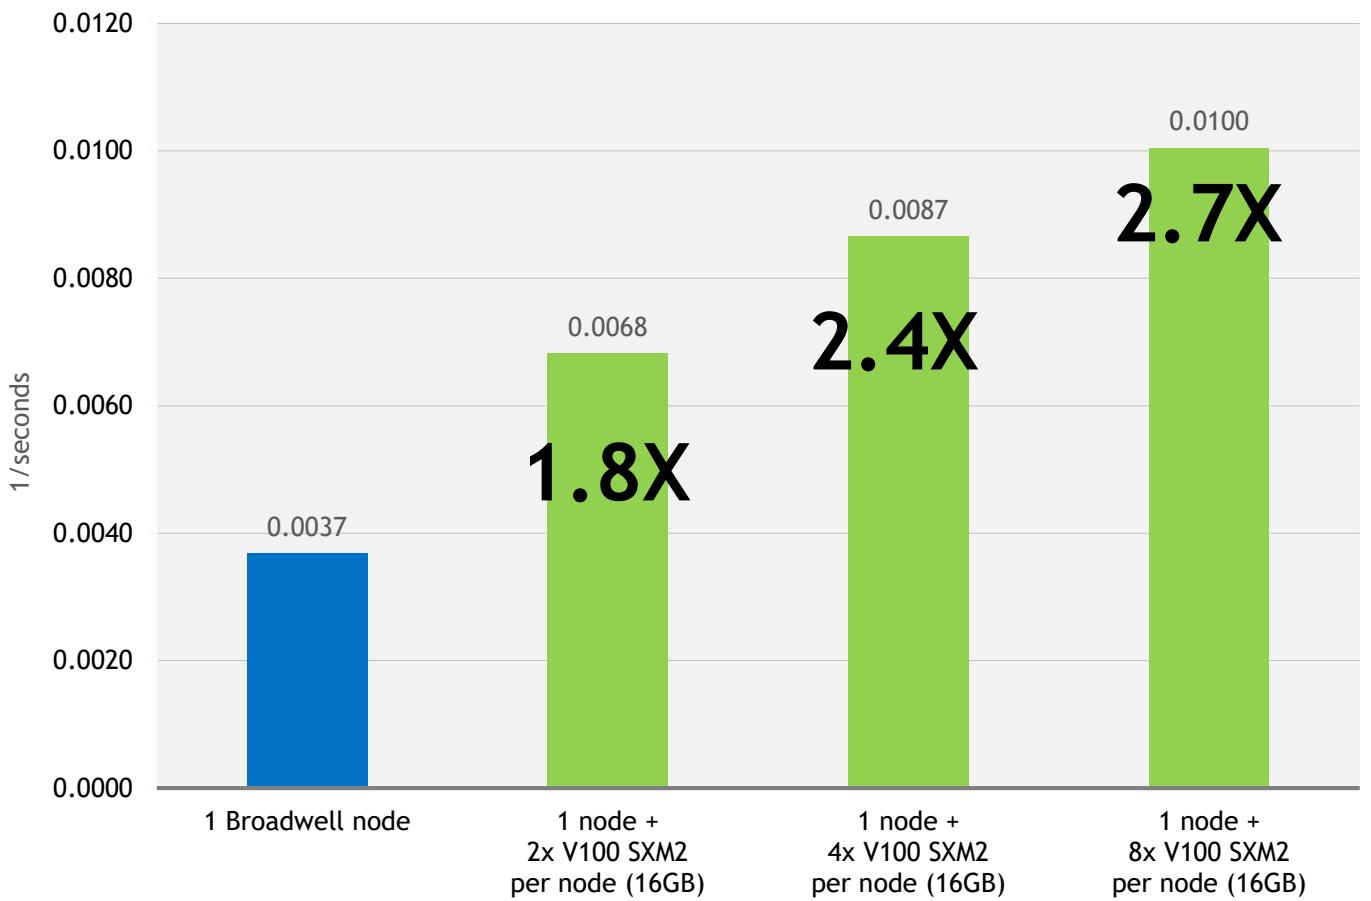

(Untuned on Volta)
Running VASP version 5.4.4

The blue node contains Dual Intel Xeon E5-2698 v4@2.2GHz [3.6GHz Turbo] (Broadwell) CPUs

The green nodes contain Dual Intel Xeon E5-2698 v4@2.2GHz [3.6GHz Turbo] (Broadwell) CPUs + Tesla V100 SXM2 (16GB) GPUs

267 ions
788 bands
762048 plane waves
ALGO = Fast (Davidson + RMM-DIIS)

NiAl-MD on V100s PCIe

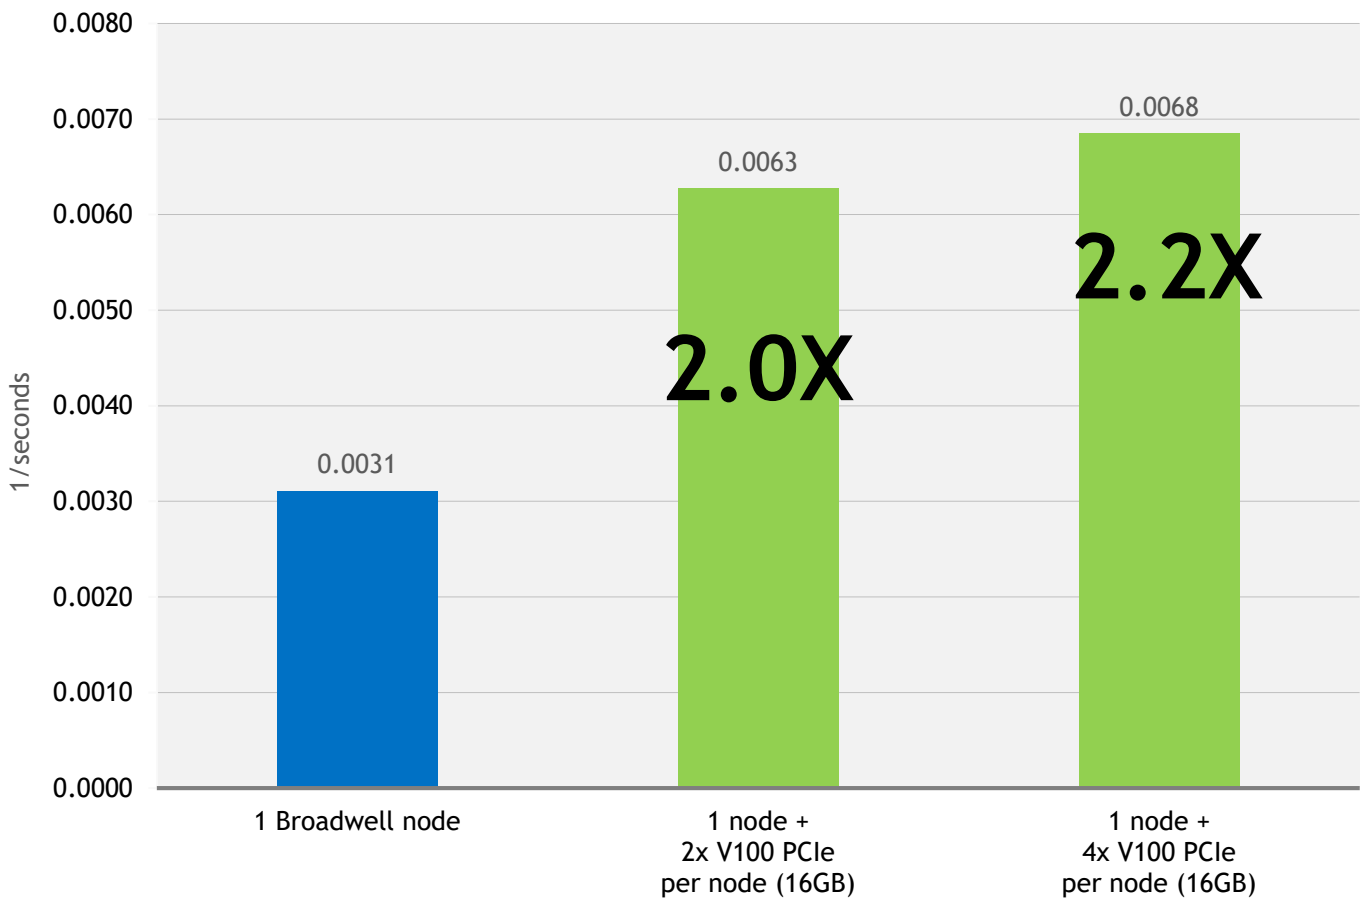

(Untuned on Volta)
Running VASP version 5.4.4

The blue node contains Dual Intel Xeon E5-2690 v4@2.6GHz [3.5GHz Turbo] (Broadwell) CPUs

The green nodes contain Dual Intel Xeon E5-2690 v4@2.6GHz [3.5GHz Turbo] (Broadwell) CPUs + Tesla V100 PCIe (16GB) GPUs

500 ions
3200 bands
729000 plane waves
ALGO = Fast (Davidson + RMM-DIIS)

NiAl-MD on V100s SXM2

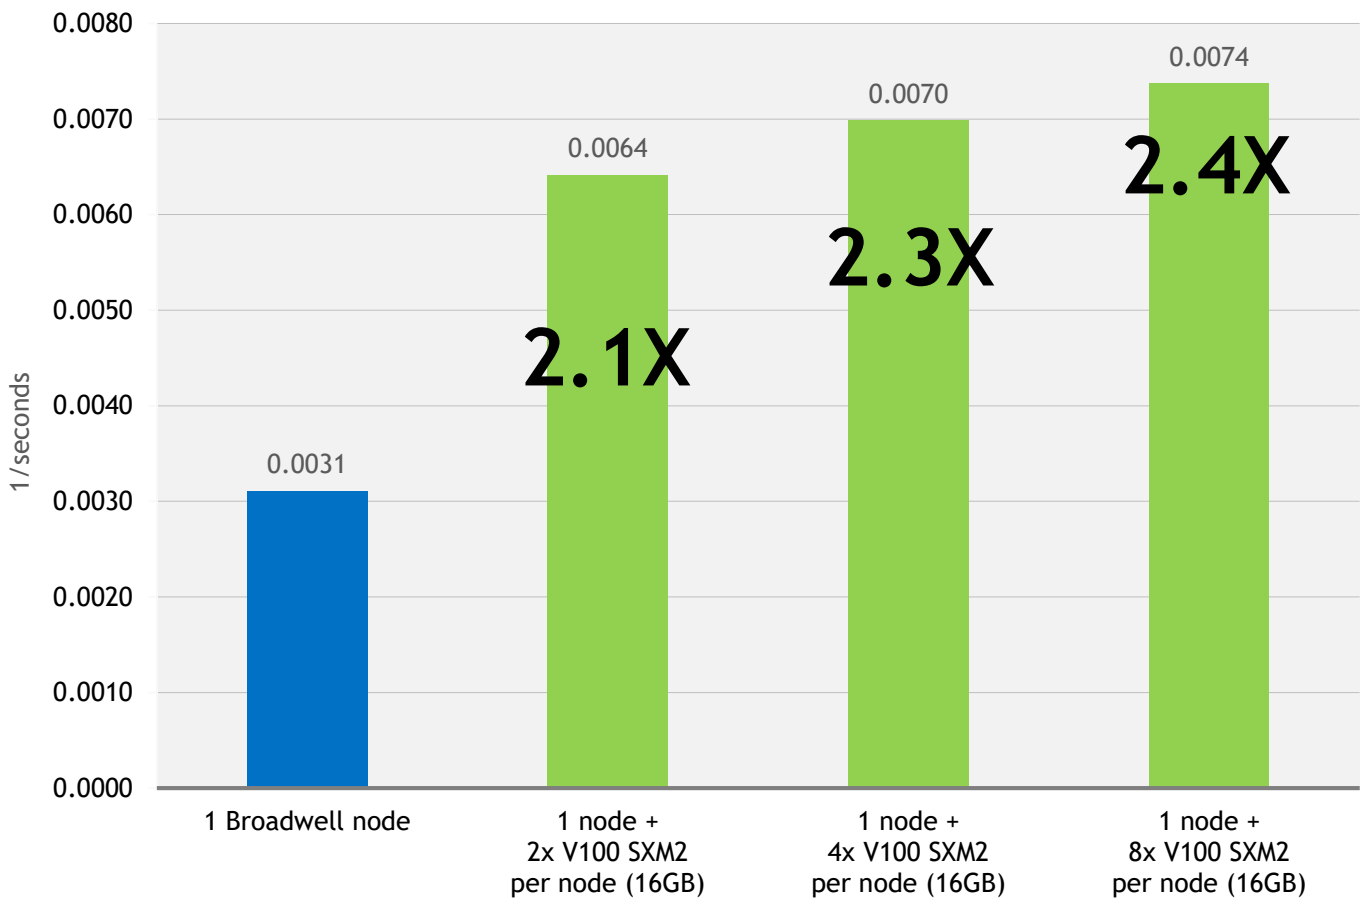

(Untuned on Volta)
Running VASP version 5.4.4

The blue node contains Dual Intel Xeon E5-2698 v4@2.2GHz [3.6GHz Turbo] (Broadwell) CPUs

The green nodes contain Dual Intel Xeon E5-2698 v4@2.2GHz [3.6GHz Turbo] (Broadwell) CPUs + Tesla V100 SXM2 (16GB) GPUs

500 ions
3200 bands
729000 plane waves
ALGO = Fast (Davidson + RMM-DIIS)

B.hR105 on V100s PCIe

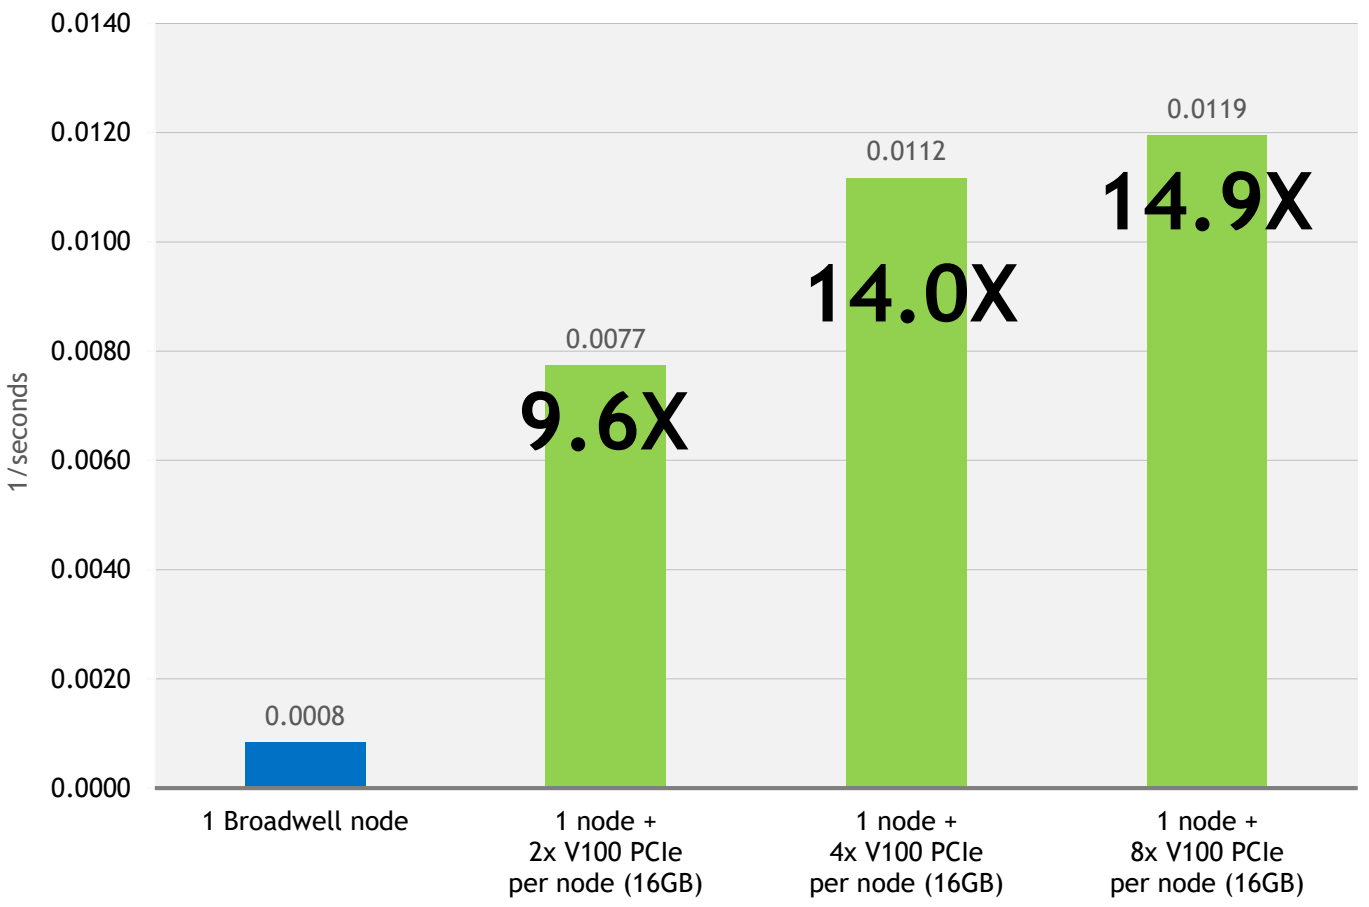

(Untuned on Volta)
Running VASP version 5.4.4

The blue node contains Dual Intel Xeon E5-2690 v4@2.6GHz [3.5GHz Turbo] (Broadwell) CPUs

The green nodes contain Dual Intel Xeon E5-2690 v4@2.6GHz [3.5GHz Turbo] (Broadwell) CPUs + Tesla V100 PCIe (16GB) GPUs

105 Boron atoms (B-rhombohedral structure)
216 bands
110592 plane waves
Hybrid Functional with blocked Davicson (ALGO=Normal)
LHFCALC=.True. (Exact Exchange)

B.hR105 on V100s SXM2

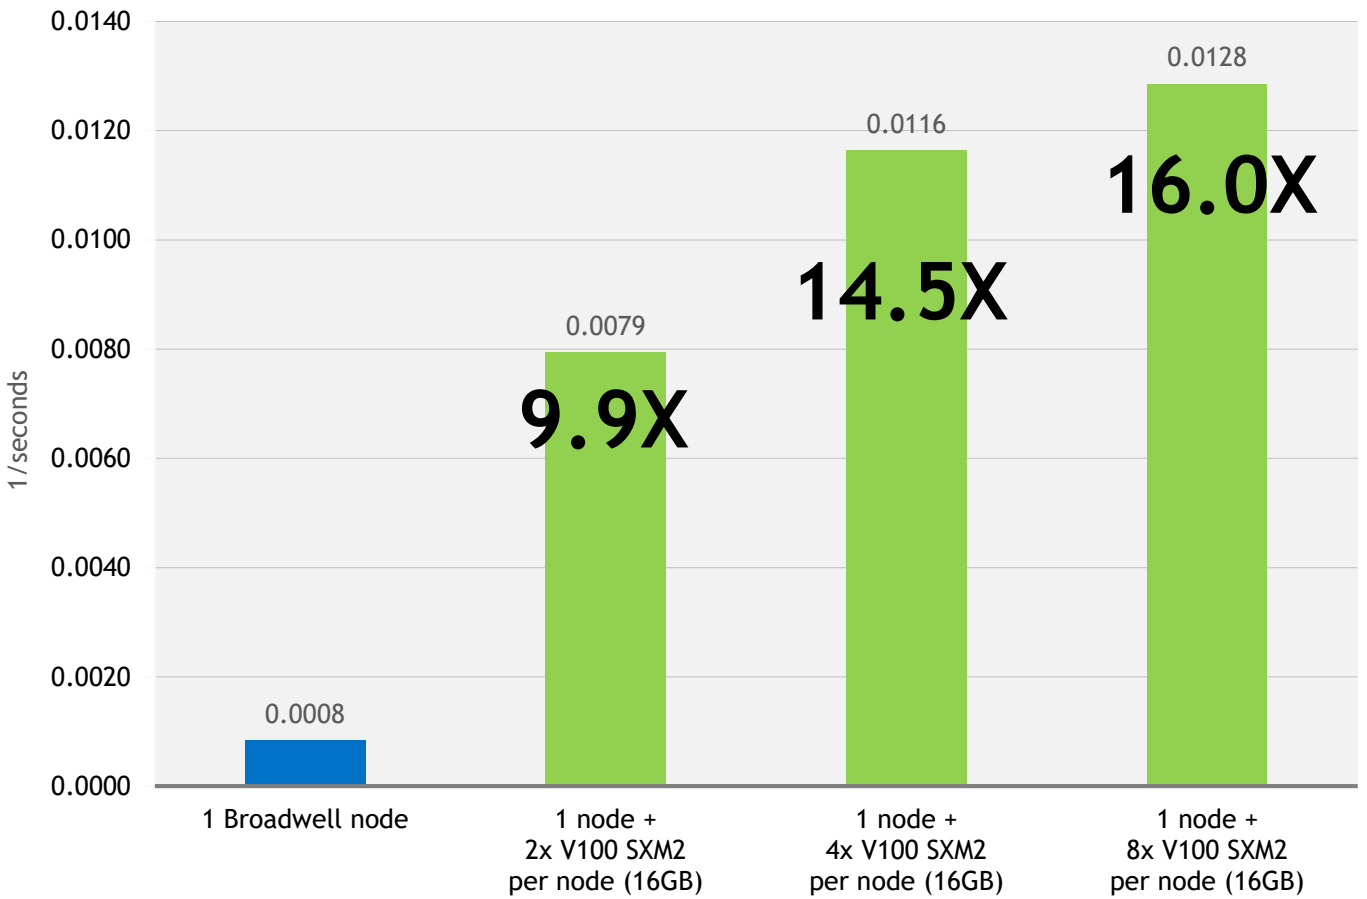

(Untuned on Volta)
Running VASP version 5.4.4

The blue node contains Dual Intel Xeon E5-2698 v4@2.2GHz [3.6GHz Turbo] (Broadwell) CPUs

The green nodes contain Dual Intel Xeon E5-2698 v4@2.2GHz [3.6GHz Turbo] (Broadwell) CPUs + Tesla V100 SXM2 (16GB) GPUs

105 Boron atoms (B-rhombohedral structure)
216 bands
110592 plane waves
Hybrid Functional with blocked Davicson (ALGO=Normal)
LHFCALC=.True. (Exact Exchange)

B.aP107 on V100s PCIe

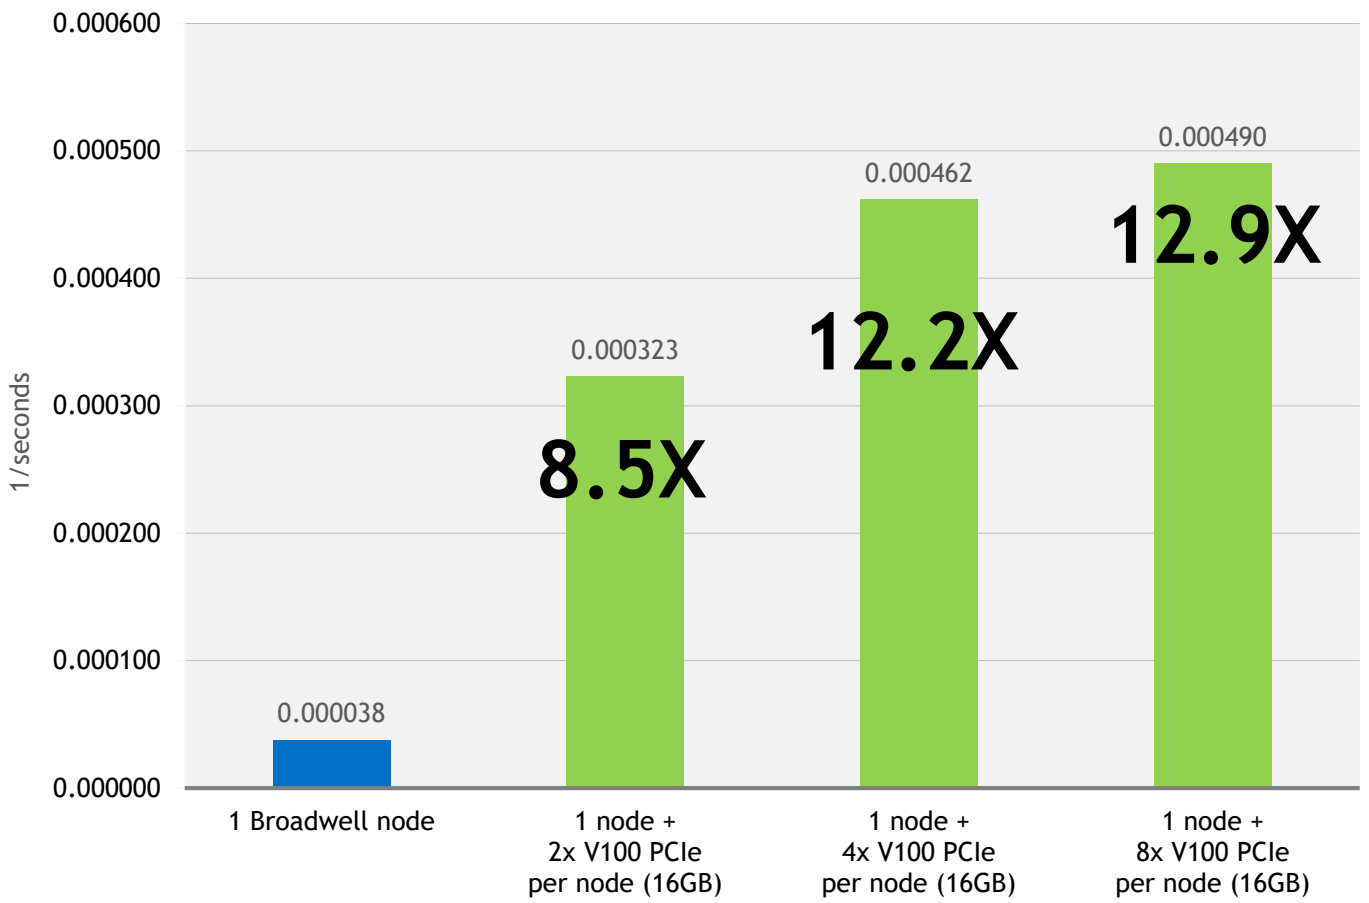

(Untuned on Volta)
Running VASP version 5.4.4

The blue node contains Dual Intel Xeon E5-2690 v4@2.6GHz [3.5GHz Turbo] (Broadwell) CPUs

The green nodes contain Dual Intel Xeon E5-2690 v4@2.6GHz [3.5GHz Turbo] (Broadwell) CPUs + Tesla V100 PCIe (16GB) GPUs

107 Boron atoms (symmetry broken 107-atom B' variant)
216 bands
110592 plane waves
Hybrid functional calculation (exact exchange) with blocked Davidson. No KPoint parallelization.
Hybrid Functional with blocked Davidson (ALGO=Normal)
LHFCALC=.True. (Exact Exchange)

B.aP107 on V100s SXM2

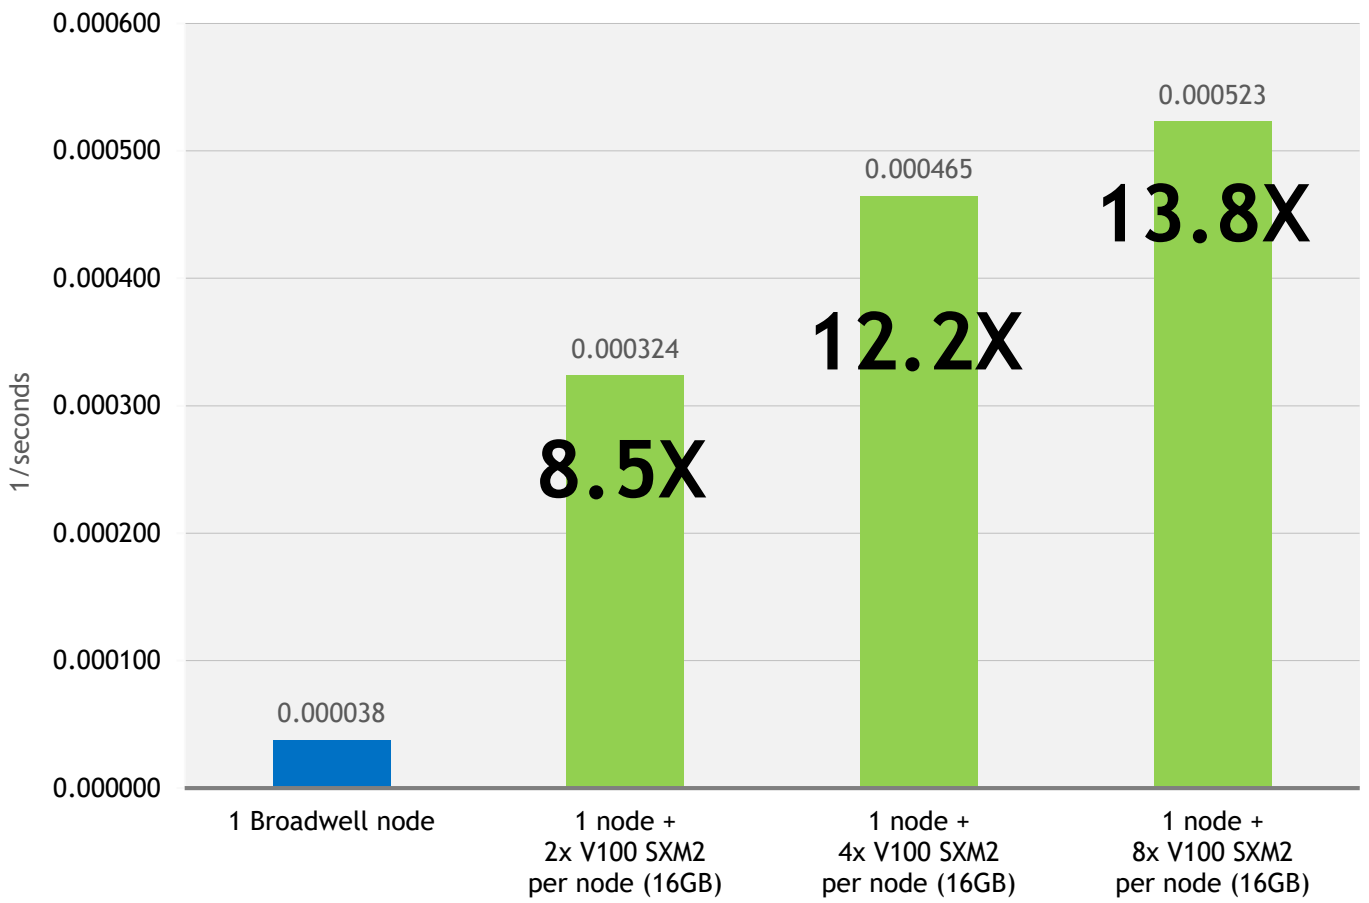

(Untuned on Volta)
Running VASP version 5.4.4

The blue node contains Dual Intel Xeon E5-2698 v4@2.2GHz [3.6GHz Turbo] (Broadwell) CPUs

The green nodes contain Dual Intel Xeon E5-2698 v4@2.2GHz [3.6GHz Turbo] (Broadwell) CPUs + Tesla V100 SXM2 (16GB) GPUs

107 Boron atoms (symmetry broken 107-atom B' variant)
216 bands
110592 plane waves
Hybrid functional calculation (exact exchange) with blocked Davidson. No KPoint parallelization.
Hybrid Functional with blocked Davidson (ALGO=Normal)
LHFCALC=.True. (Exact Exchange)

Quantum Chemistry (QC) on GPUs

March 2018

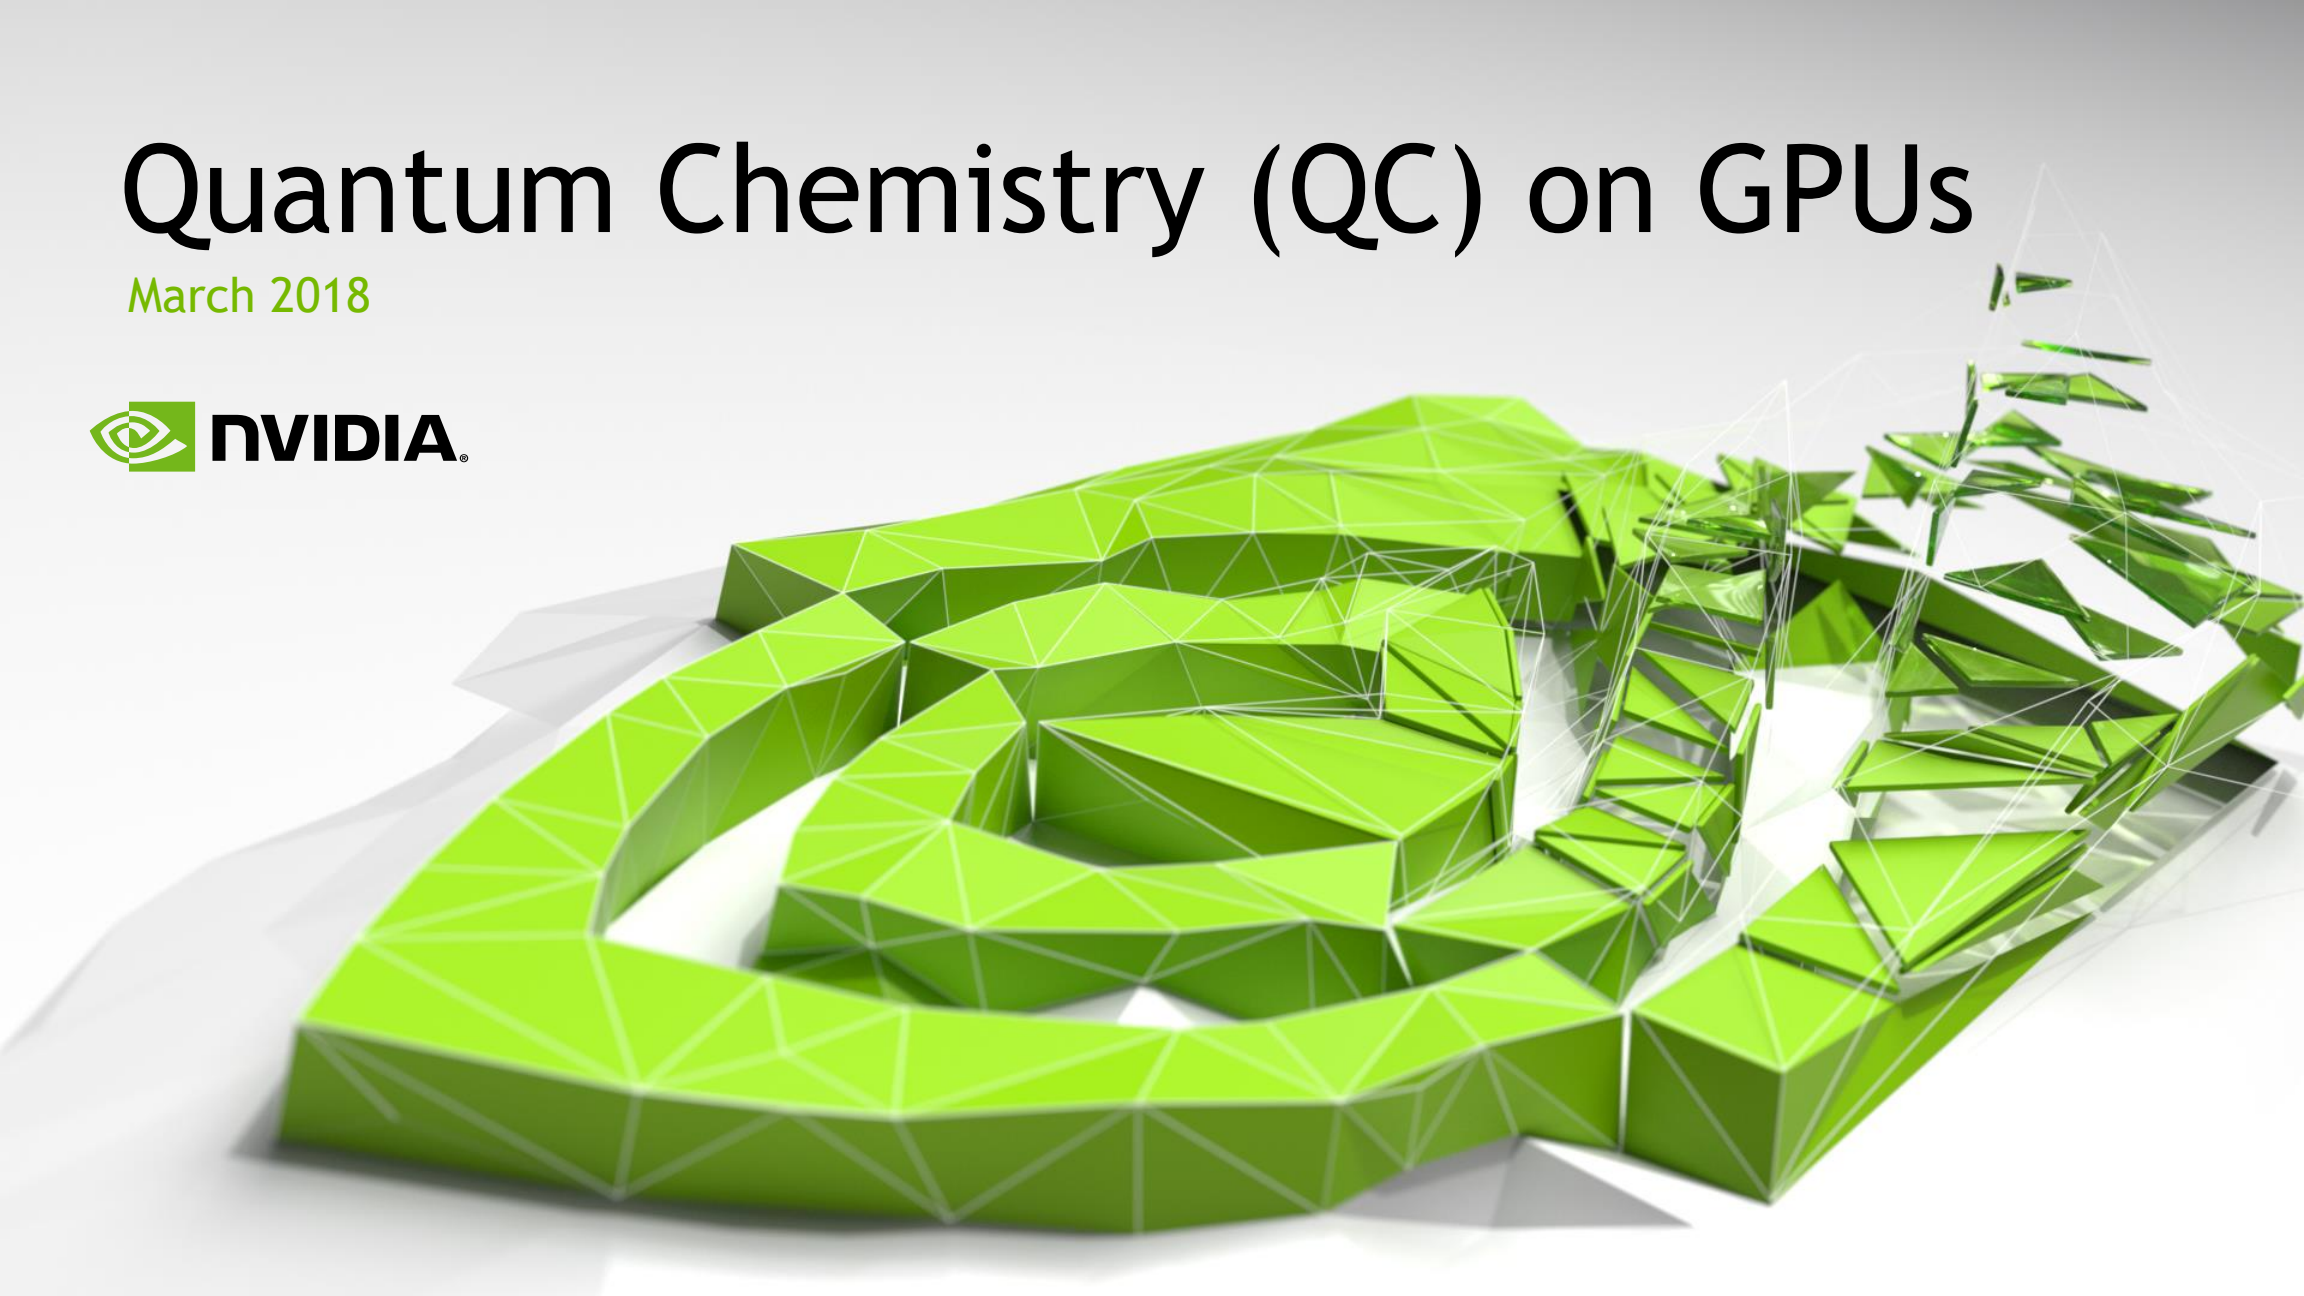

GPU-Accelerated Molecular Dynamics Apps

Green Lettering Indicates Performance Slides Included

- ACEMD
- AMBER
- CHARMM
- DESMOND
- ESPResSO
- Folding@Home
- GENESIS
- GPUGrid.net
- GROMACS
- HALMD
- HOOMD-Blue
- HTMD
- LAMMPS
- mdcore
- MELD
- NAMD
- OpenMM
- PolyFTS

GPU Perf compared against dual multi-core x86 CPU socket.