Molecular Dynamics (MD) on GPUs

May 5, 2016

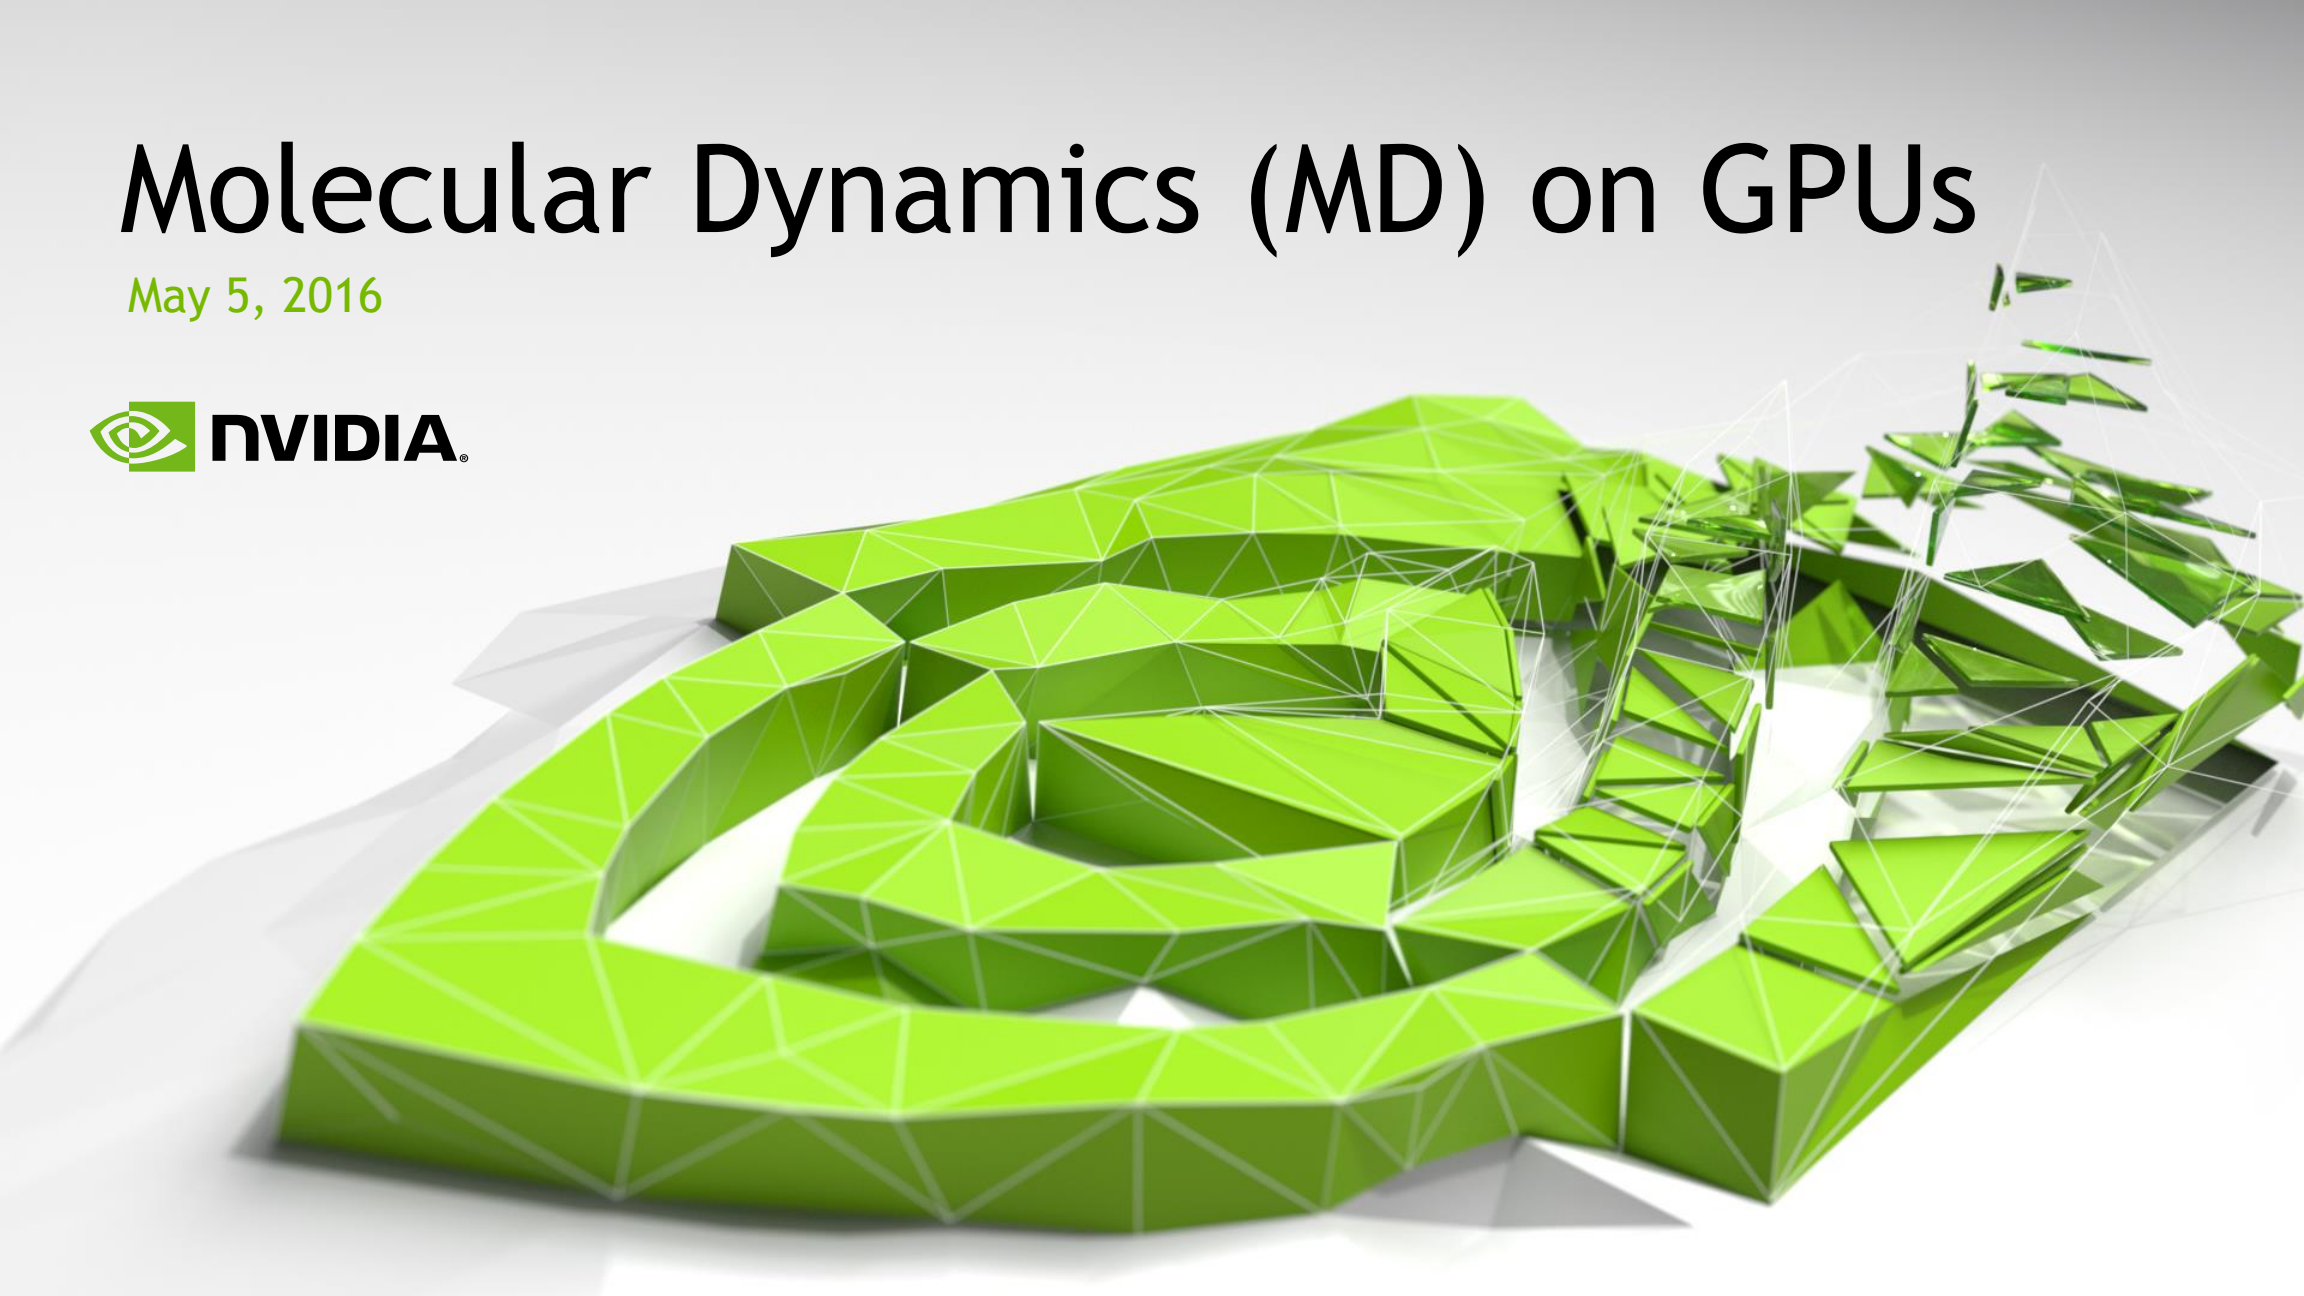

Accelerating Discoveries

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.

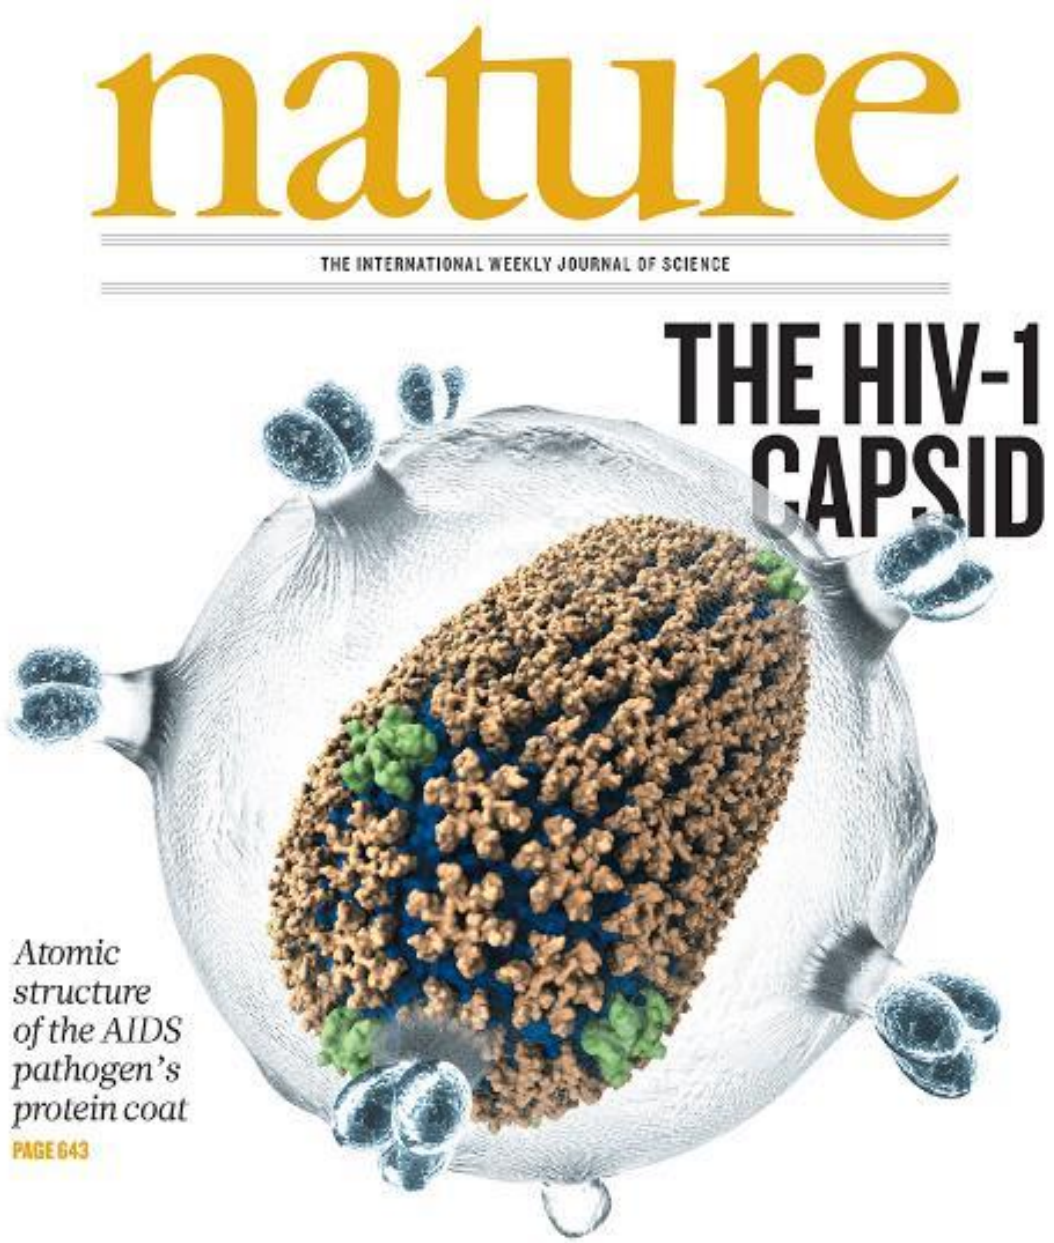

Overview of Life & Material Accelerated Apps

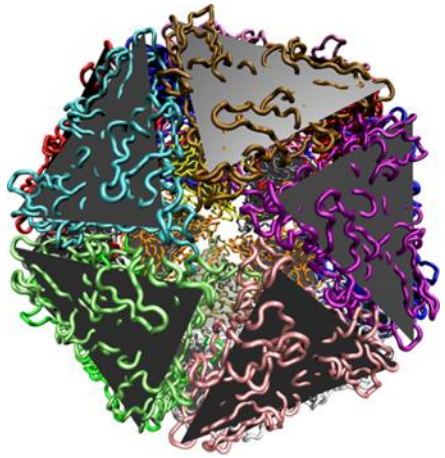

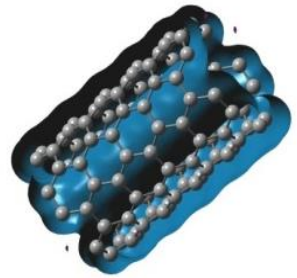

MD: All key codes are GPU-accelerated

- 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, 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	Double precision is important
Uses cuBLAS, cuFFT, CUDA	Uses cuBLAS, cuFFT, OpenACC
Geforce (Accademics), Tesla (Servers)	Tesla recommended
ECC off	ECC on

GPU-Accelerated Molecular Dynamics Apps

Green Lettering Indicates Performance Slides Included

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

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

Benefits of MD GPU-Accelerated Computing

Why wouldn't you want to turbocharge your research?

- 3x-8x Faster than CPU only systems in all tests (on average)
- Most major compute intensive aspects of classical MD ported
- Large performance boost with marginal price increase
- Energy usage cut by more than half
- GPUs scale well within a node and/or over multiple nodes
- K80 GPU is our fastest and lowest power high performance GPU yet

Try GPU accelerated MD apps for free – www.nvidia.com/GPUTestDrive

ACEMD

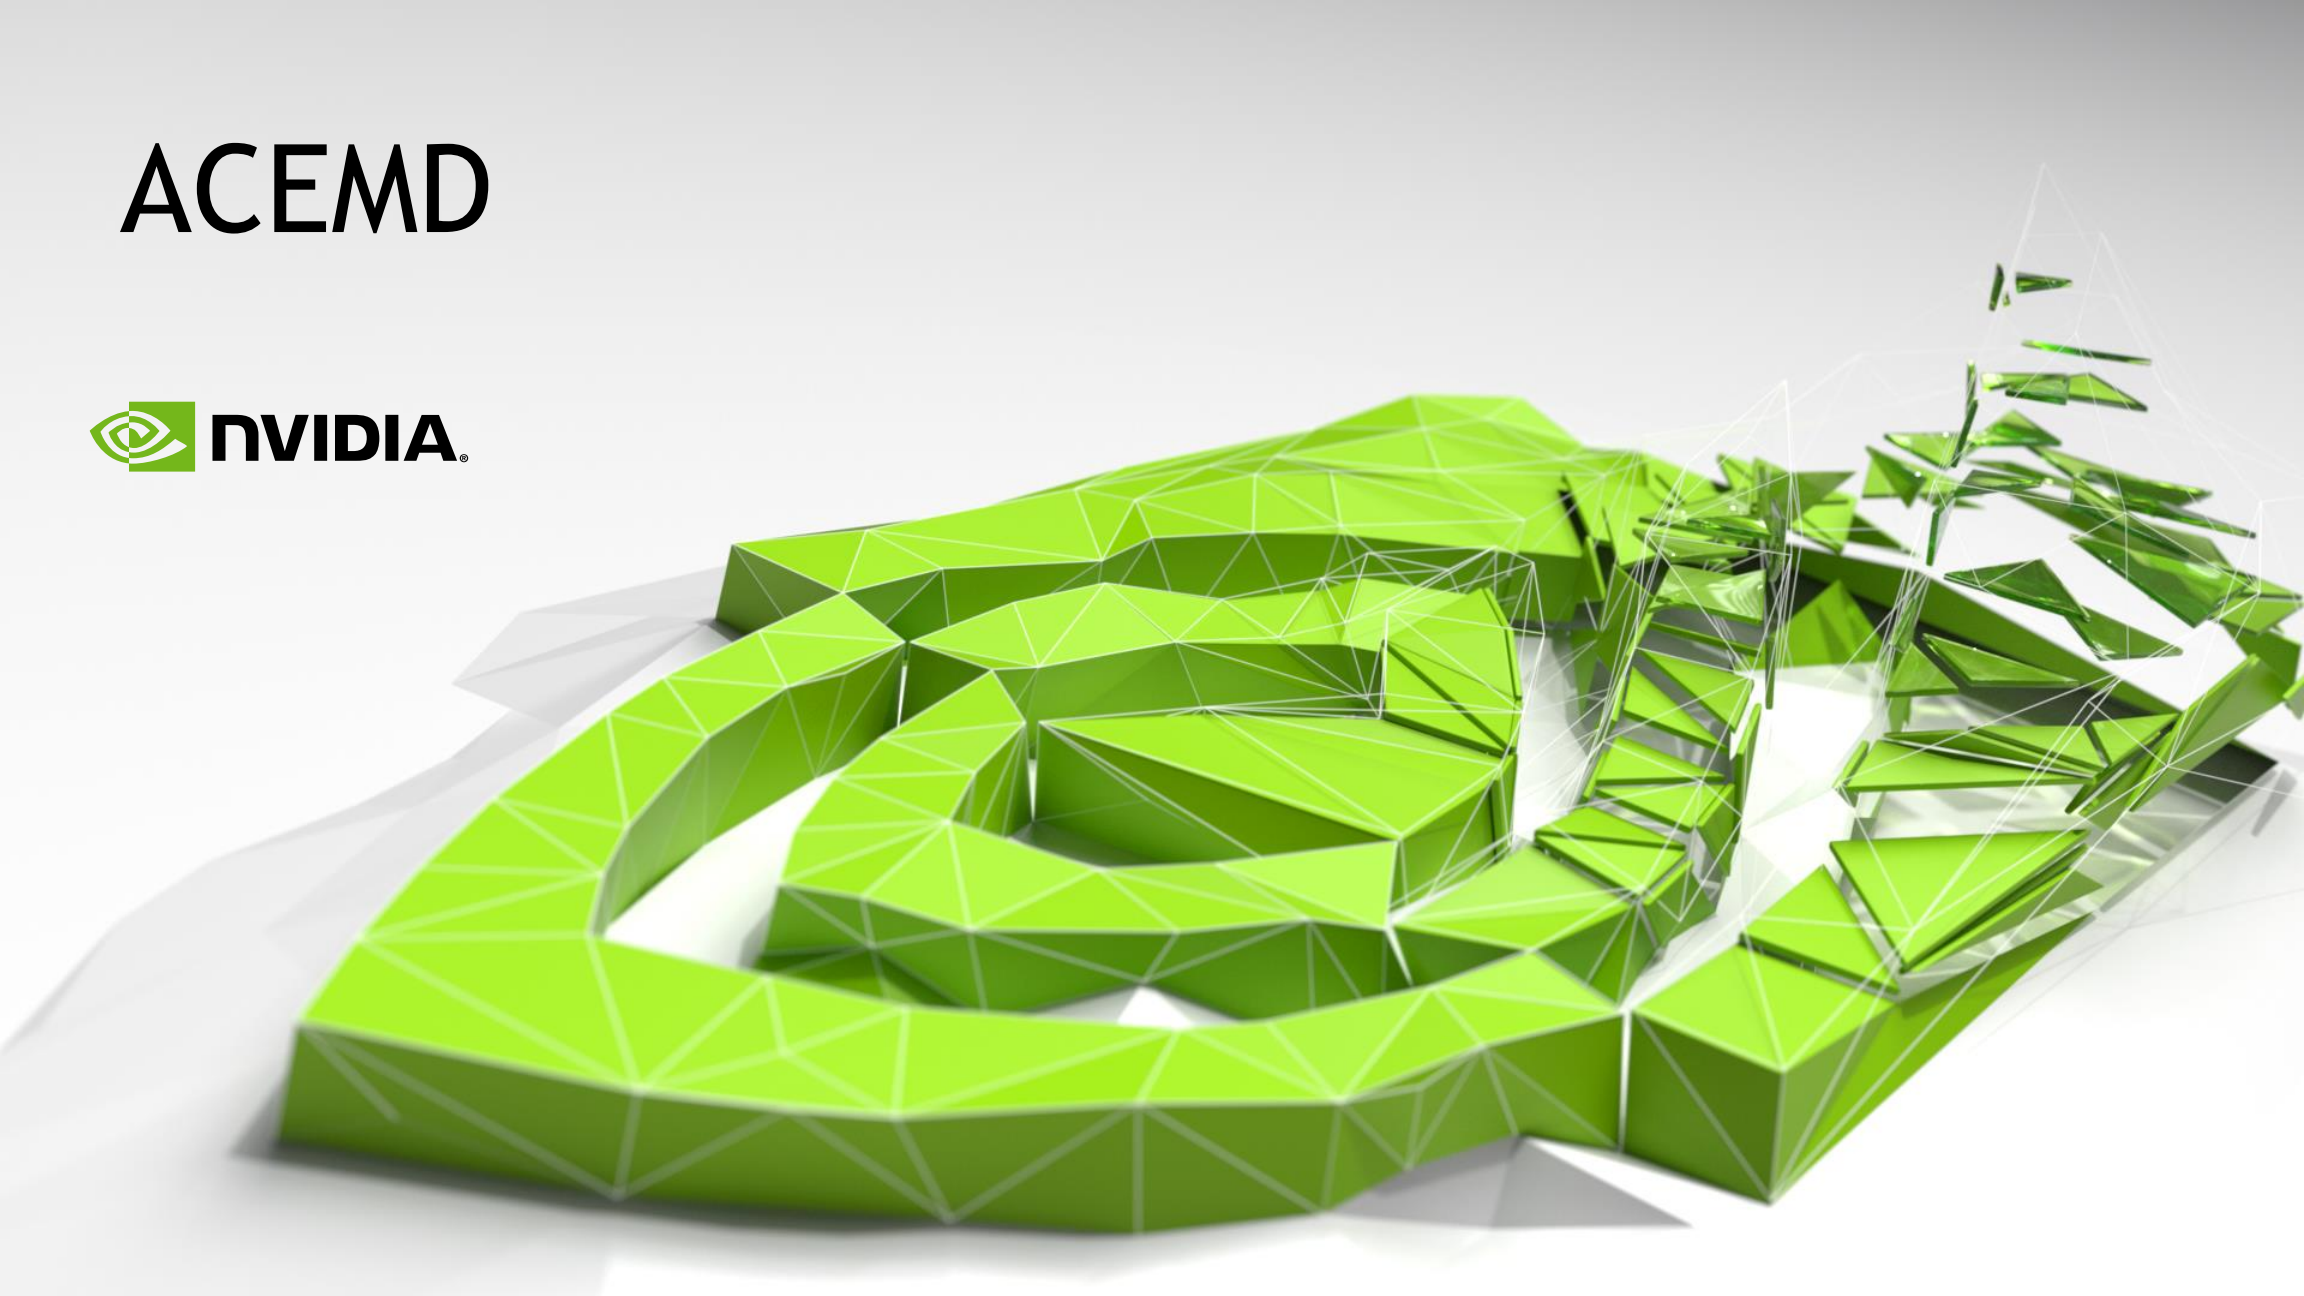

www.acellera.com

ACEMD: Extremely efficient and robust MD software built on GPUs

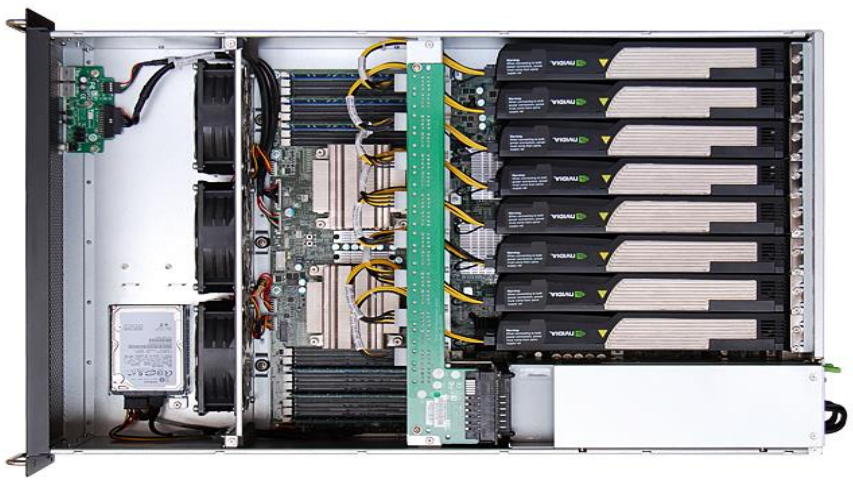

470 ns/day on 1 GPU for L-Iduronic acid (1362 atoms)

116 ns/day on 1 GPU for DHFR (23K atoms)

M. Harvey, G. Giupponi and G. De Fabritiis, *ACEMD: Accelerated molecular dynamics simulations in the microseconds timescale*, J. Chem. Theory and Comput. 5, 1632 (2009)

www.acellera.com

ACEMD

- **Standardised and easy to use**: ACEMD reads CHARMM/NAMD and AMBER input files and uses similar syntax to other MD software.
- **Fully featured:** NVT, NPT, PME, TCL, PLUMED, CAMSHIFT[1]
- **Robust**: ACEMD is a proven computational engine and is used in one of the largest distributed computing projects worldwide: GPUGRID.
- **Compatible**: ACEMD works with CUDA and OpenCL, the new standard framework for parallel and high-performance computing.
- **Validated**: ACEMD is used in reputable academic and industrial institutions. Results describing its applications have appeared in peer-reviewed journals of high impact such as PNAS, PLoS and JACS.[2]

1 M. J. Harvey and G. De Fabritiis, *An implementation of the smooth particle-mesh Ewald (PME) method on GPU hardware*, J. Chem. Theory Comput., 5, 2371–2377 (2009)

2 For a list of selected references see http://www.acellera.com/acemd/publications

AMBER 14

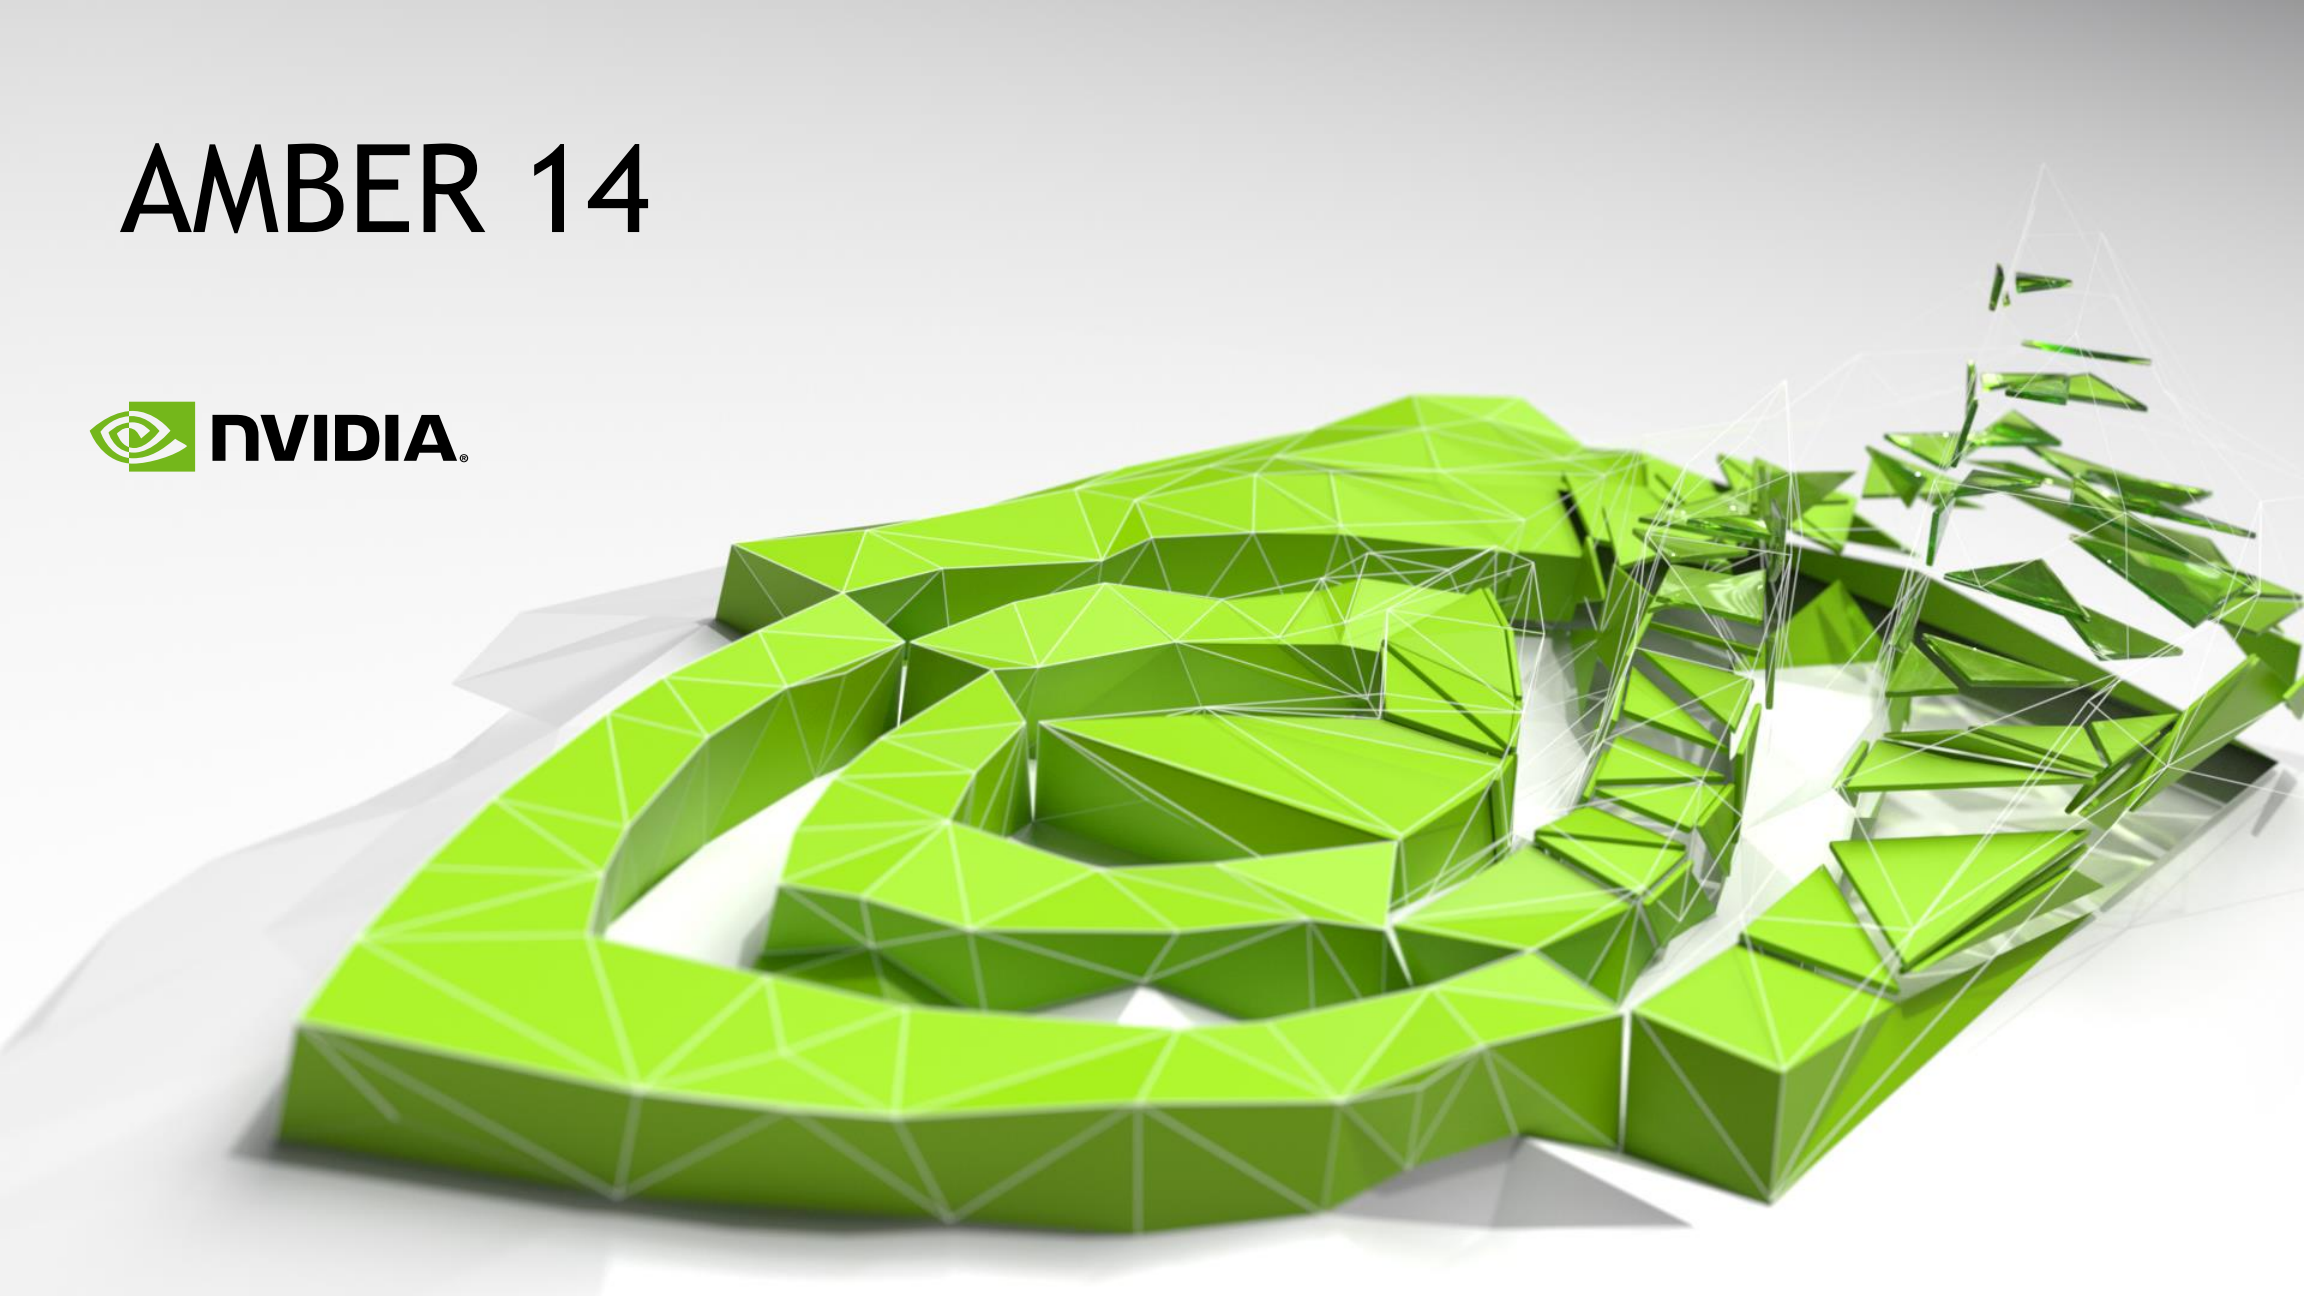

Ross Walker (AMBER developer) video

AMBER 14 vs. AMBER 12

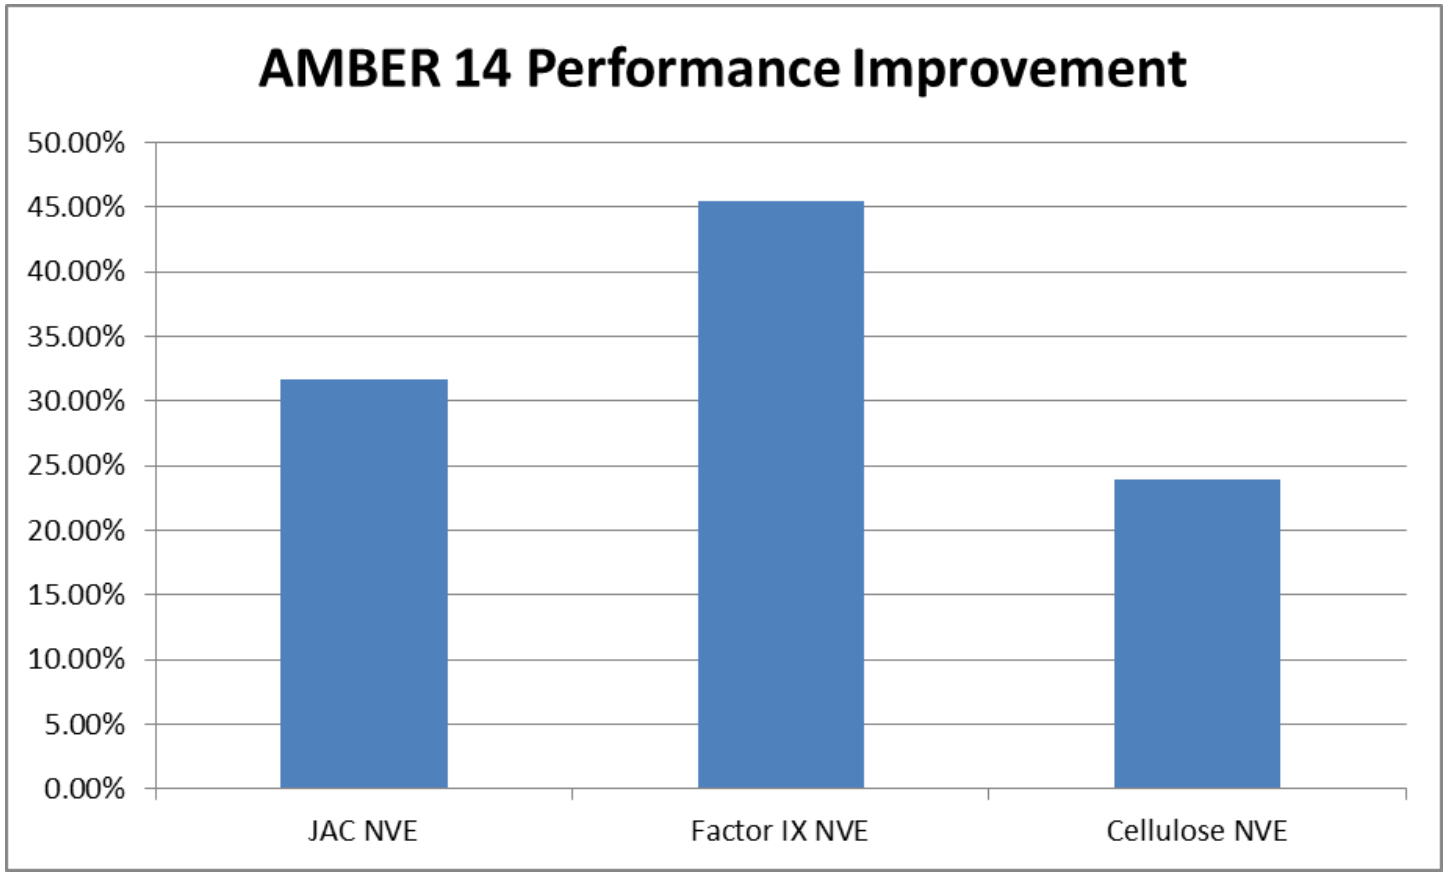

Courtesy of Scott Le Grand From GTC 2014 presentation

AMBER 14; large P2P and small Boost Clocks impacts

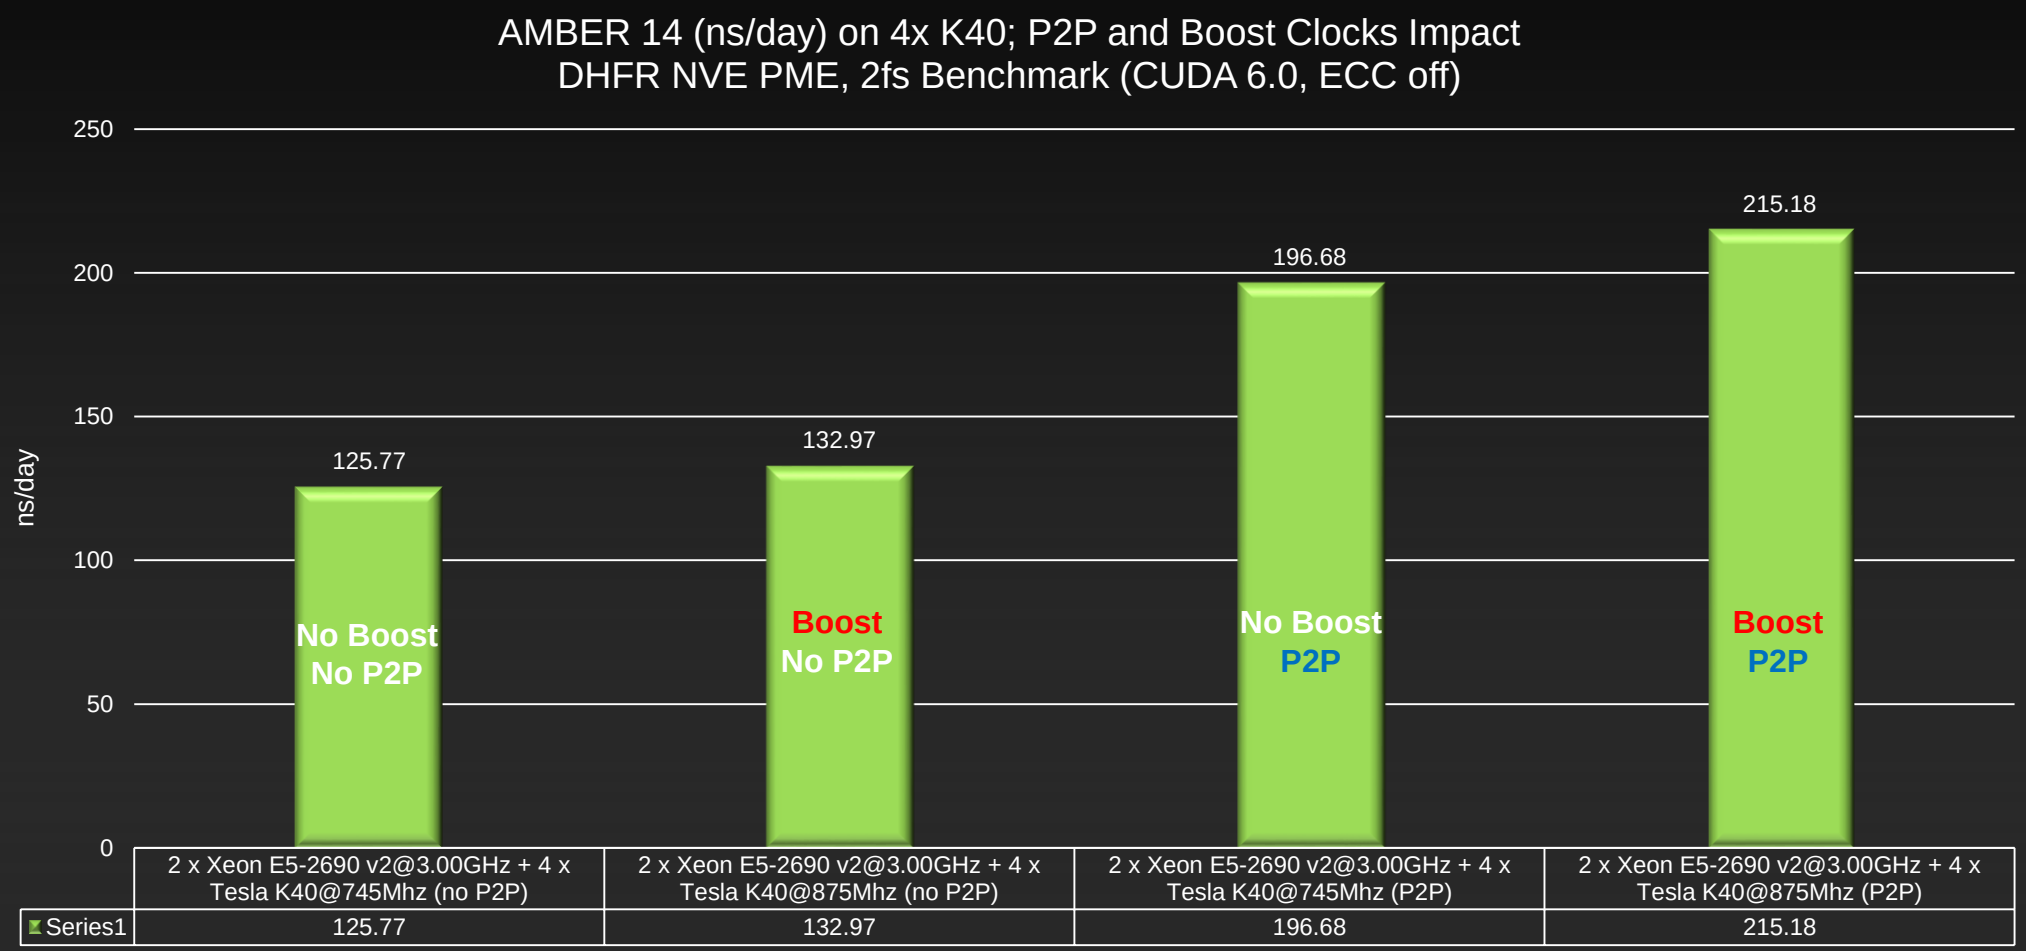

AMBER Performance Over Time

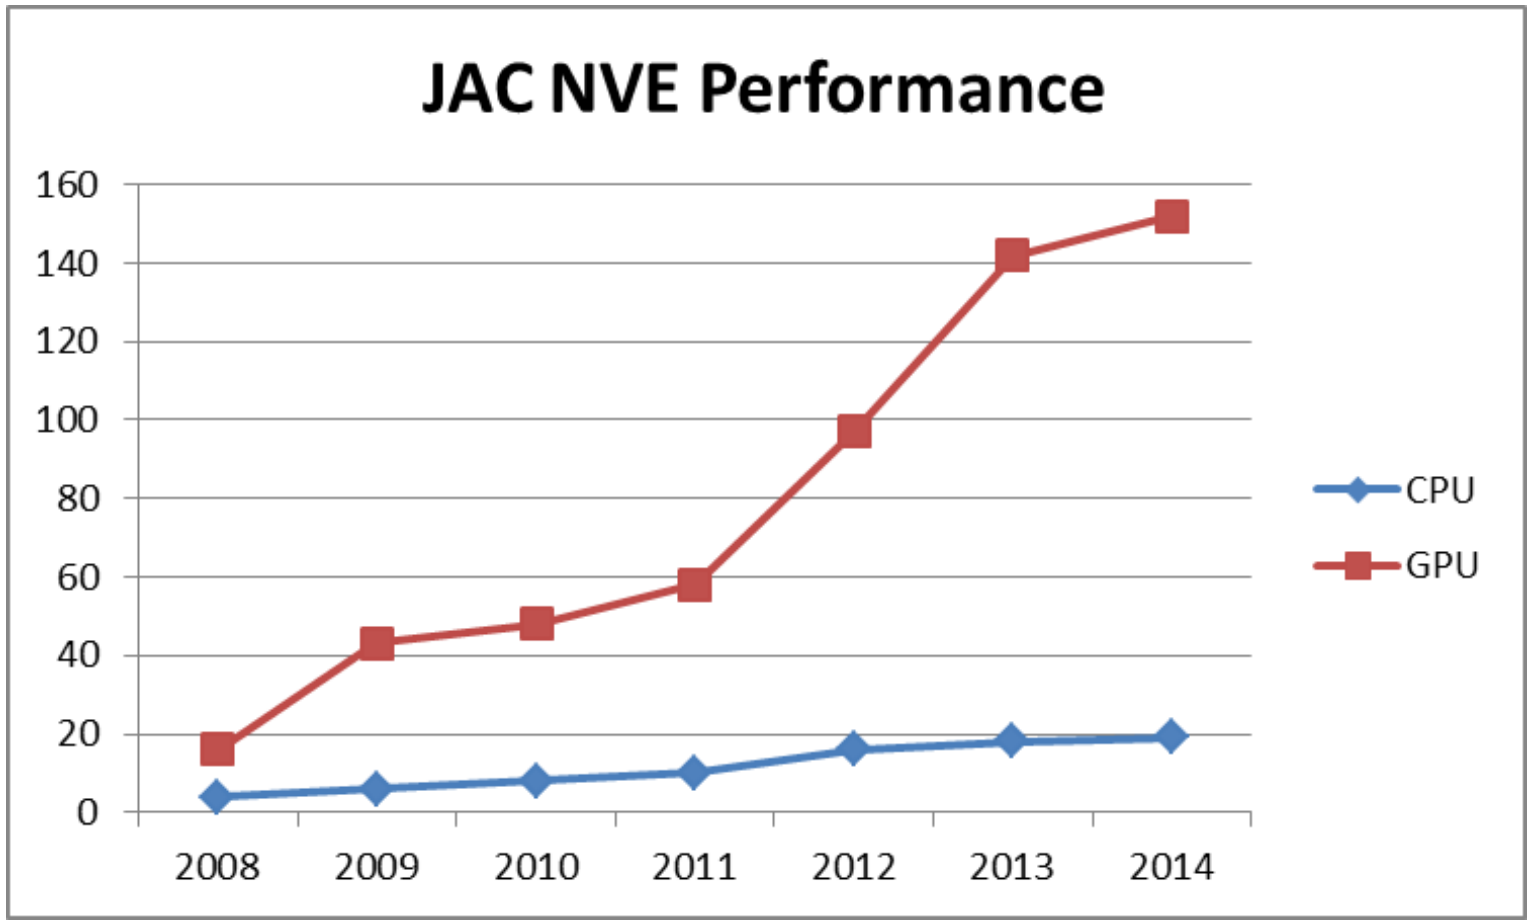

Courtesy of Scott Le Grand From GTC 2014 presentation

Protein Folding Simulation With AMBER Accelerated By GPUs

2.95x Faster

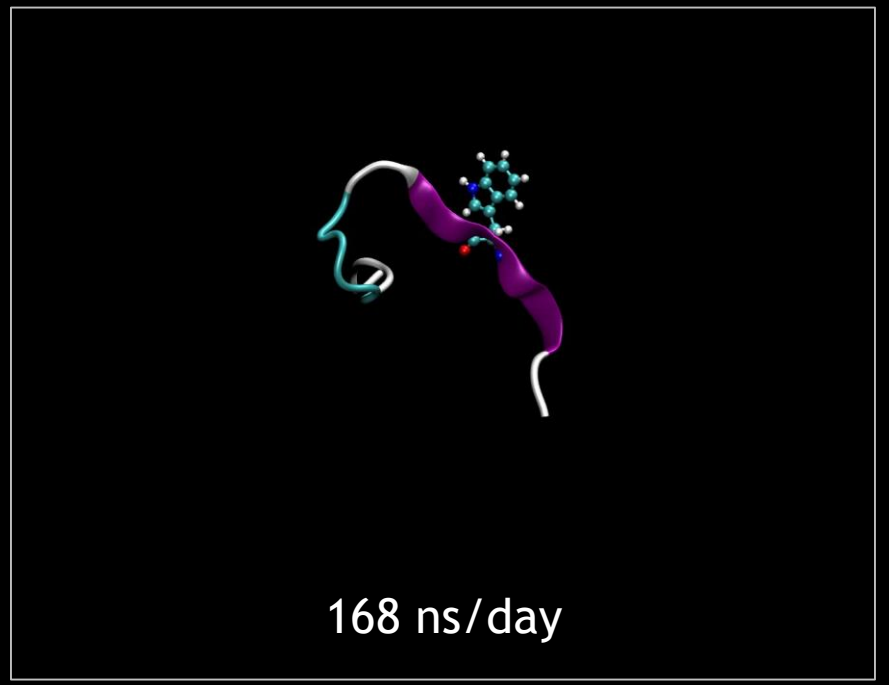

168 ns/day

2x8xE5-2670 CPU

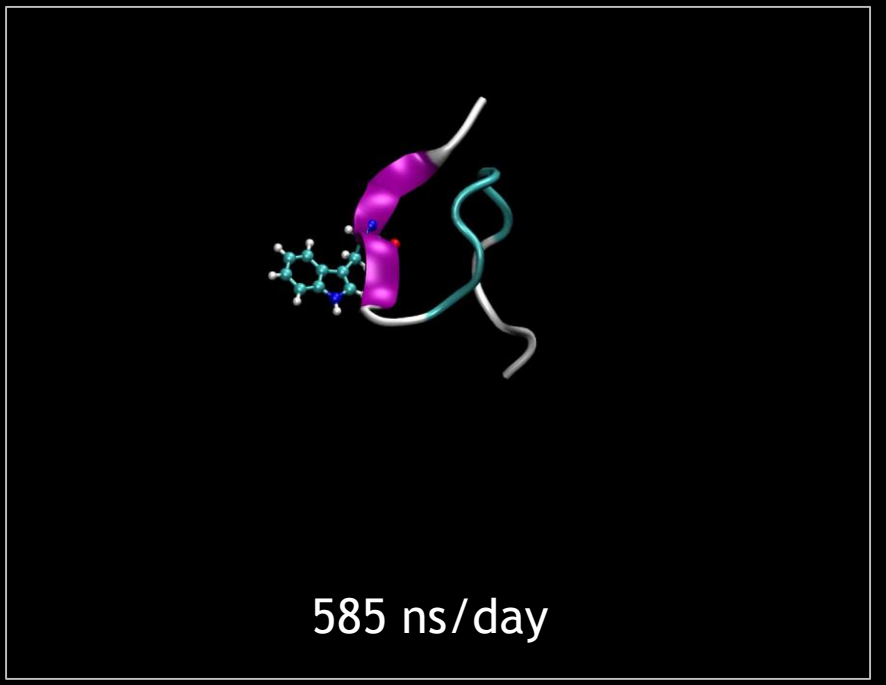

585 ns/day

2x8xE5-2670 CPU + Tesla K20X GPU

Data courtesy of AMBER.org

AMBER 14 with P2P, Higher Density Nodes

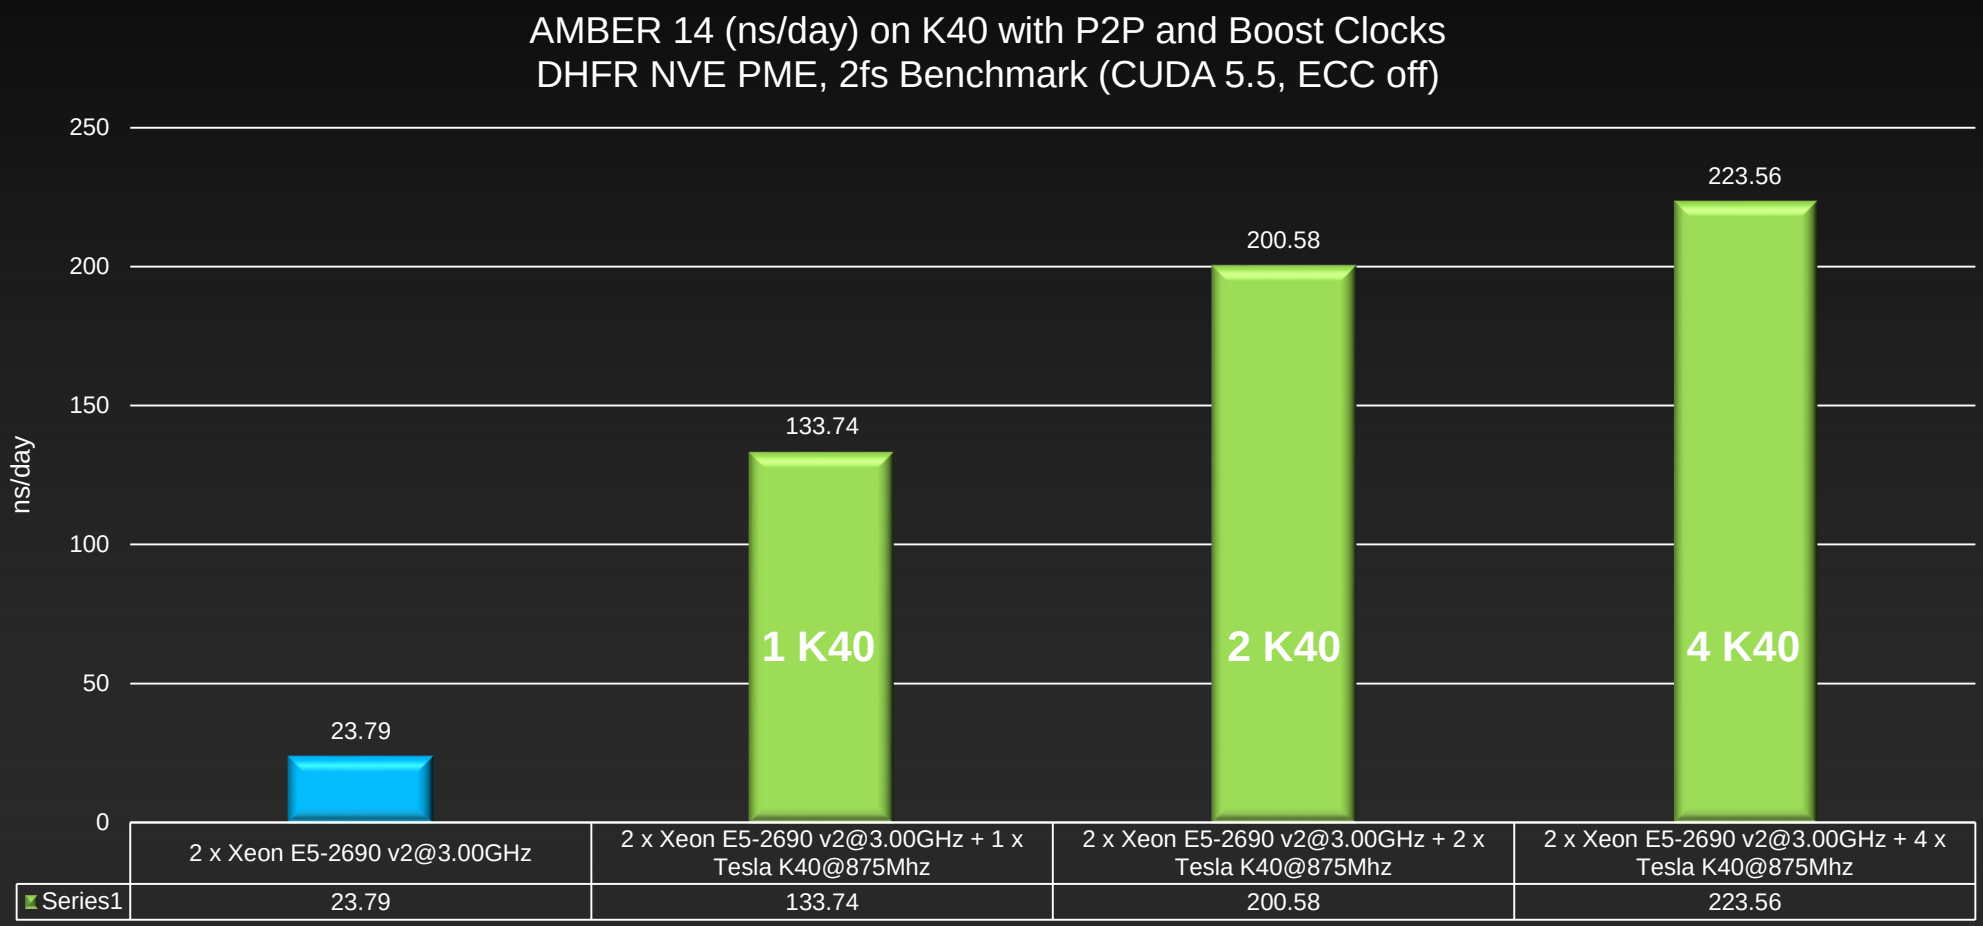

AMBER 14 and K40 with P2P, fastest GPU yet!

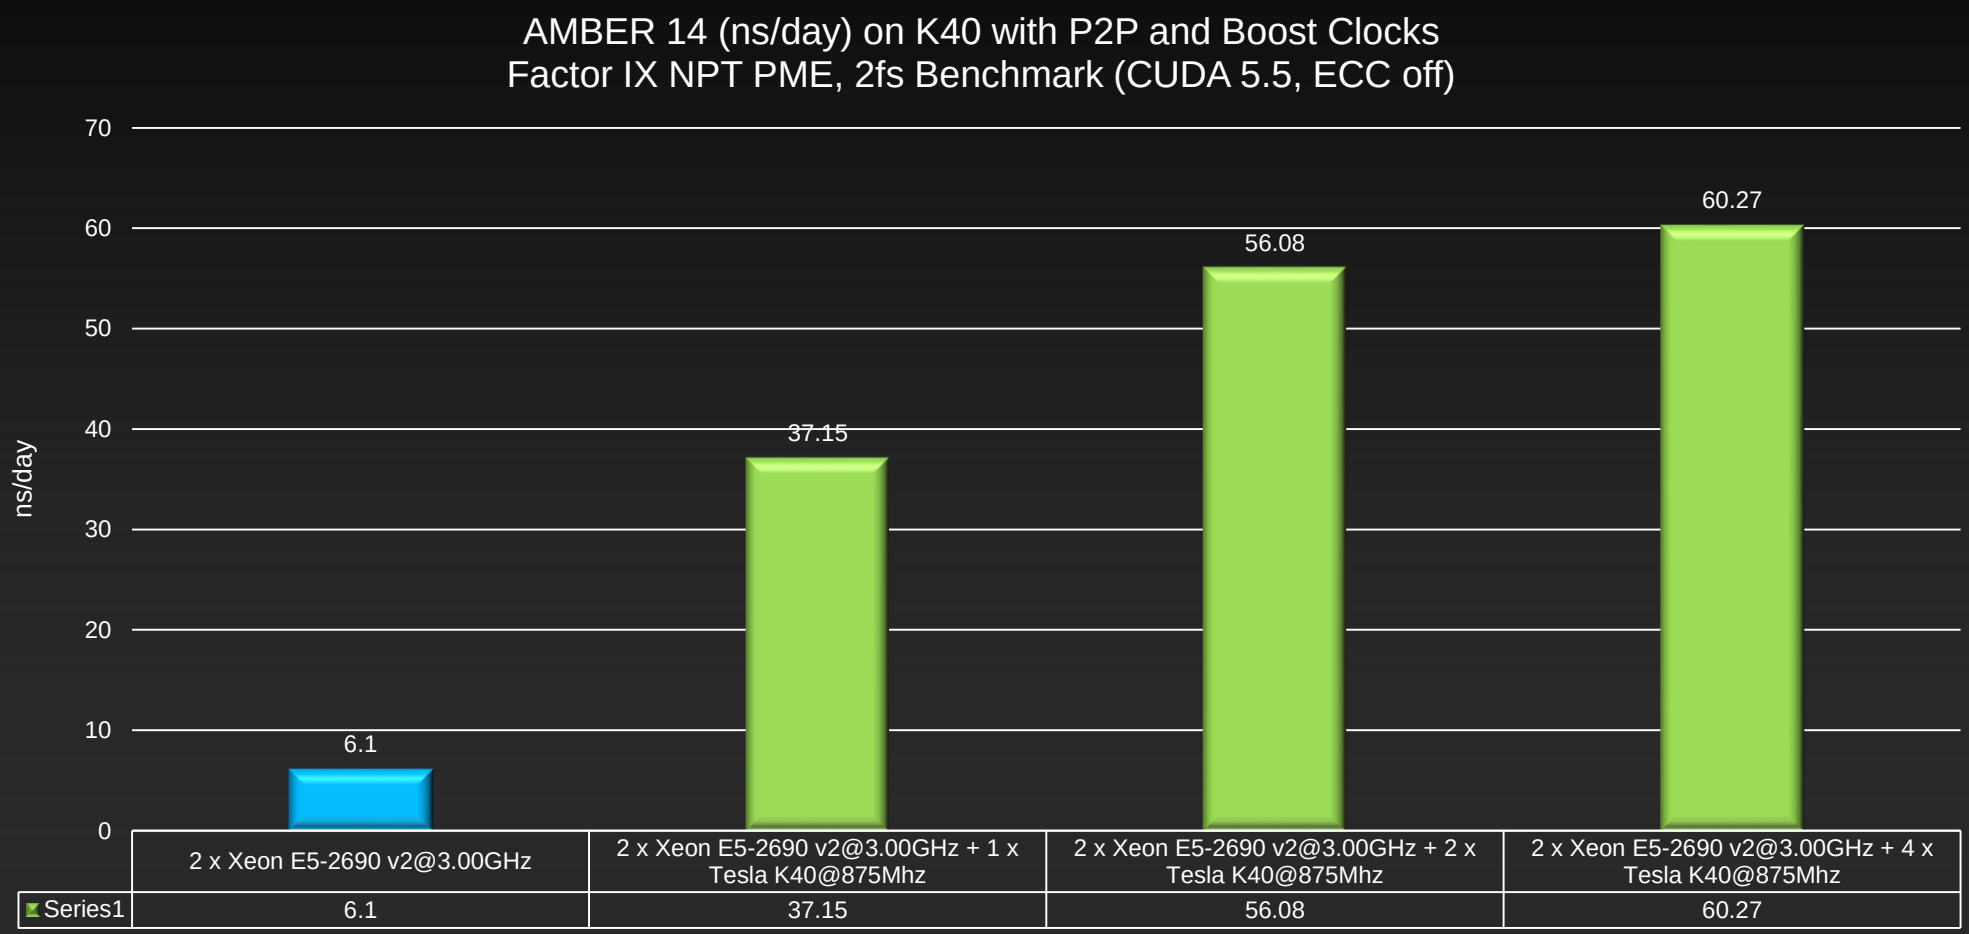

AMBER 14 and K40 with P2P, fastest GPU yet!

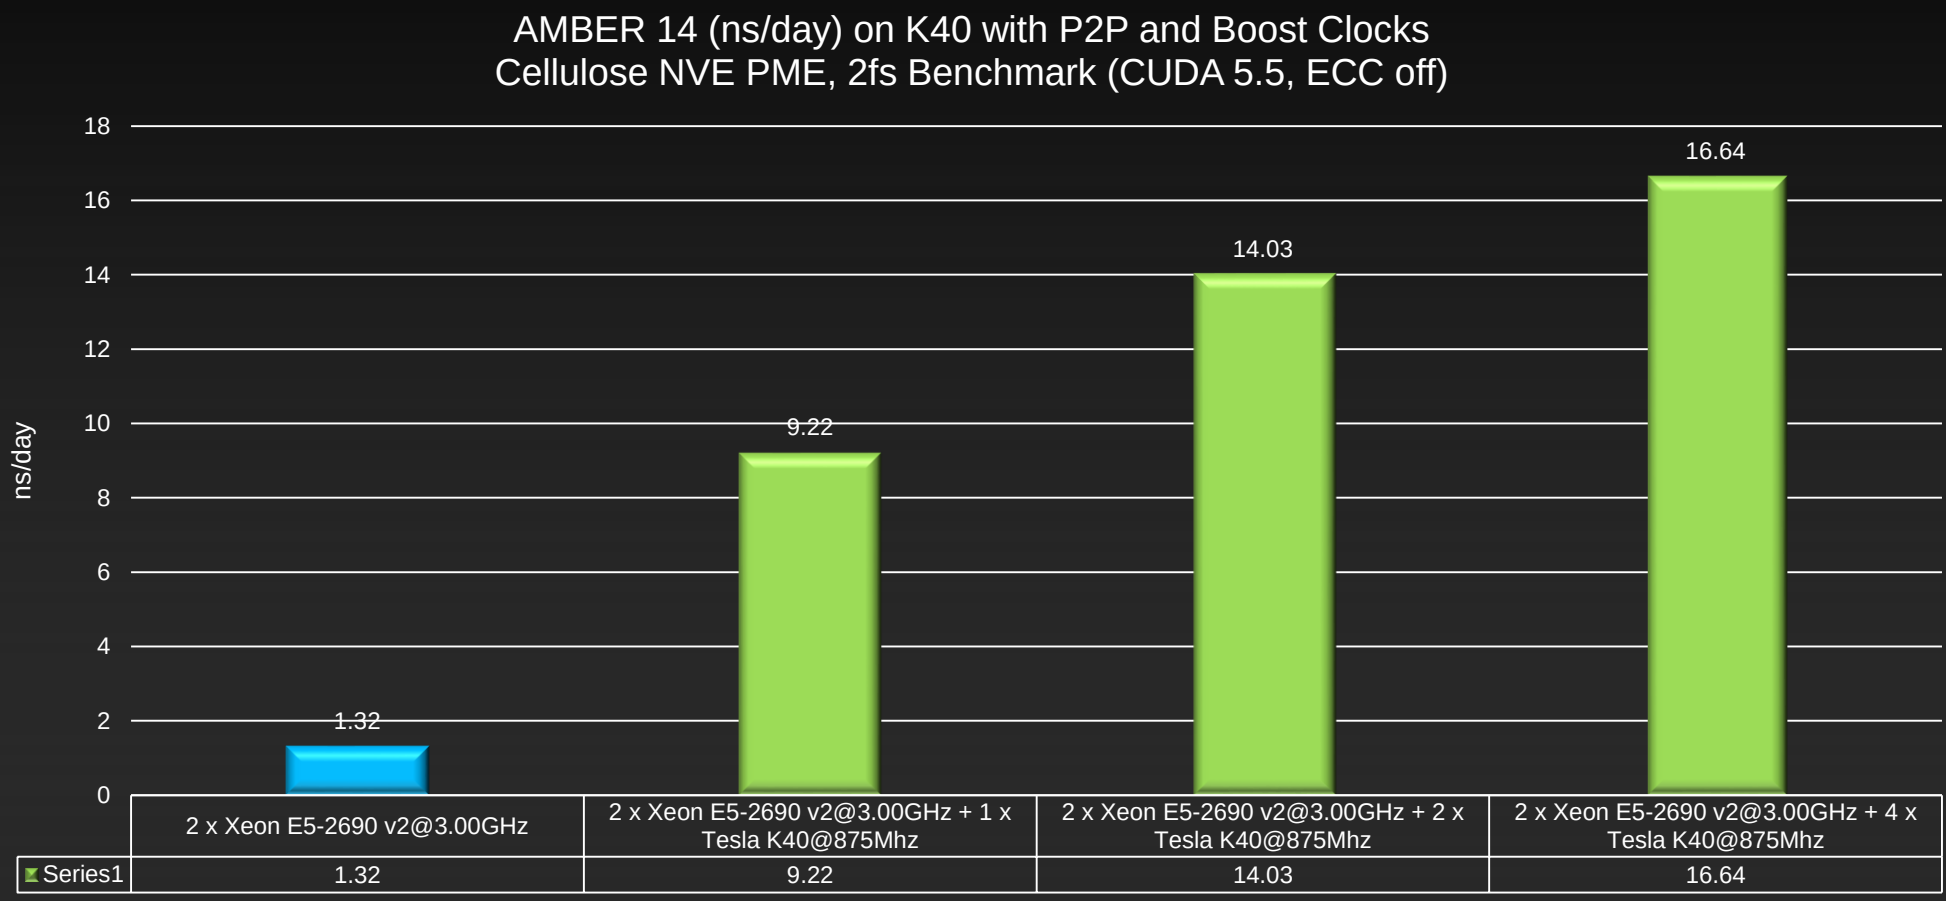

AMBER 14 and K40 with P2P, fastest GPU yet!

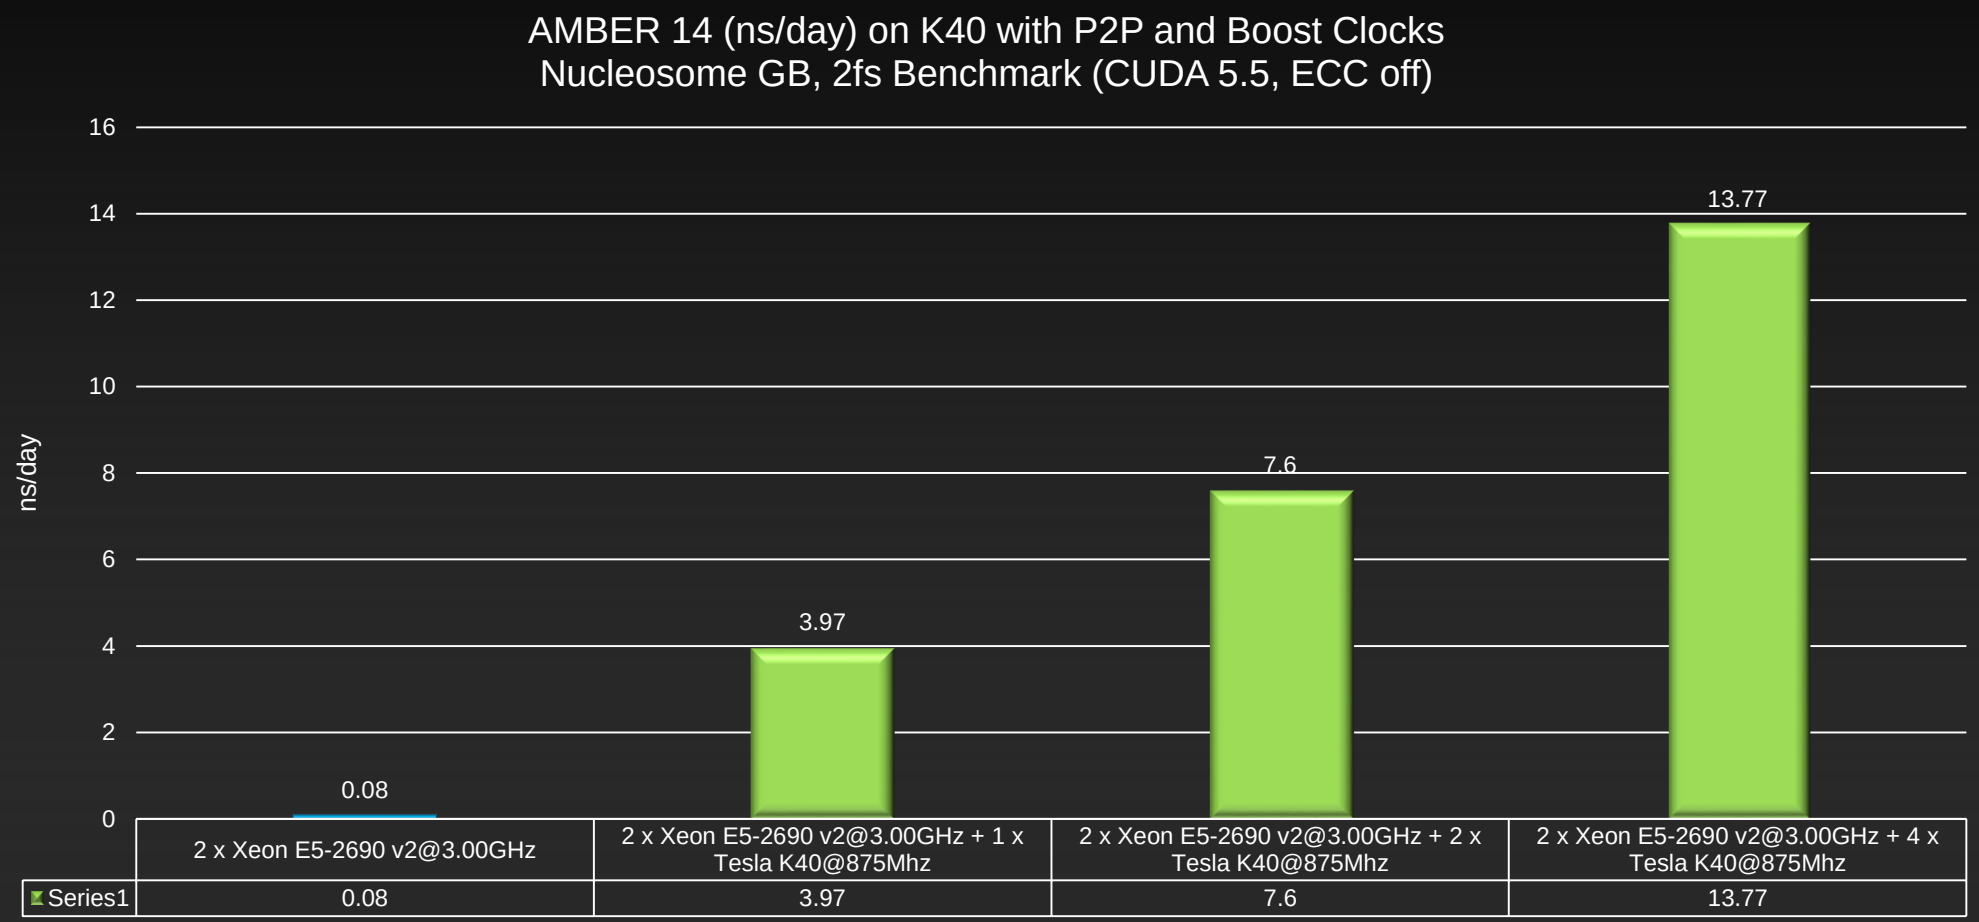

Cellulose on K40s, K80s and M6000s

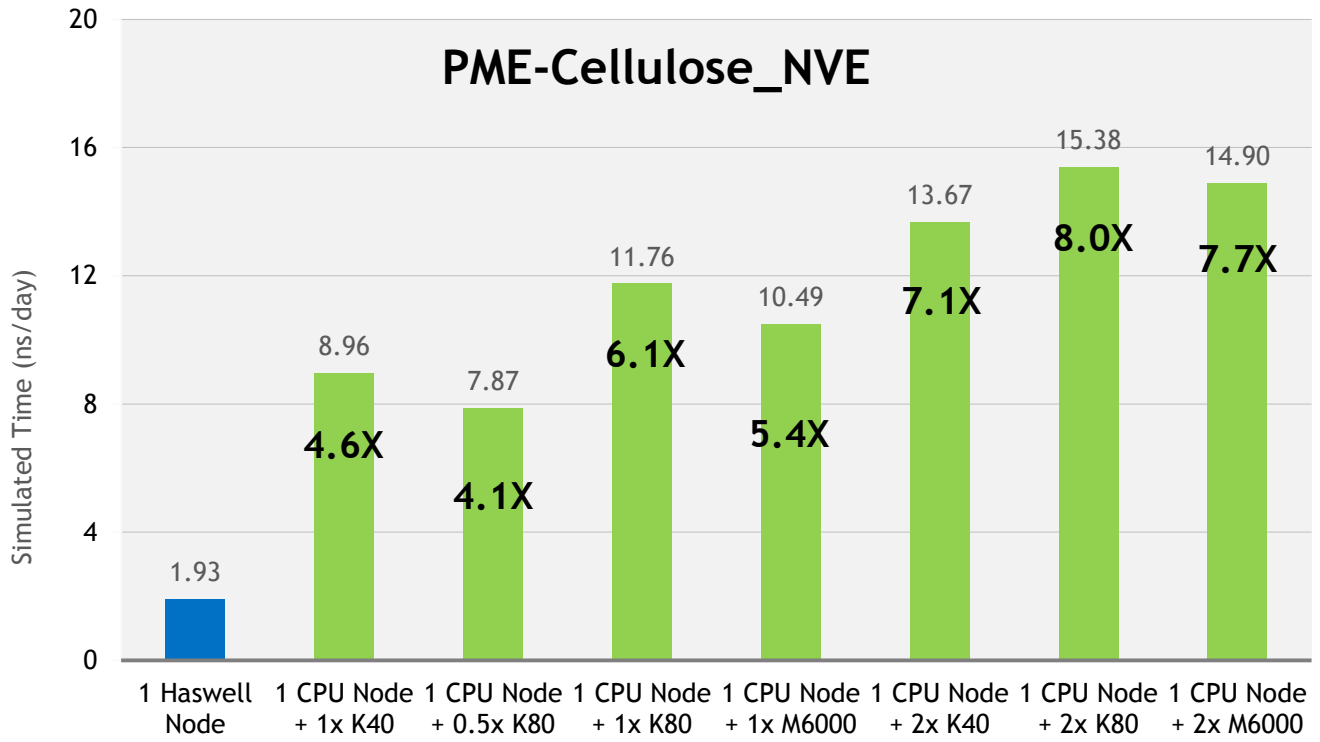

Running AMBER version 14

The blue node contains Dual Intel E5-2698 v3@2.3GHz, 3.6GHz Turbo CPUs

The green nodes contain Dual Intel E5-2698 v3@2.3GHz, 3.6GHz Turbo CPUs + either NVIDIA Tesla K40@875Mhz, Tesla K80@562Mhz (autoboost), or Quadro M6000@987Mhz GPUs

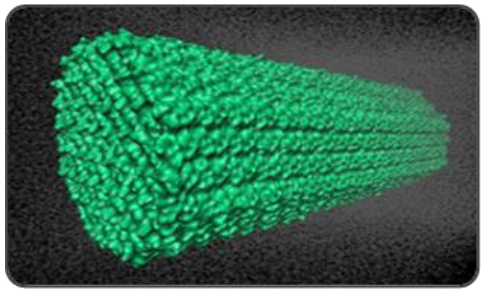

Factor IX on K40s, K80s and M6000s

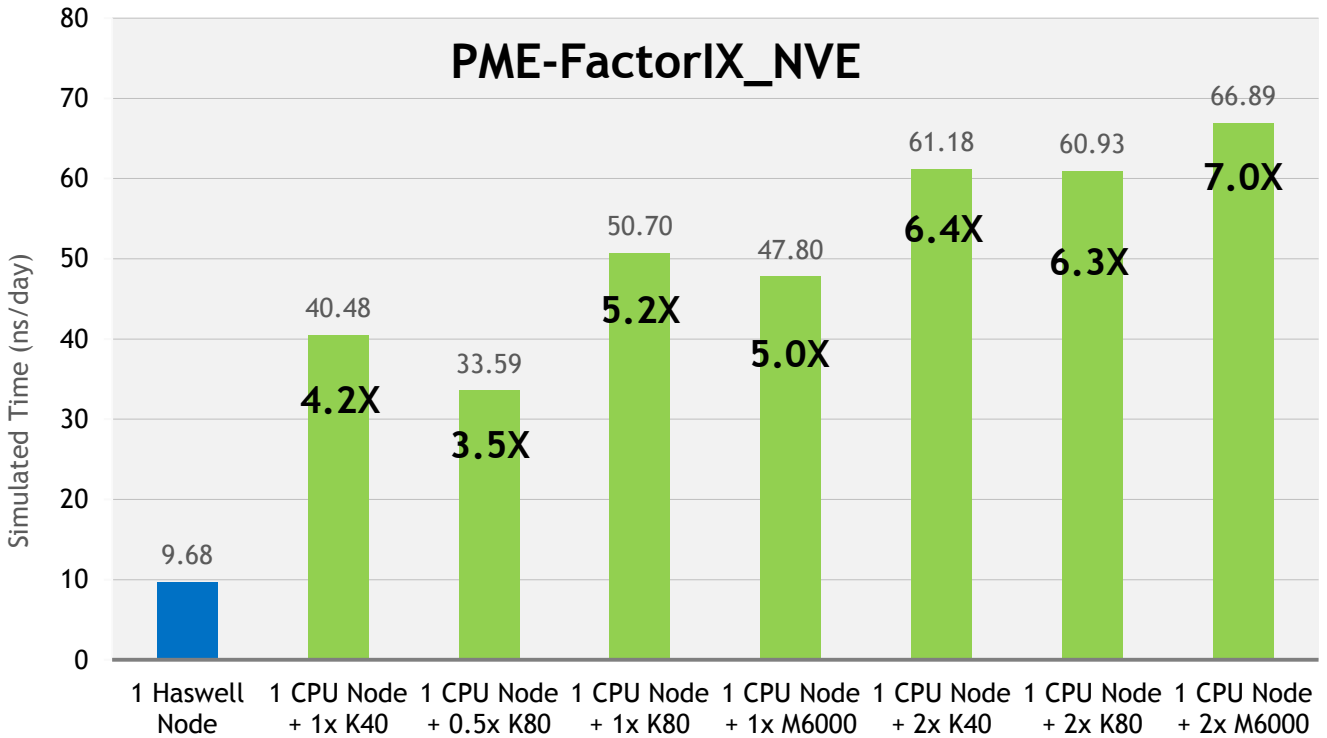

Running AMBER version 14

The blue node contains Dual Intel E5-2698 v3@2.3GHz, 3.6GHz Turbo CPUs

The green nodes contain Dual Intel E5-2698 v3@2.3GHz, 3.6GHz Turbo CPUs + either NVIDIA Tesla K40@875Mhz, Tesla K80@562Mhz (autoboost), or Quadro M6000@987Mhz GPUs

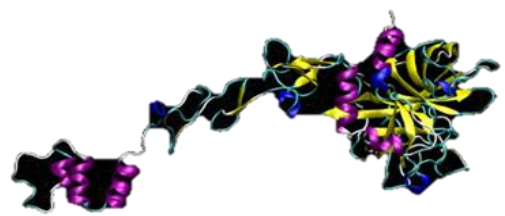

JAC on K40s, K80s and M6000s

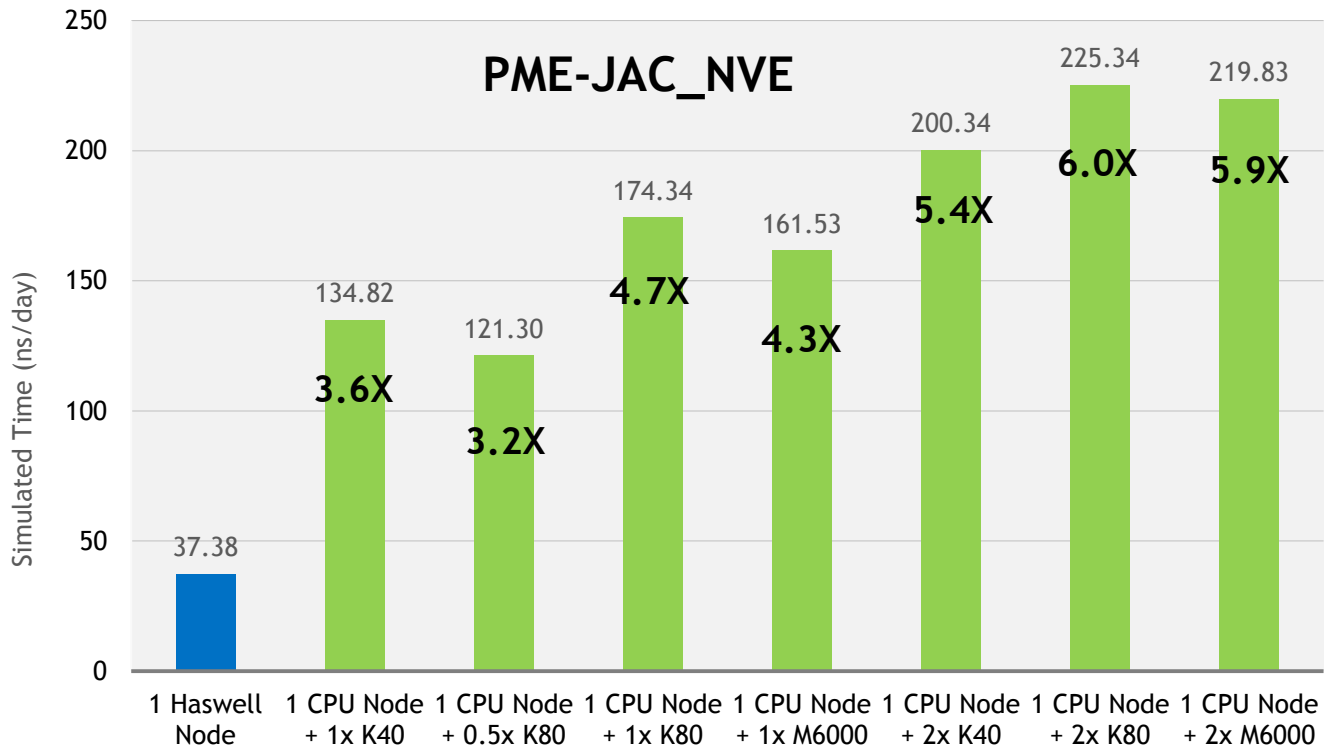

Running AMBER version 14

The blue node contains Dual Intel E5-2698 v3@2.3GHz, 3.6GHz Turbo CPUs

The green nodes contain Dual Intel E5-2698 v3@2.3GHz, 3.6GHz Turbo CPUs + either NVIDIA Tesla K40@875Mhz, Tesla K80@562Mhz (autoboost), or Quadro M6000@987Mhz GPUs

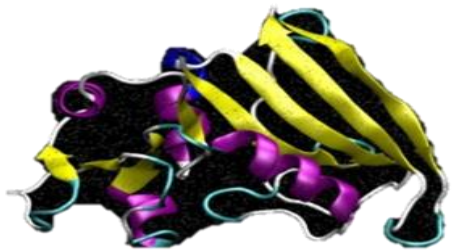

AMBER 14

March 2016

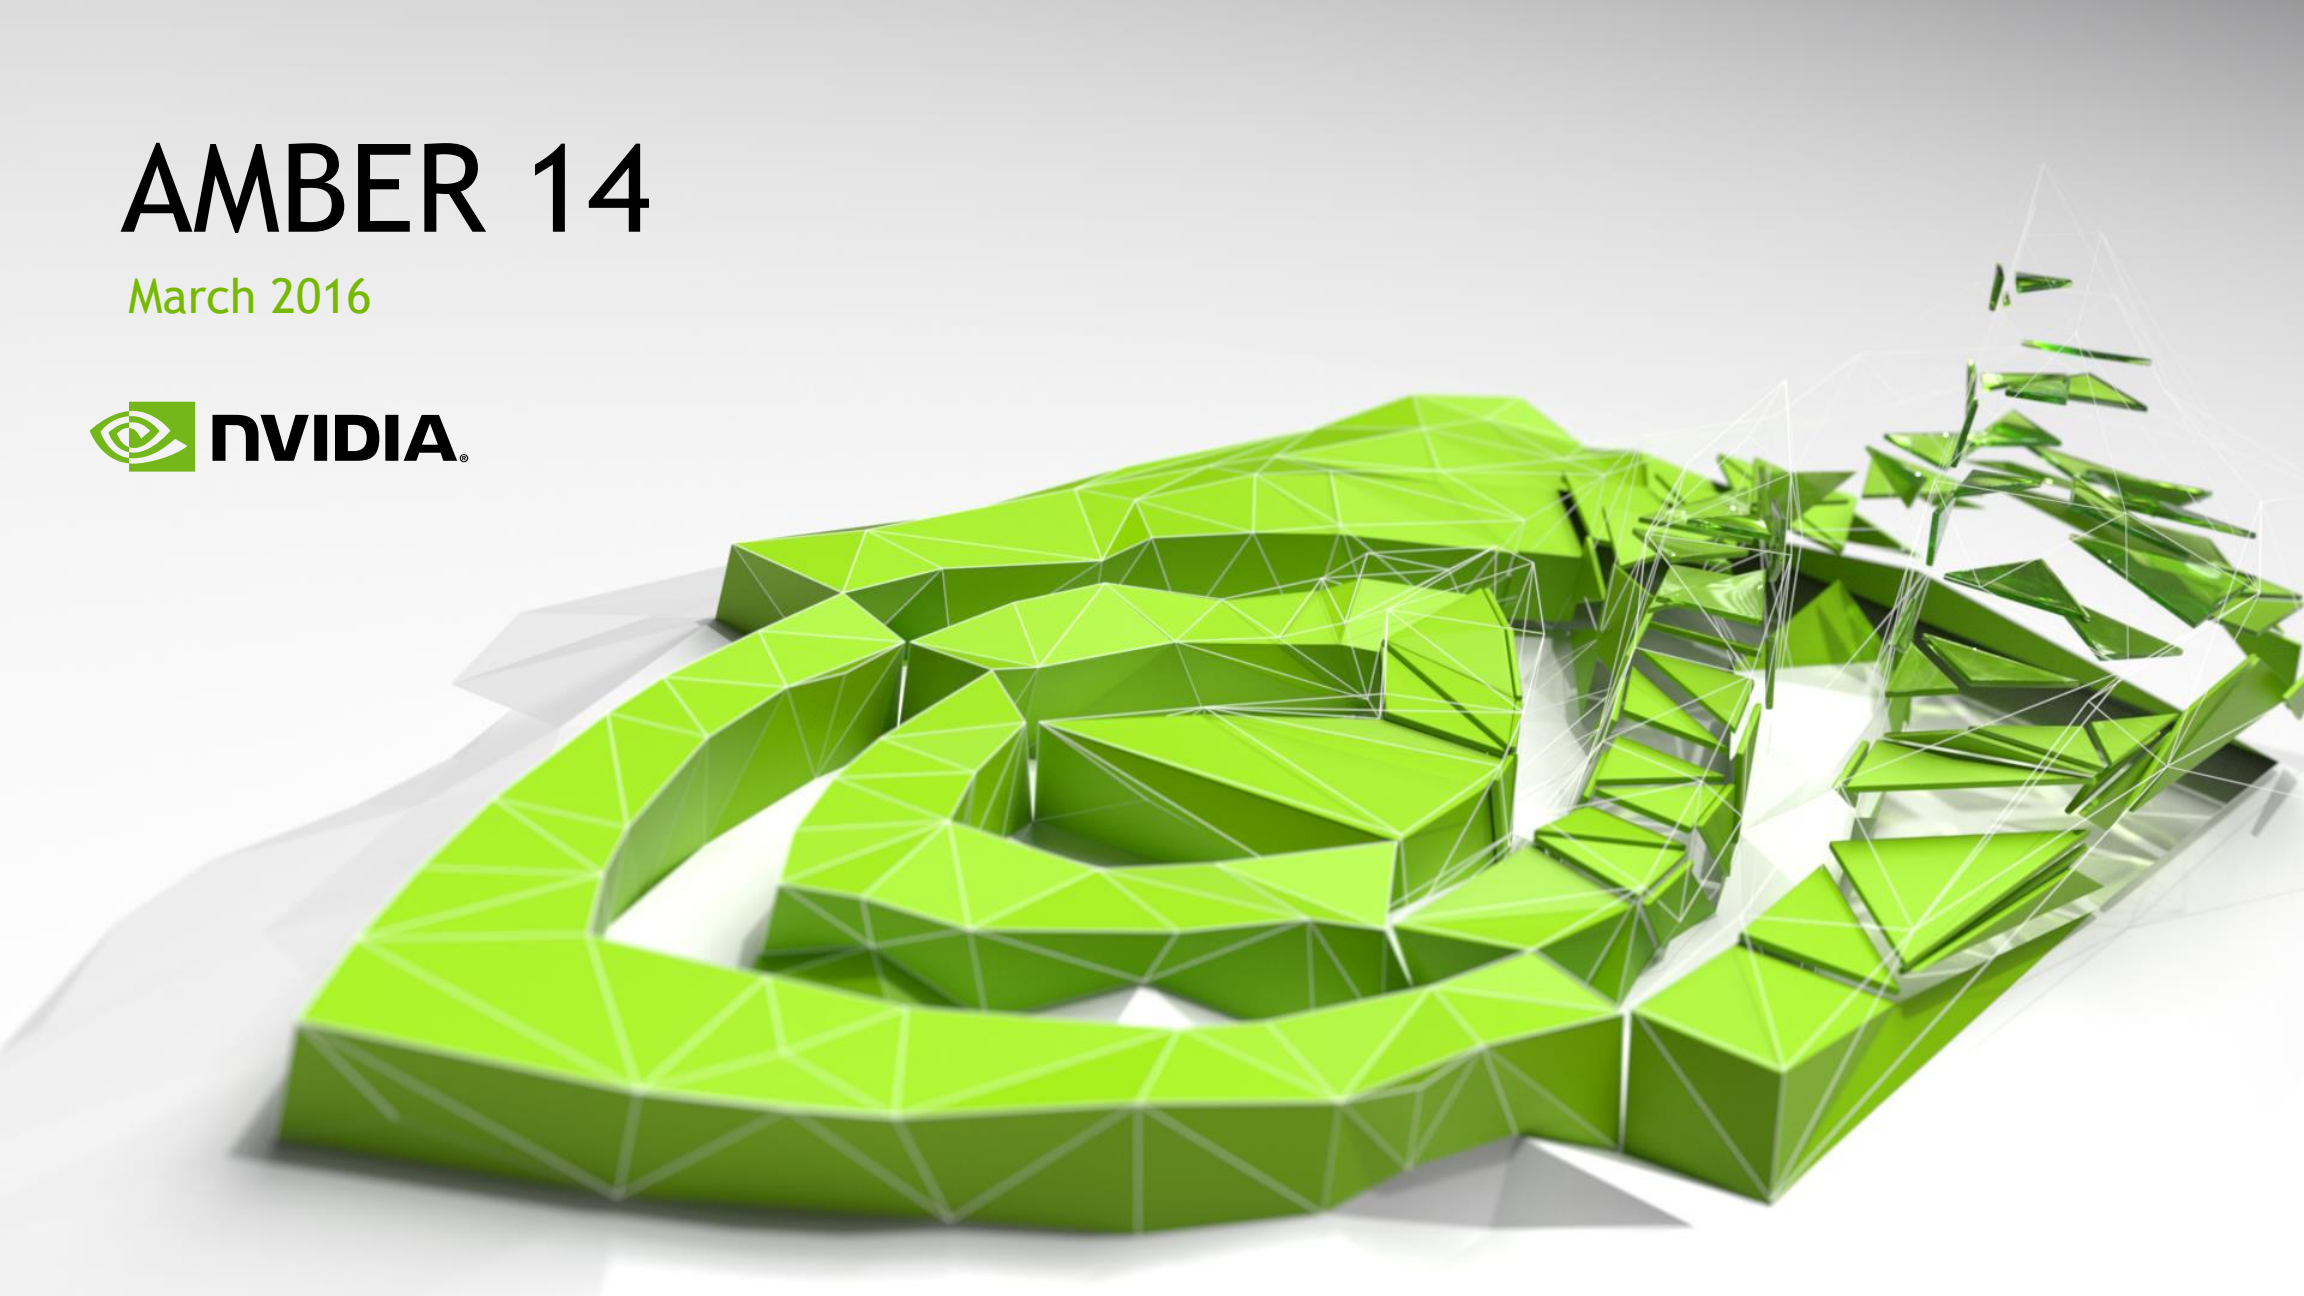

Cellulose on M40s

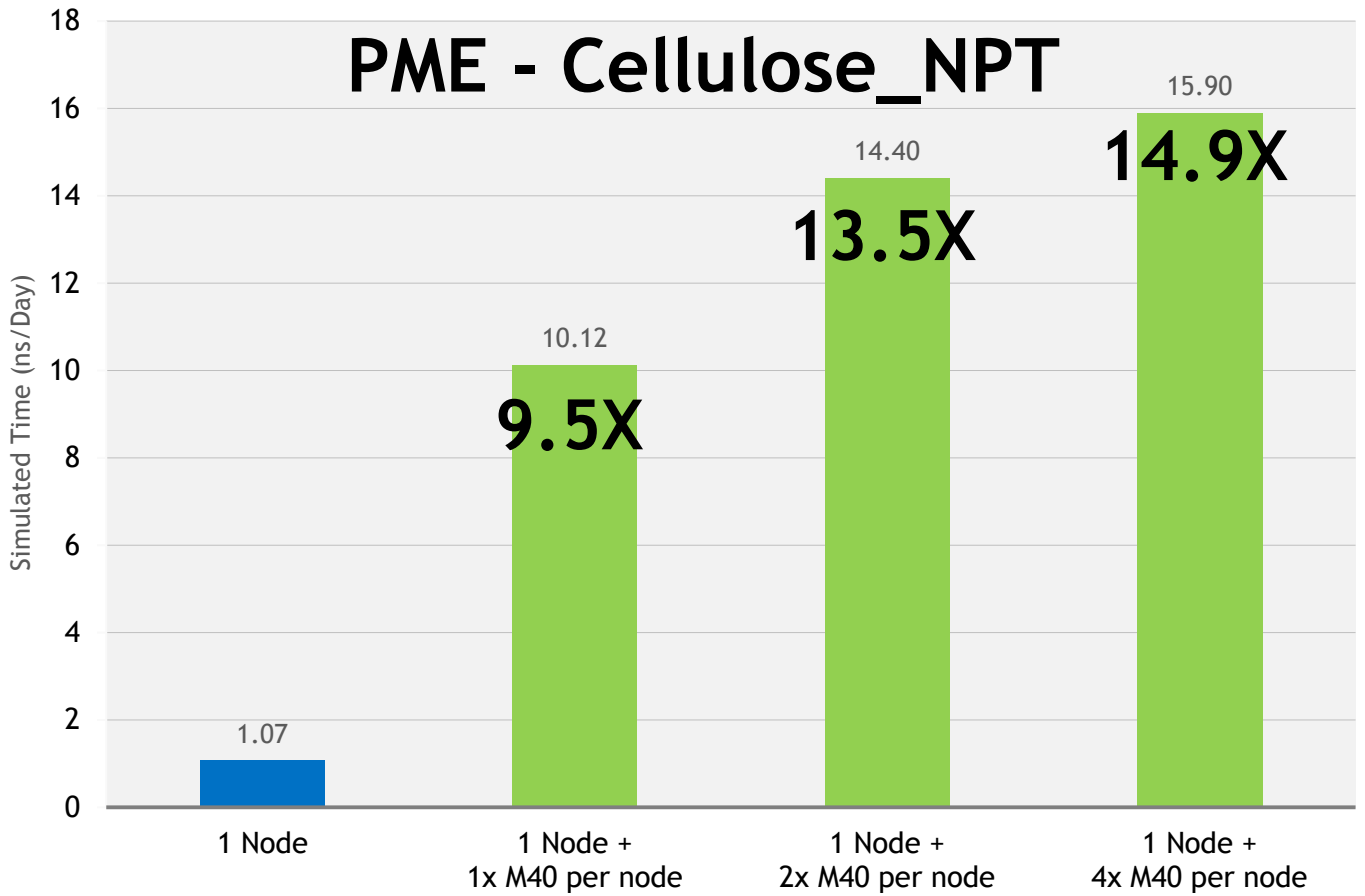

Running AMBER version 14

The blue node contain Single Intel Xeon E5-2698 v3@2.30GHz (Haswell) CPUs

The green nodes contain Single Intel Xeon E5-2697 v2@2.70GHz (IvyBridge) CPUs + Tesla M40 (autoboost) GPUs

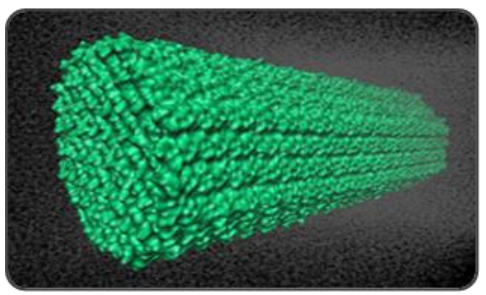

Cellulose on M40s

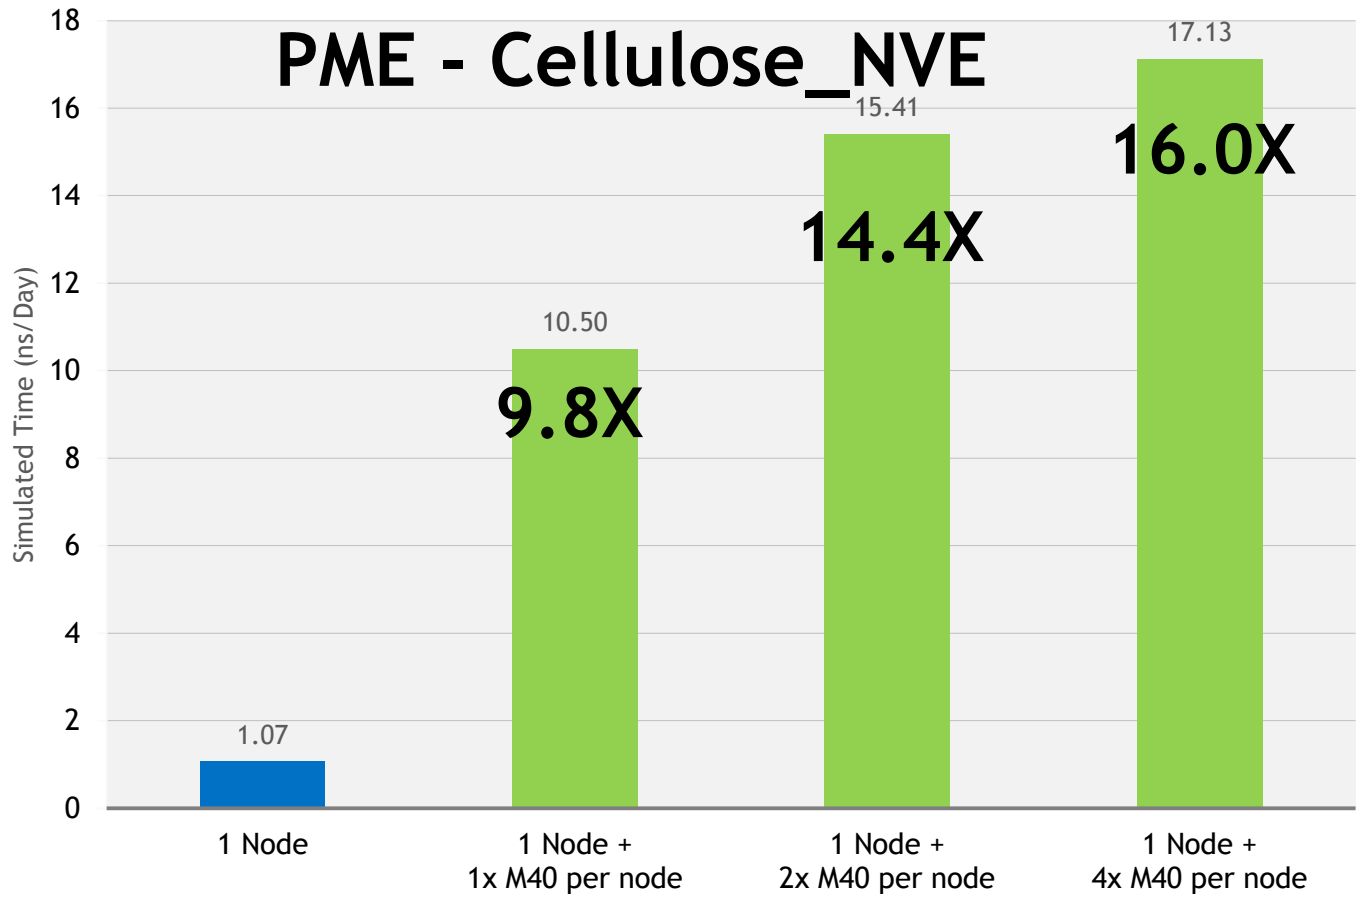

Running AMBER version 14

The blue node contain Single Intel Xeon E5-2698 v3@2.30GHz (Haswell) CPUs

The green nodes contain Single Intel Xeon E5-2697 v2@2.70GHz (IvyBridge) CPUs + Tesla M40 (autoboost) GPUs

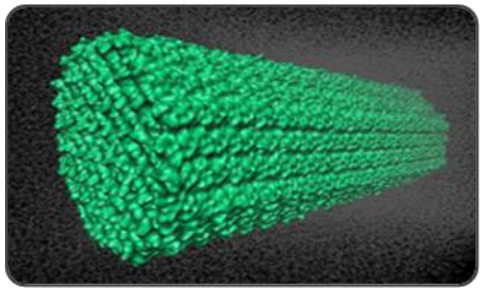

FactorIX on M40s

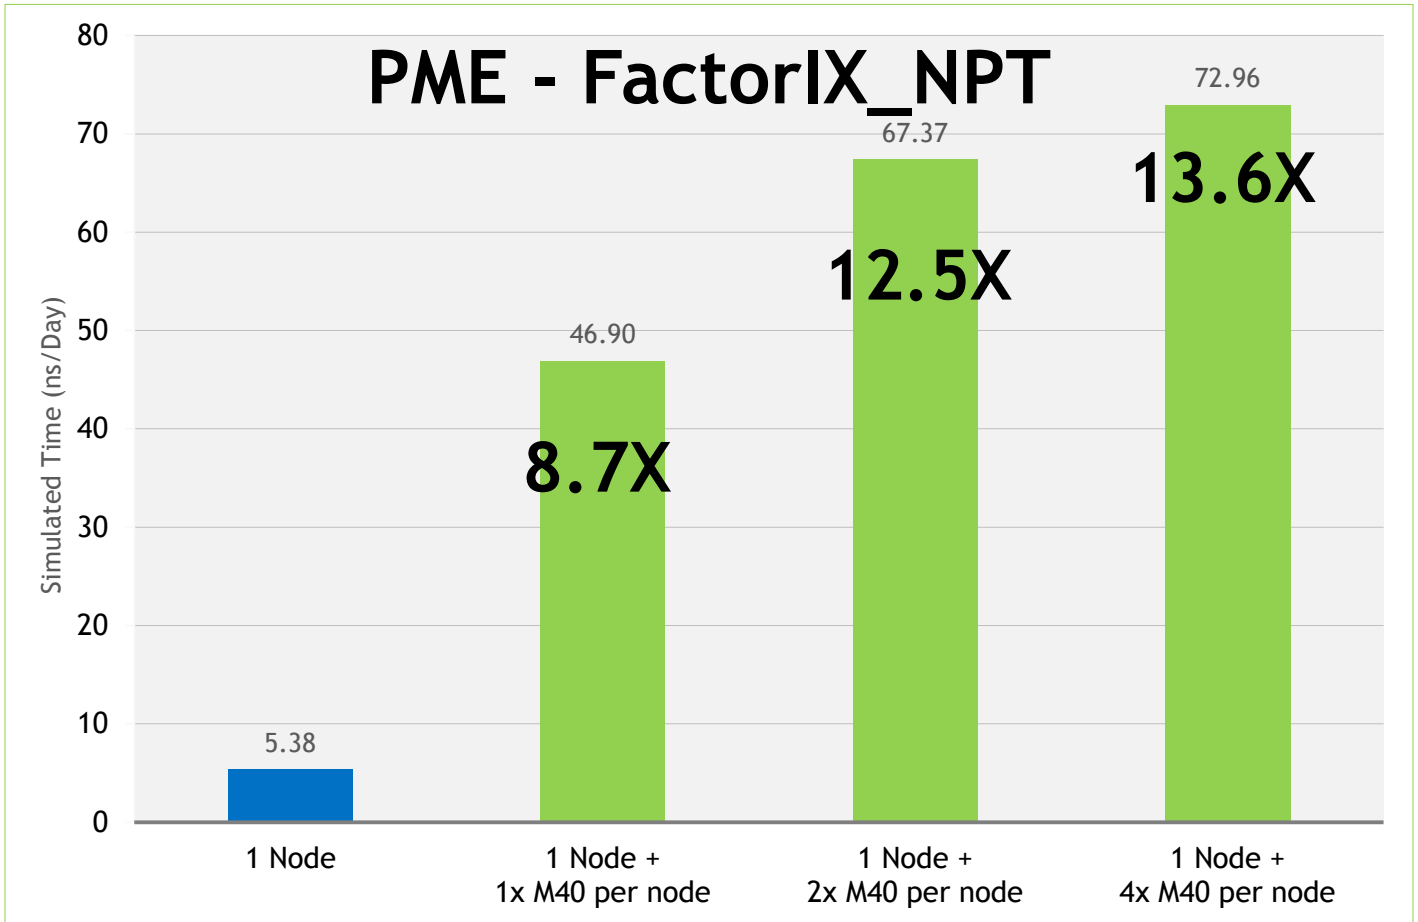

Running AMBER version 14

The blue node contain Single Intel Xeon E5-2698 v3@2.30GHz (Haswell) CPUs

The green nodes contain Single Intel Xeon E5-2697 v2@2.70GHz (IvyBridge) CPUs + Tesla M40 (autoboost) GPUs

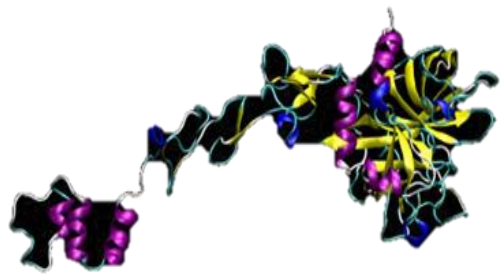

FactorIX on M40s

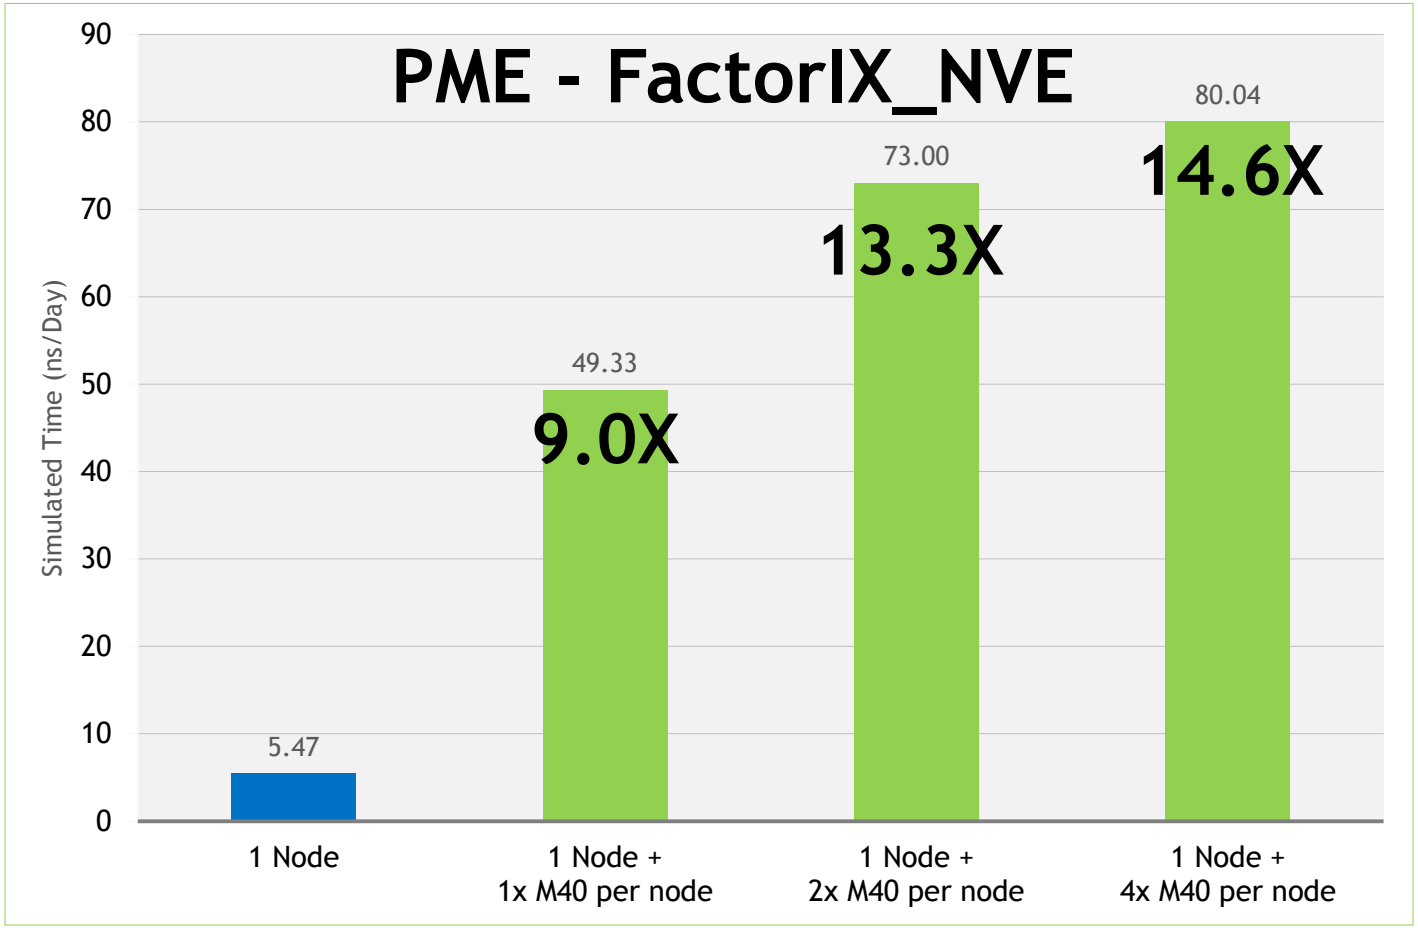

Running AMBER version 14

The blue node contain Single Intel Xeon E5-2698 v3@2.30GHz (Haswell) CPUs

The green nodes contain Single Intel Xeon E5-2697 v2@2.70GHz (IvyBridge) CPUs + Tesla M40 (autoboost) GPUs

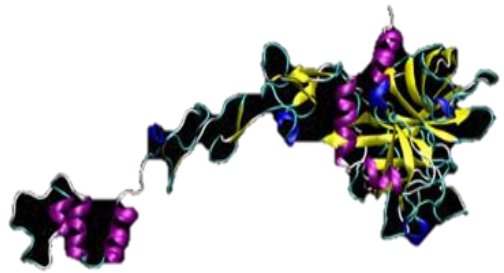

JAC on M40s

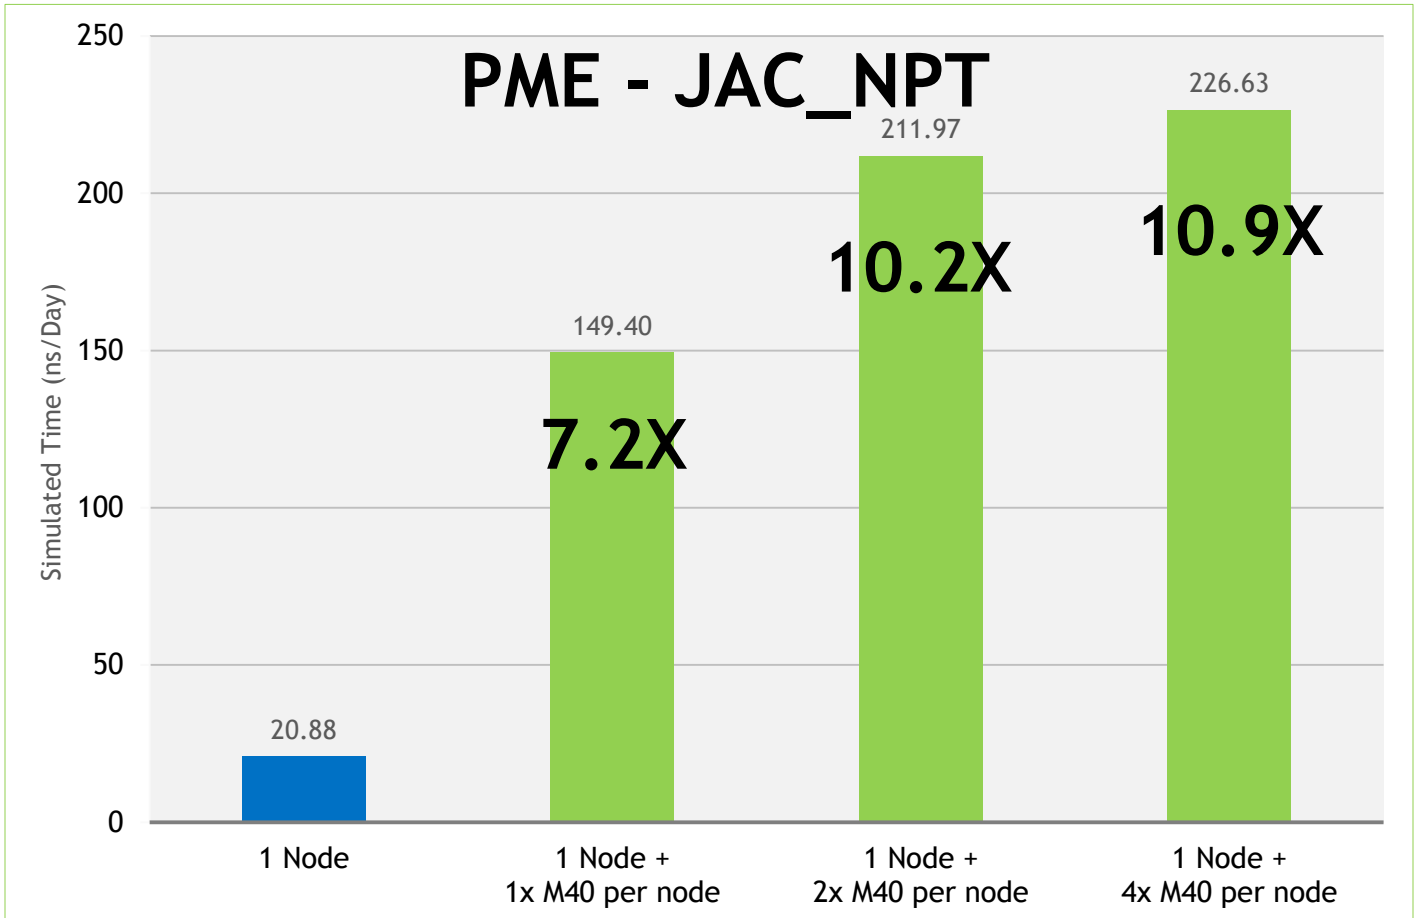

Running AMBER version 14

The blue node contain Single Intel Xeon E5-2698 v3@2.30GHz (Haswell) CPUs

The green nodes contain Single Intel Xeon E5-2697 v2@2.70GHz (IvyBridge) CPUs + Tesla M40 (autoboost) GPUs

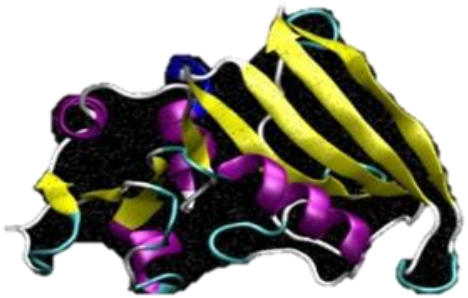

JAC on M40s

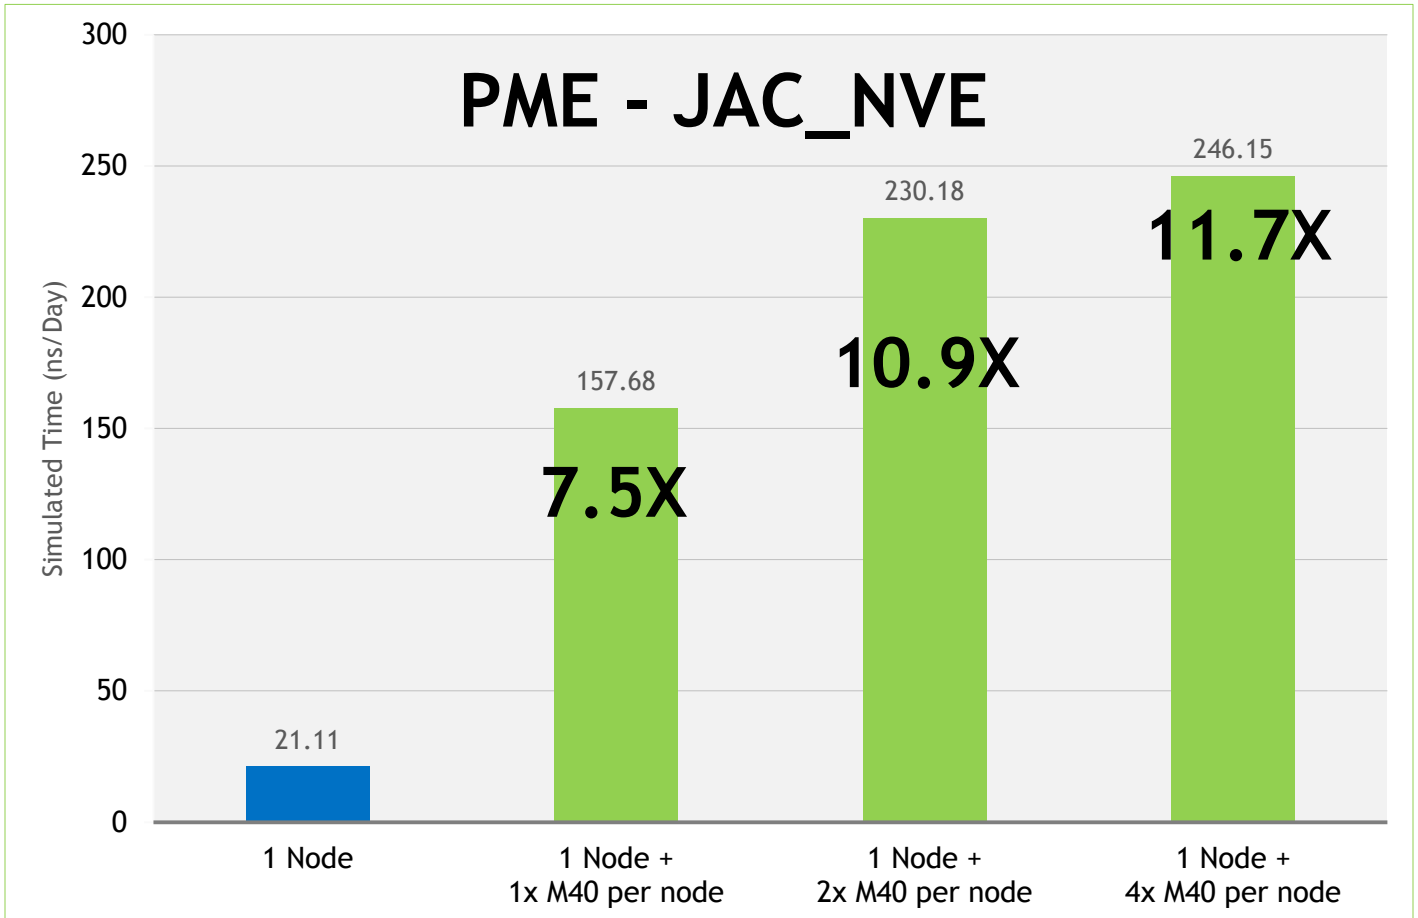

Running AMBER version 14

The blue node contain Single Intel Xeon E5-2698 v3@2.30GHz (Haswell) CPUs

The green nodes contain Single Intel Xeon E5-2697 v2@2.70GHz (IvyBridge) CPUs + Tesla M40 (autoboost) GPUs

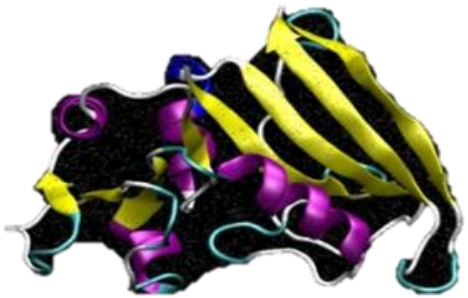

Myoglobin on M40s

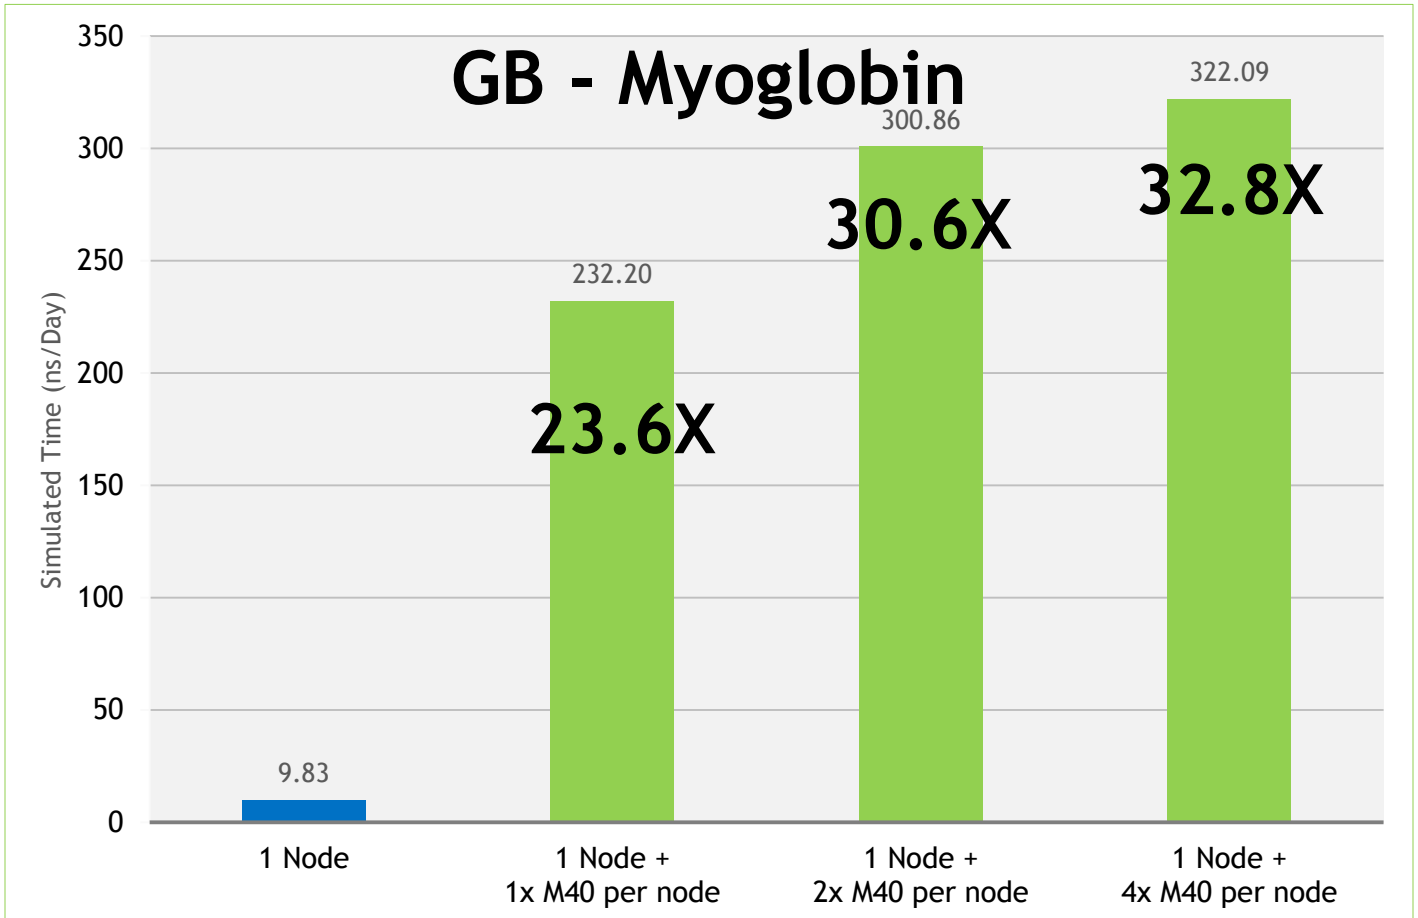

Running AMBER version 14

The blue node contain Single Intel Xeon E5-2698 v3@2.30GHz (Haswell) CPUs

The green nodes contain Single Intel Xeon E5-2697 v2@2.70GHz (IvyBridge) CPUs + Tesla M40 (autoboost) GPUs

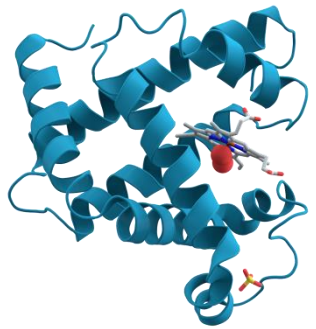

Nucleosome on M40s

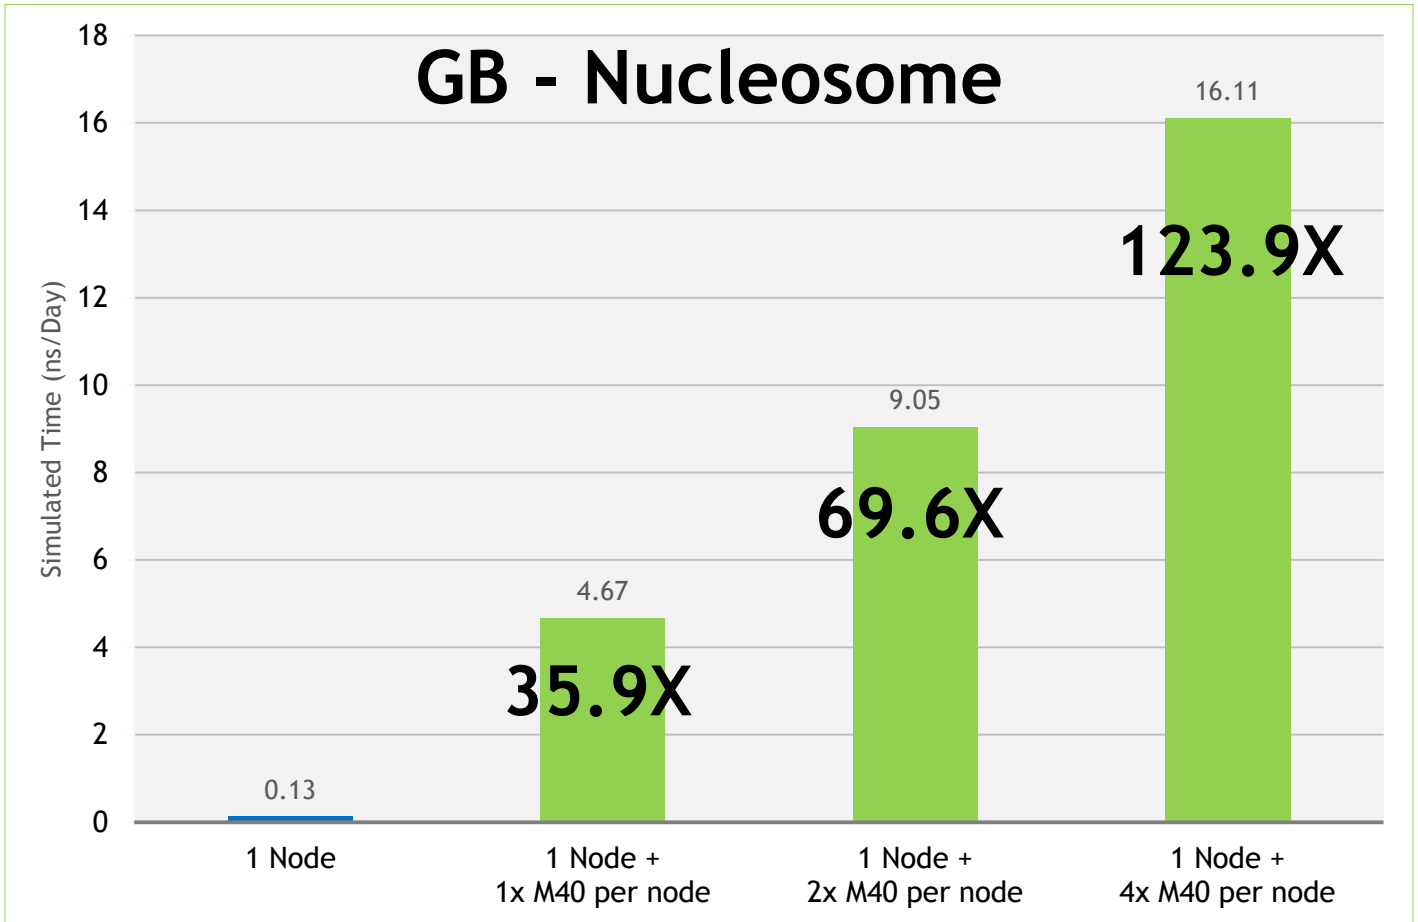

Running AMBER version 14

The blue node contain Single Intel Xeon E5-2698 v3@2.30GHz (Haswell) CPUs

The green nodes contain Single Intel Xeon E5-2697 v2@2.70GHz (IvyBridge) CPUs + Tesla M40 (autoboost) GPUs

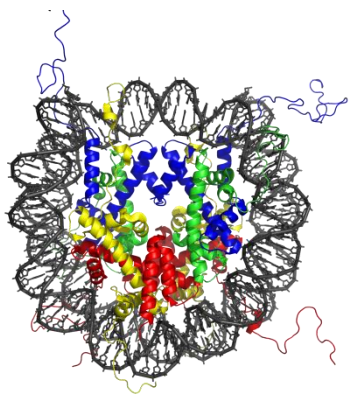

TrpCage on M40s

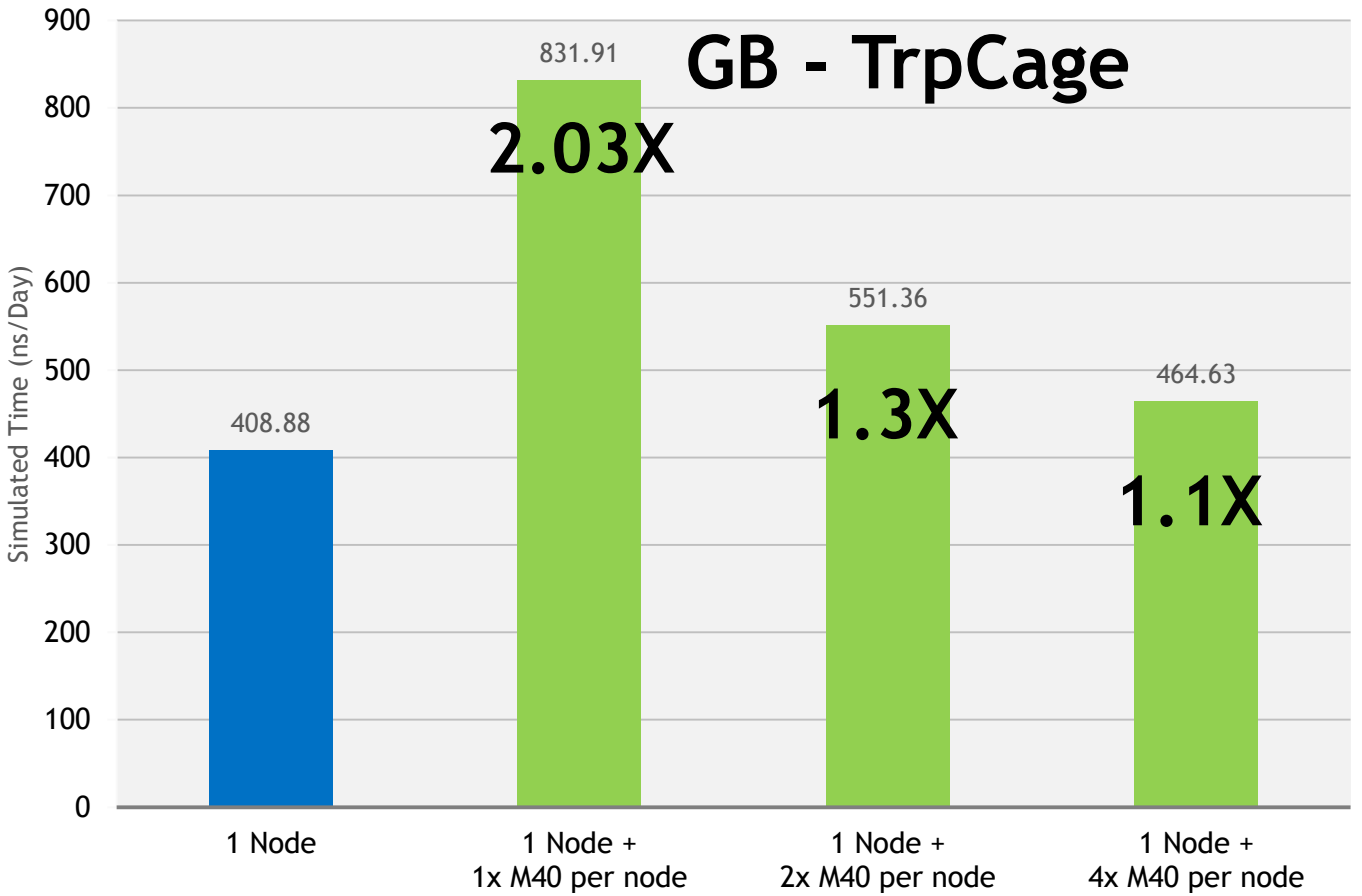

Running AMBER version 14

The blue node contain Single Intel Xeon E5-2698 v3@2.30GHz (Haswell) CPUs

The green nodes contain Single Intel Xeon E5-2697 v2@2.70GHz (IvyBridge) CPUs + Tesla M40 (autoboost) GPUs

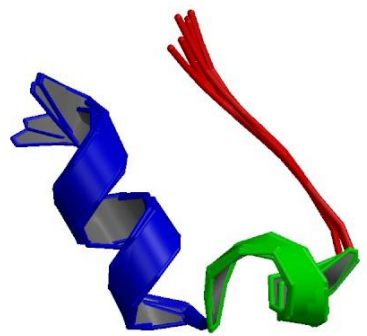

JAC on K40s and K80s

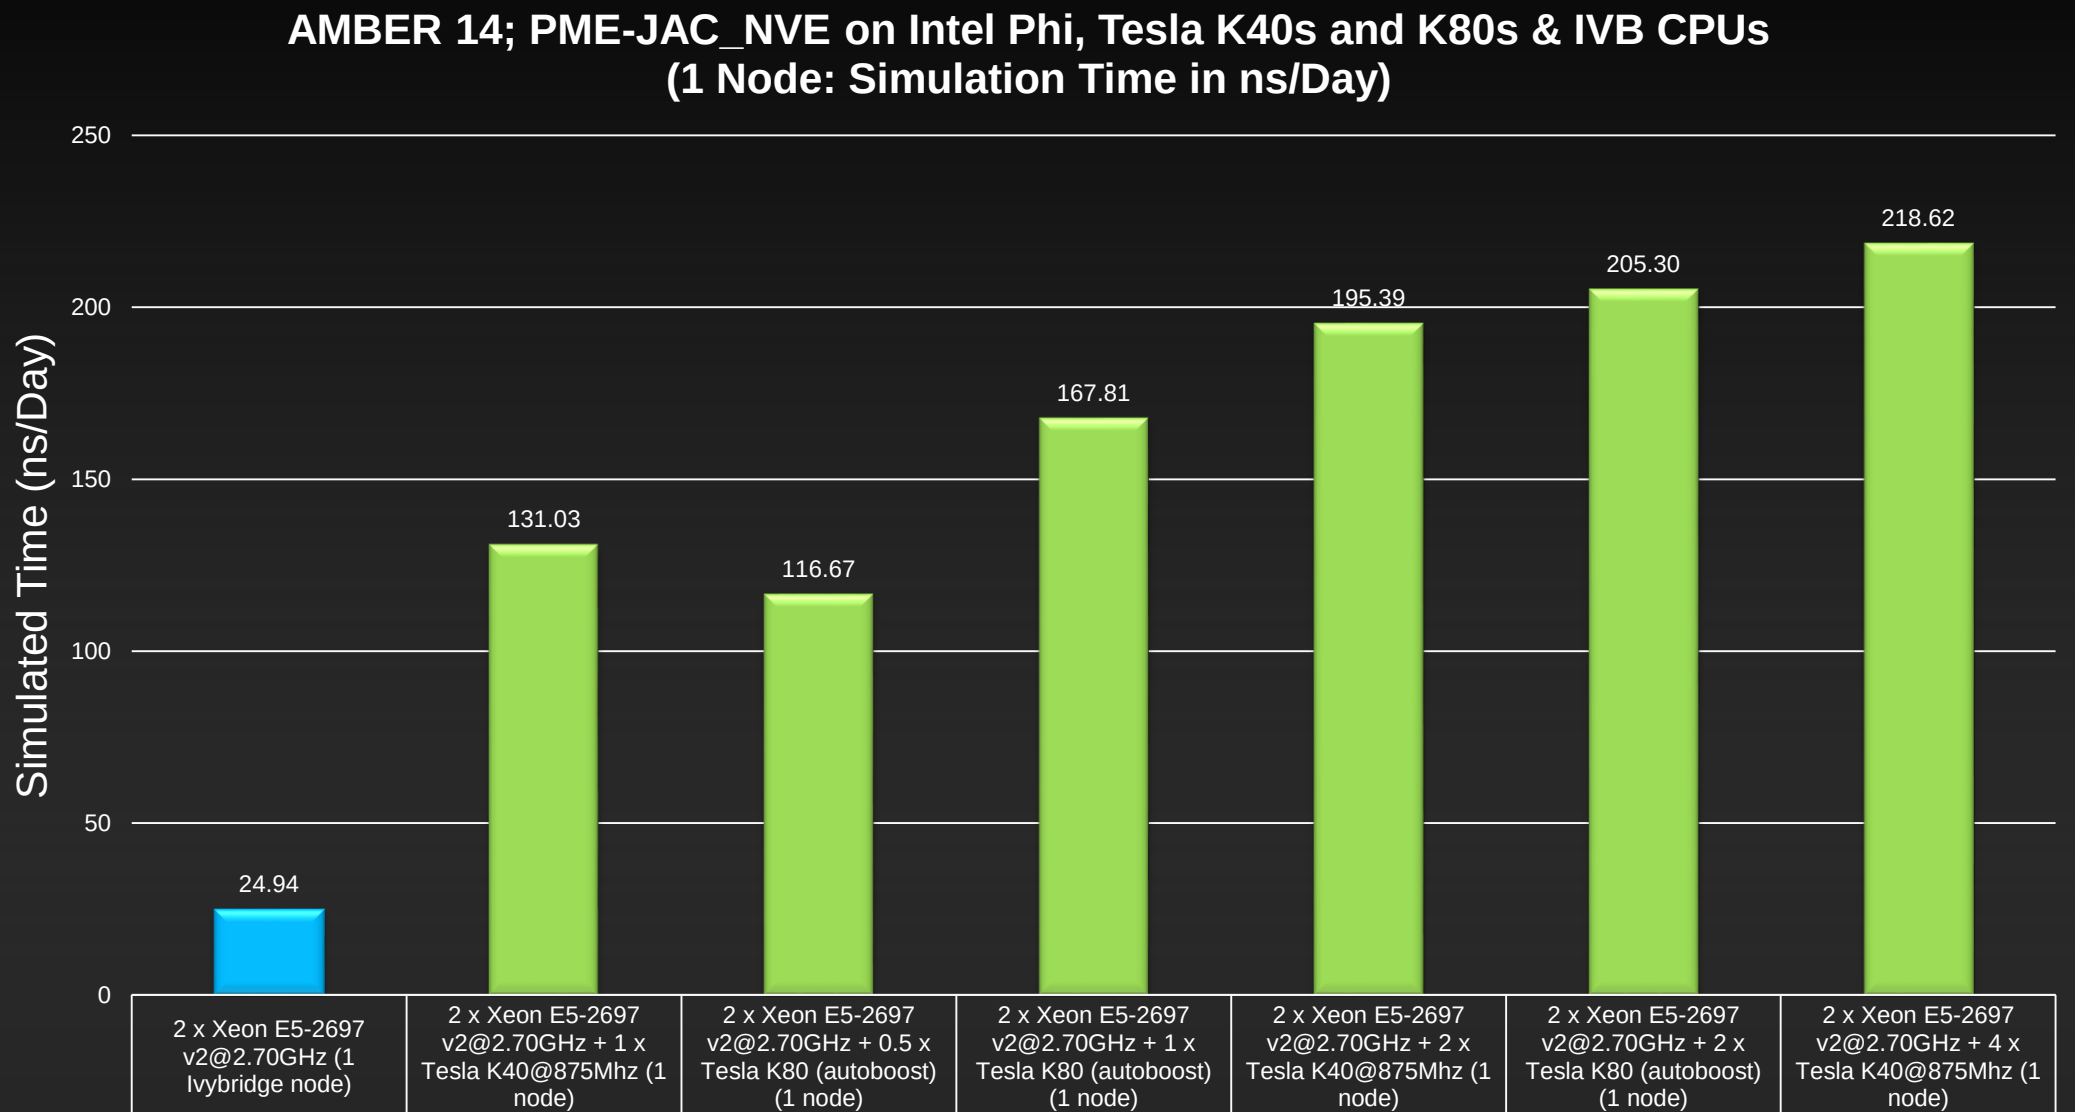

FactorIX on K40s and K80s

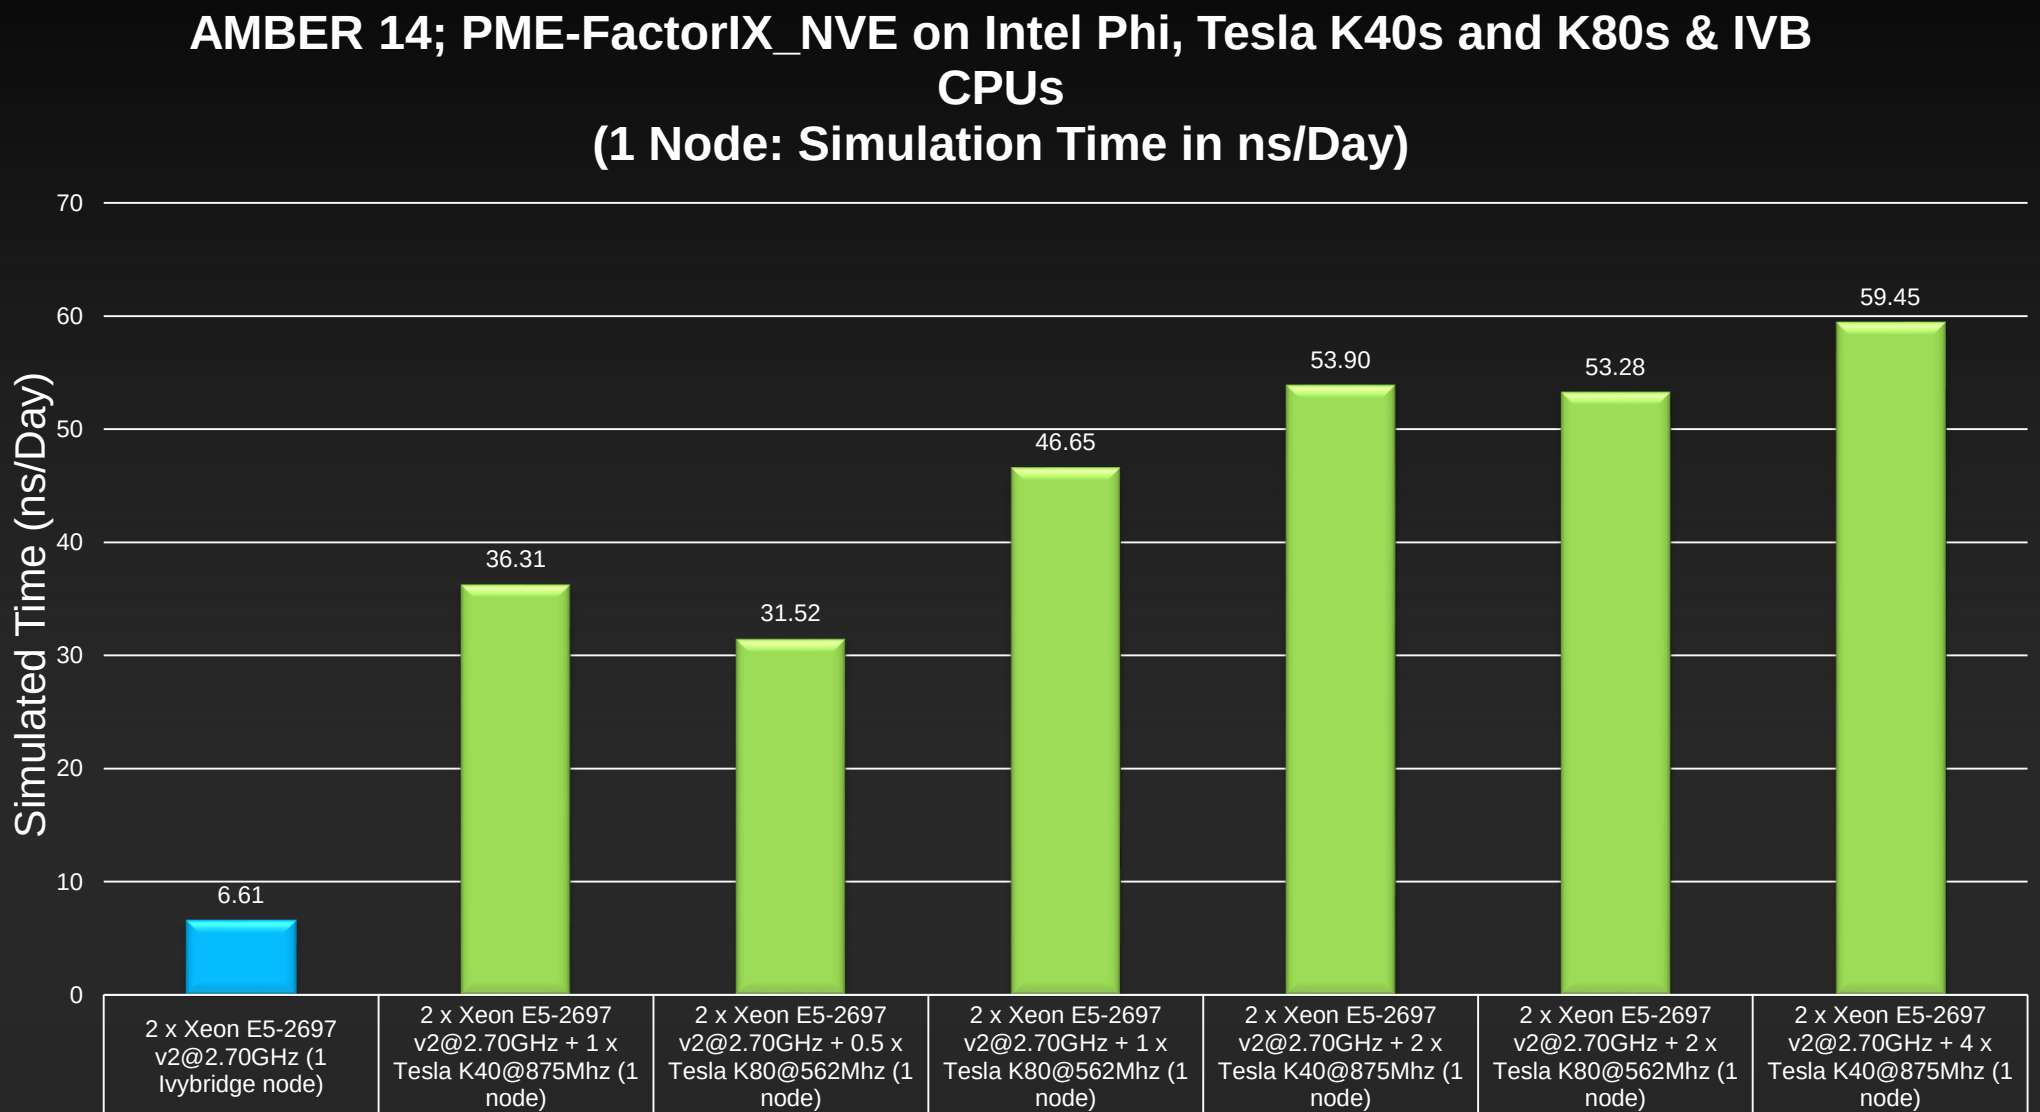

Cellulose on K40s and K80s

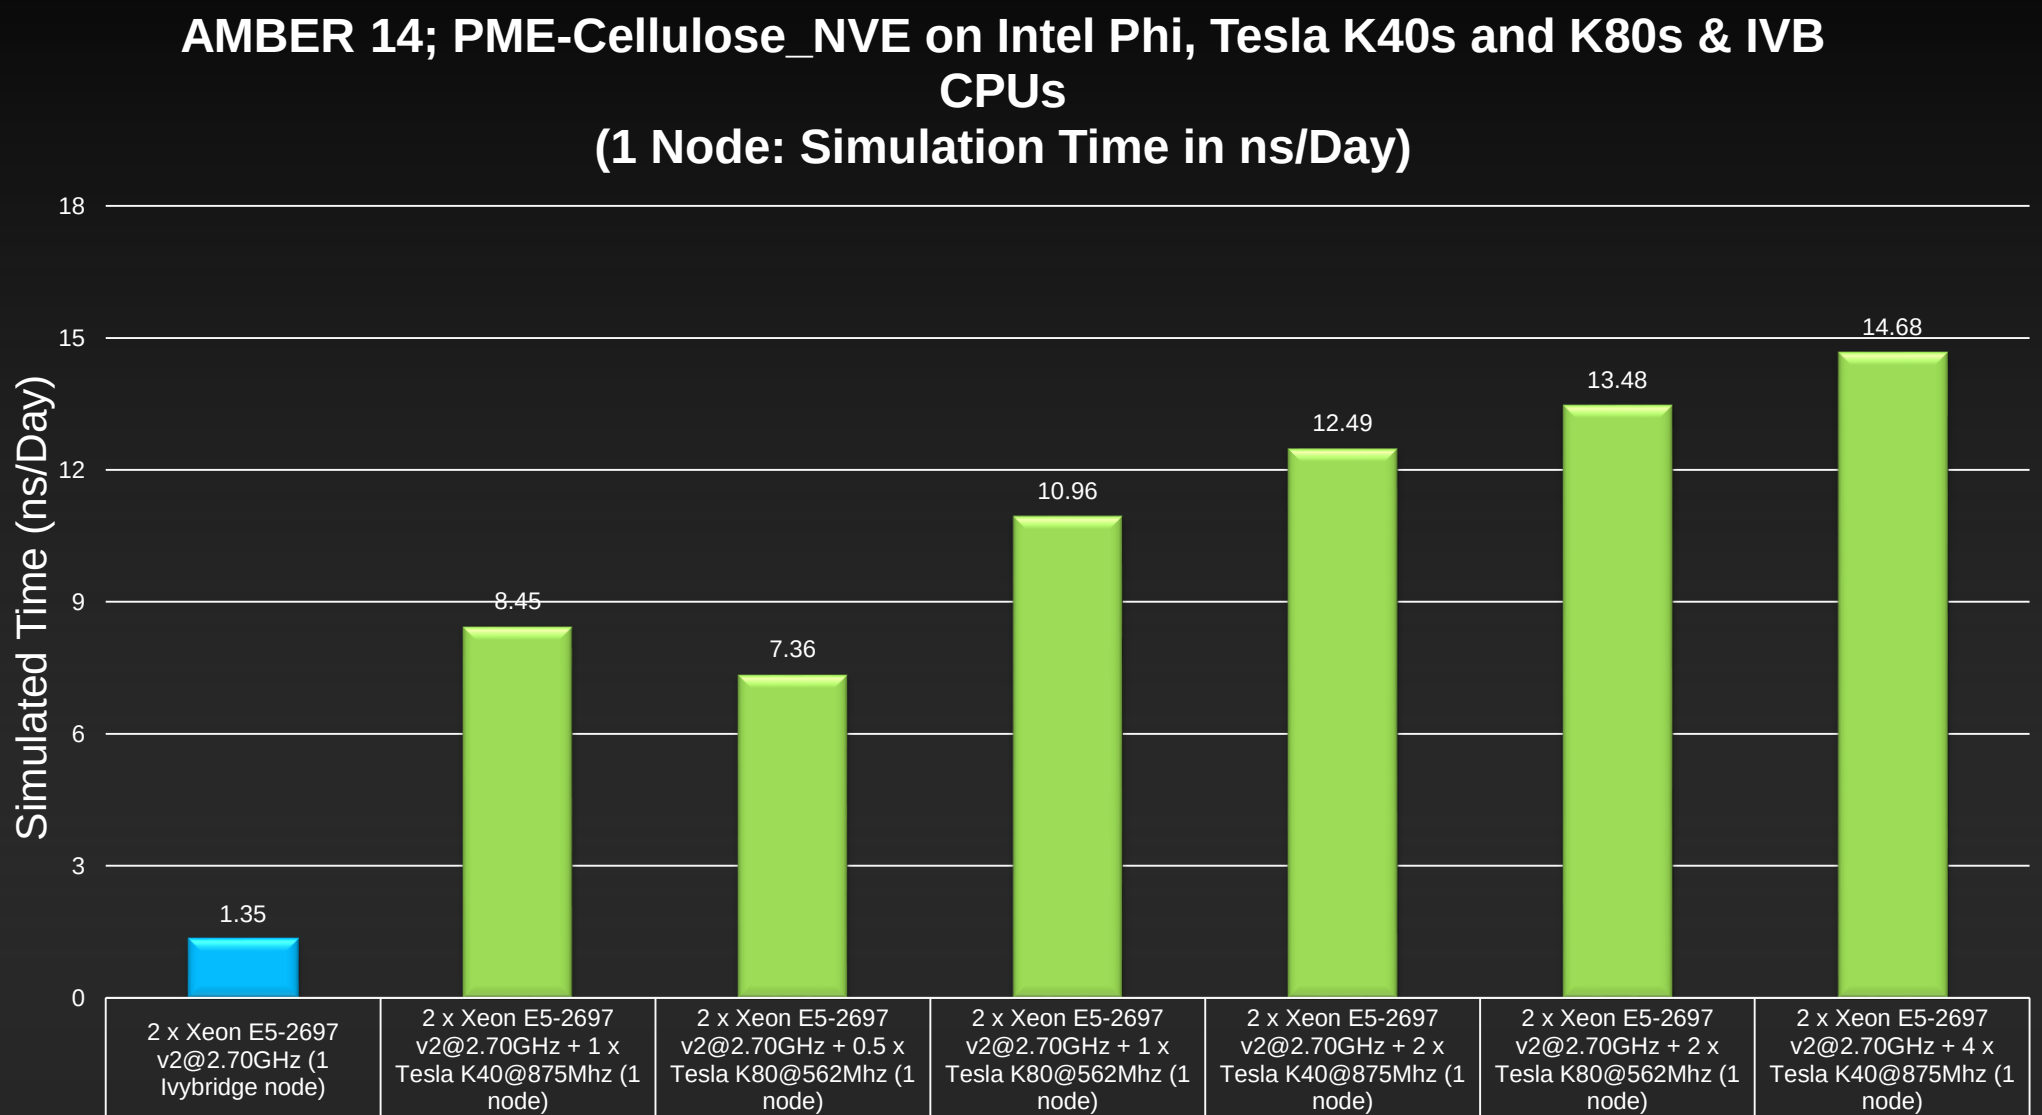

Kepler - Our Fastest Family of GPUs Yet

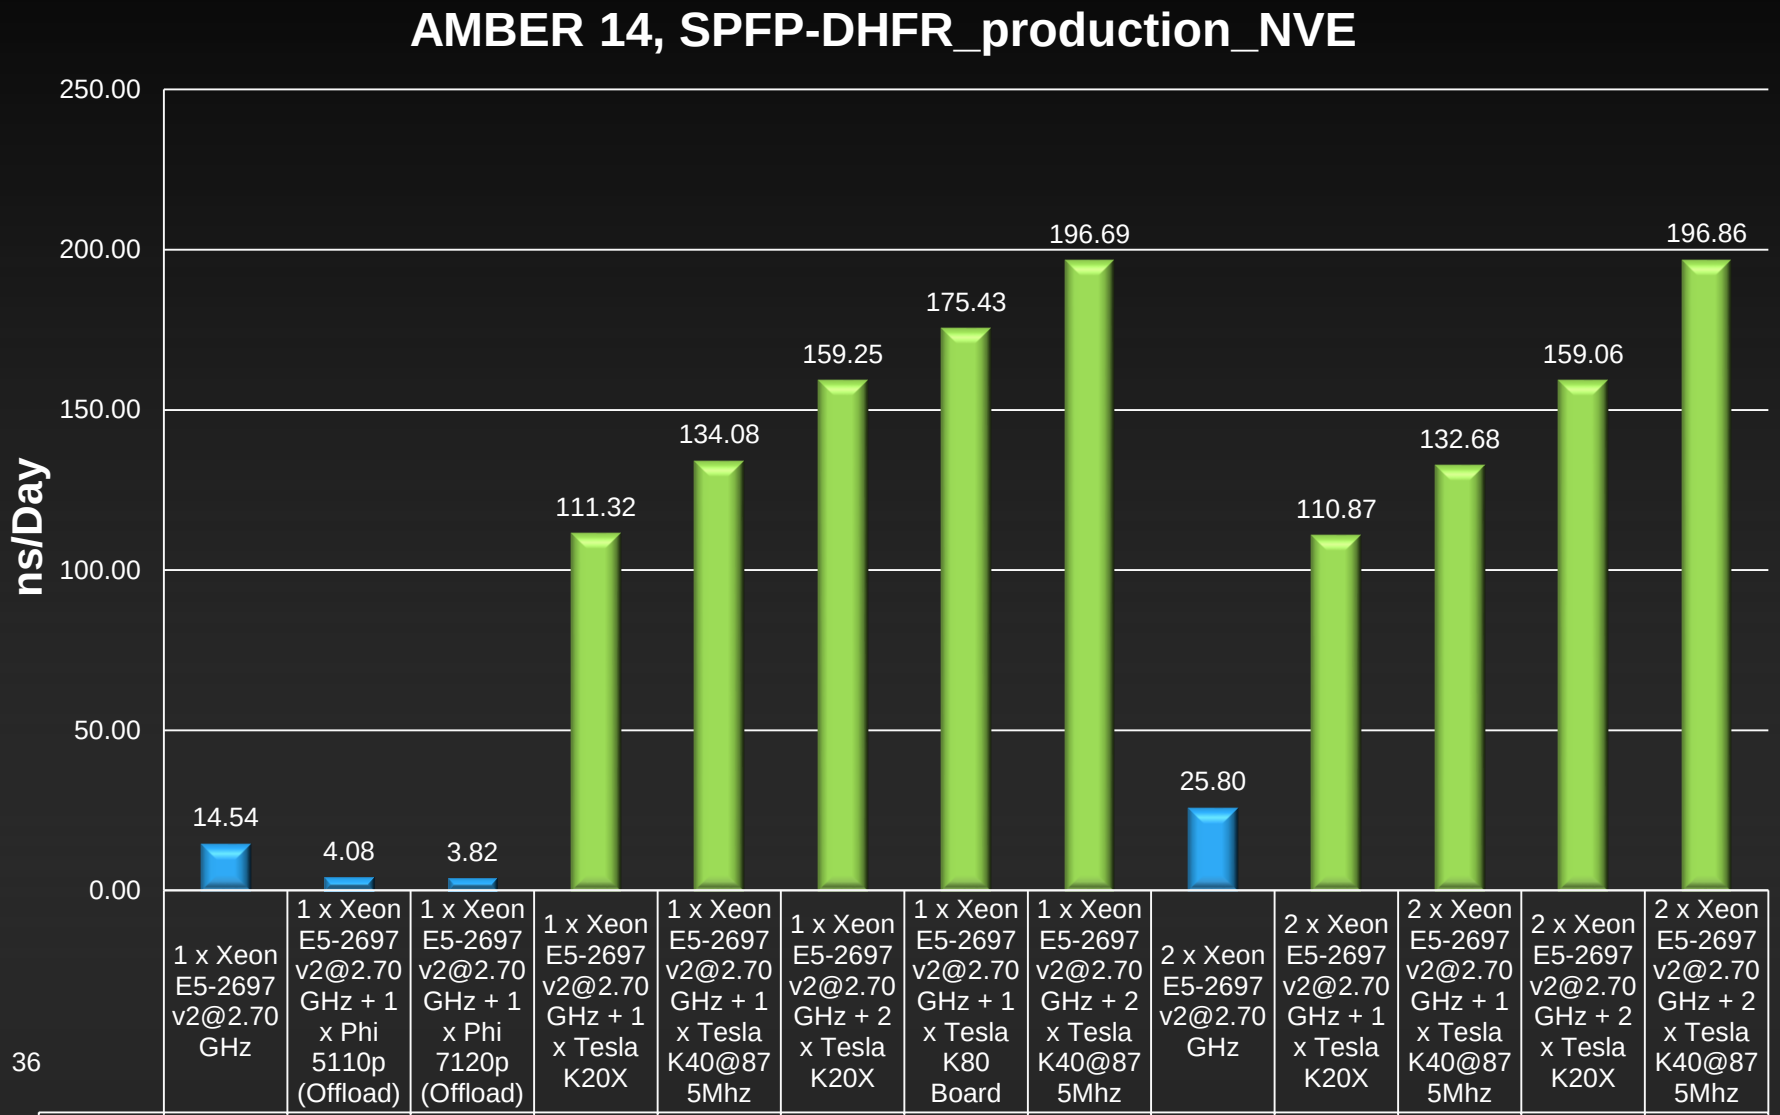

Running AMBER 14

The blue node contains Dual E5-2697 CPUs (12 Cores per CPU).

The green nodes contain Dual E5-2697 CPUs (12 Cores per CPU) and either 1x or 2x NVIDIA K20X, K40 or K80 for the GPU

Kepler - Our Fastest Family of GPUs Yet

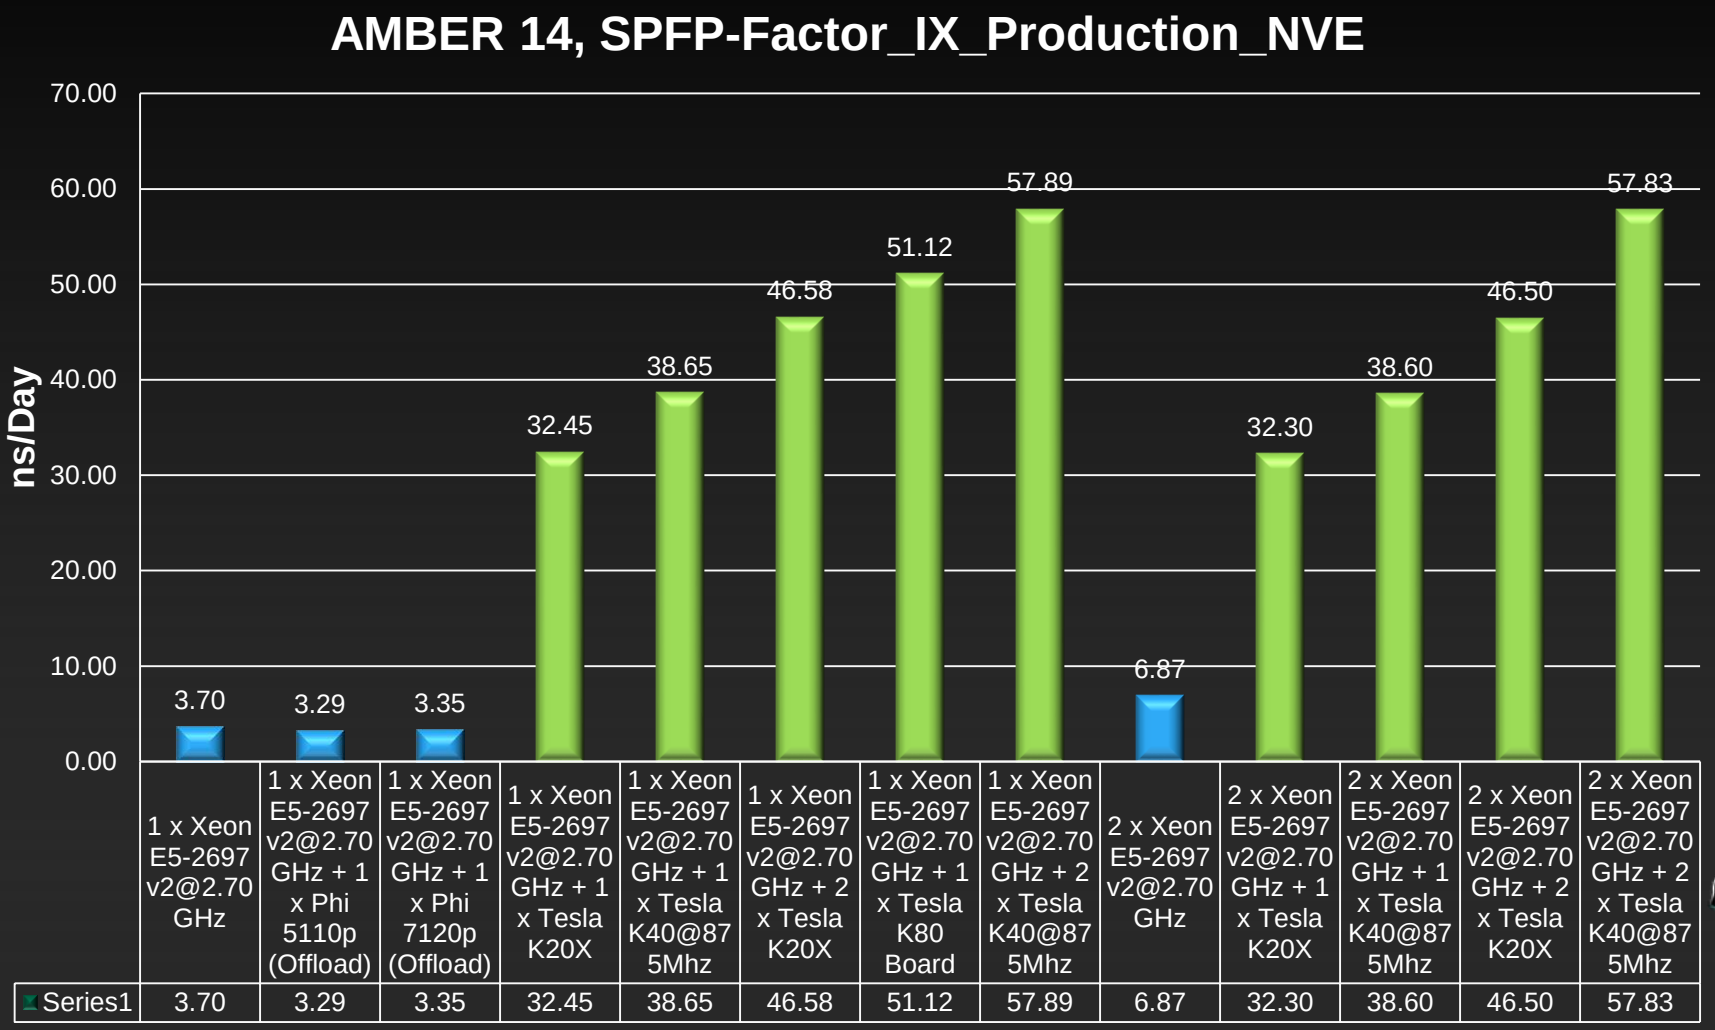

Running AMBER 14

The blue node contains Dual E5-2697 CPUs (12 Cores per CPU).

The green nodes contain Dual E5-2697 CPUs (12 Cores per CPU) and either 1x or 2x NVIDIA K20X, K40 or K80 for the GPU

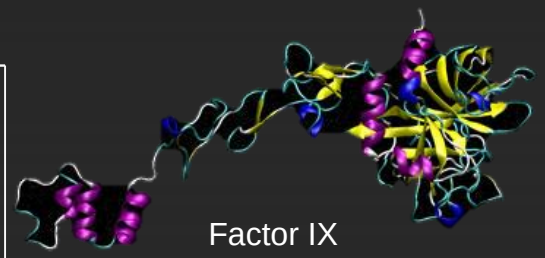

Kepler - Our Fastest Family of GPUs Yet

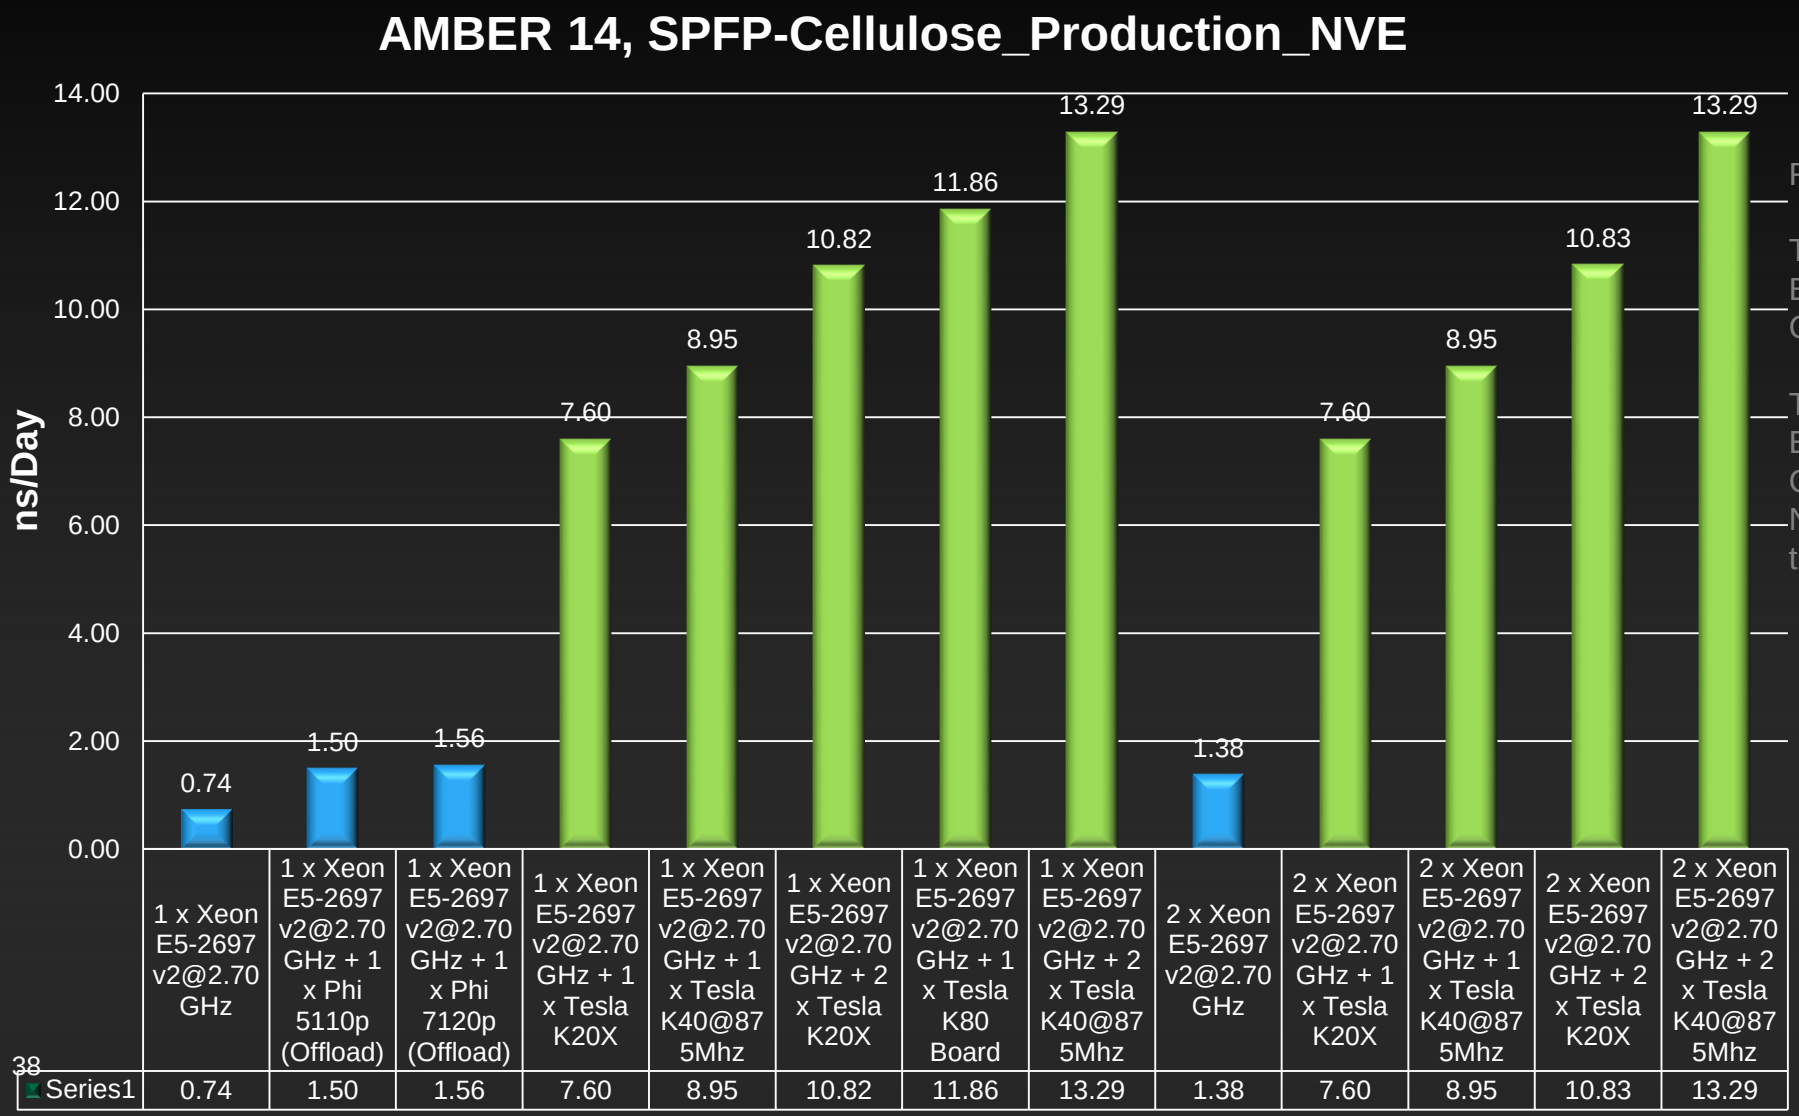

Running AMBER 14

The blue node contains Dual E5-2697 CPUs (12 Cores per CPU).

The green nodes contain Dual E5-2697 CPUs (12 Cores per CPU) and either 1x or 2x NVIDIA K20X, K40 or K80 for the GPU

Cost Comparison

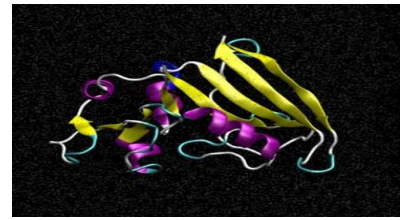

4 simultaneous simulations, 23,000 atoms, 250ns each, 5 days maximum time to solution.

	Traditional Cluster	GPU Workstation
Nodes Required	12	1 (4 GPUs)
Interconnect	QDR IB	None
Time to complete simulations	4.98 days	2.25 days
Power Consumption	5.7 kW (681.3 kWh)	1.0 kW (54.0 kWh)
System Cost (per day)	$96,800 ($88.40)	$5200 ($4.75)
Simulation Cost	(681.3 * 0.18) + (88.40 * 4.98)	(54.0 * 0.18) + (4.75 * 2.25)
	$562.87	***$20.41***

>25x cheaper AND solution obtained in less than half the time

Replace 8 Nodes with 1 K20 GPU

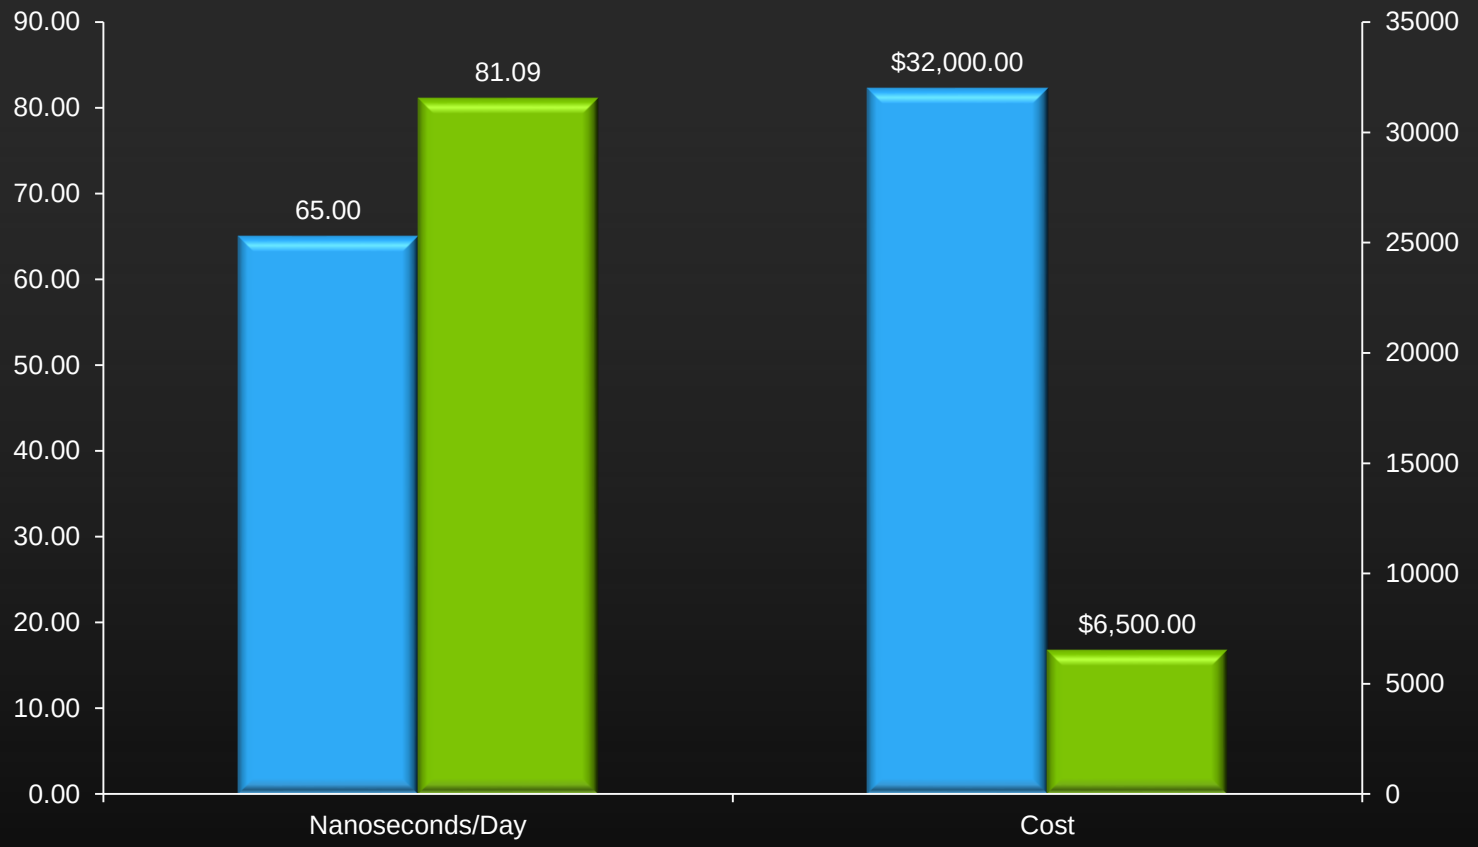

Running AMBER 12 GPU Support Revision 12.1 SPFP with CUDA 4.2.9 ECC Off

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

Each green node contains 2x Intel E5-2687W CPUs (8 Cores per CPU) plus 1x NVIDIA K20 GPU

Note: Typical CPU and GPU node pricing used. Pricing may vary depending on node configuration. Contact your preferred HW vendor for actual pricing.

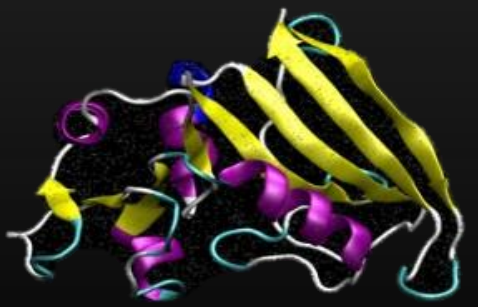

DHFR

Cut down simulation costs to ¼ and gain higher performance

Replace 7 Nodes with 1 K10 GPU

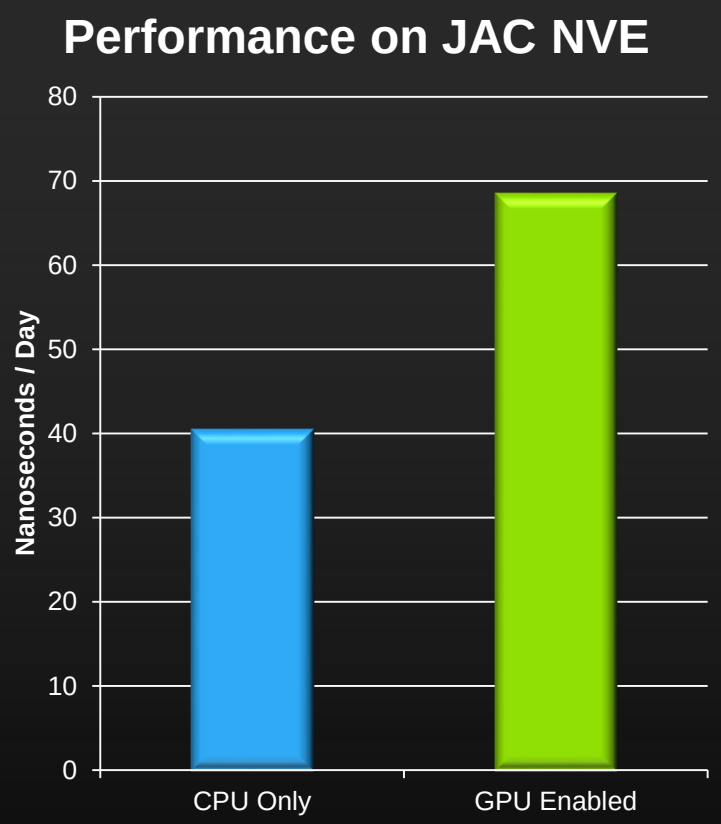

Running AMBER 12 GPU Support Revision 12.1 SPFP with CUDA 4.2.9 ECC Off

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

The green node contains 2x Intel E5-2687W CPUs (8 Cores per CPU) plus 1x NVIDIA K10 GPU

Note: Typical CPU and GPU node pricing used. Pricing may vary depending on node configuration. Contact your preferred HW vendor for actual pricing.

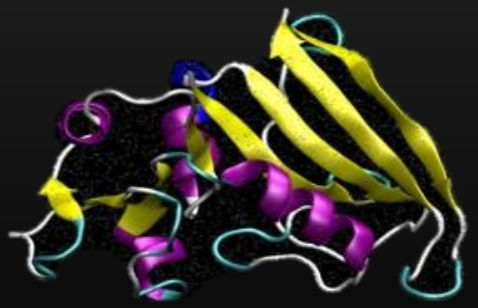

DHFR

Cut down simulation costs to ¼ and increase performance by 70%

Extra CPUs decrease Performance

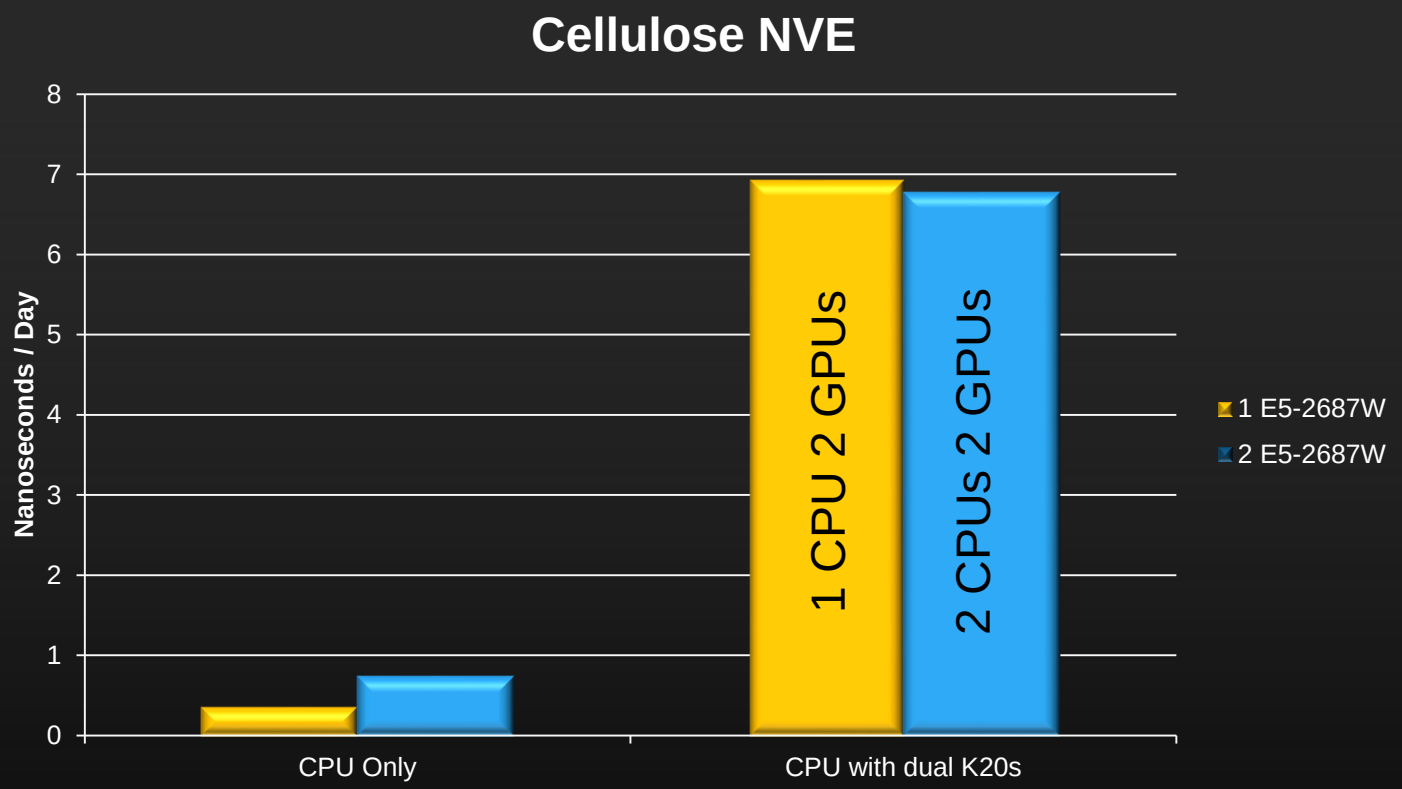

Running AMBER 12 GPU Support Revision 12.1

The orange bars contains one E5-2687W CPUs (8 Cores per CPU).

The blue bars contain Dual E5-2687W CPUs (8 Cores per CPU)

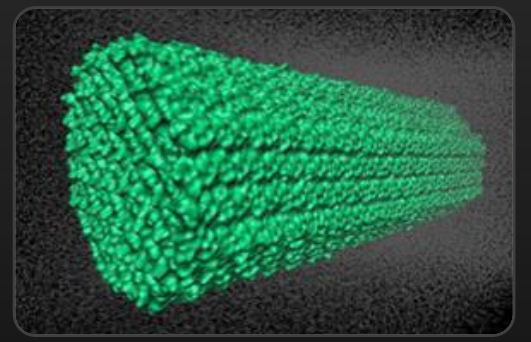

Cellulose

When used with GPUs, dual CPU sockets perform worse than single CPU sockets.

Kepler - Greener Science

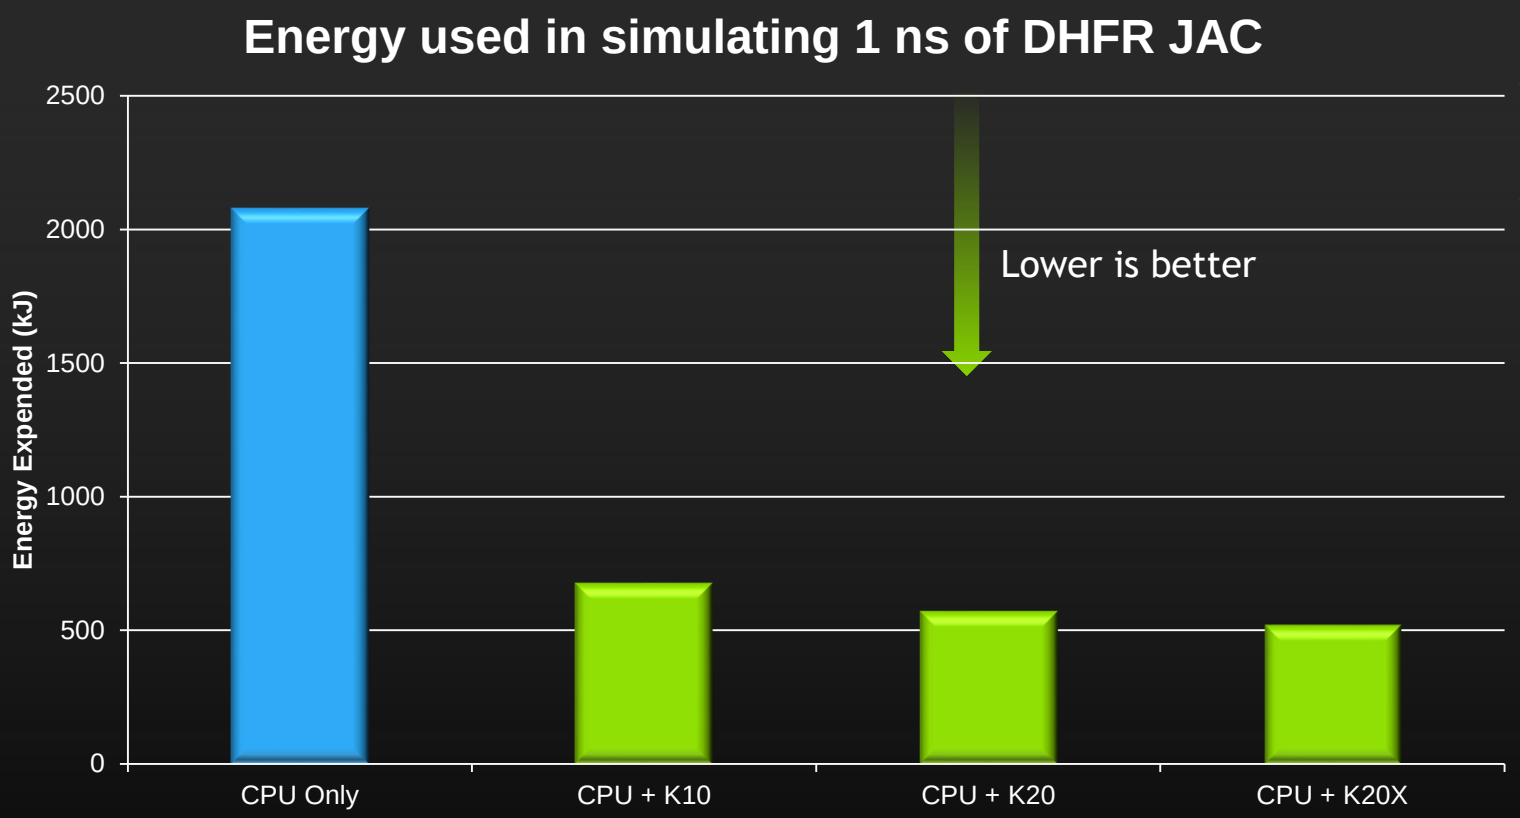

The GPU Accelerated systems use 65-75% less energy

Running AMBER 12 GPU Support Revision 12.1

The blue node contains Dual E5-2687W CPUs (150W each, 8 Cores per CPU).

The green nodes contain Dual E5-2687W CPUs (8 Cores per CPU) and 1x NVIDIA K10, K20, or K20X GPUs (235W each).

Energy Expended = Power x Time

Recommended GPU Node Configuration for AMBER Computational Chemistry

Workstation or Single Node Configuration	
# of CPU sockets	2
Cores per CPU socket	6+ (1 CPU core drives 1 GPU)
CPU speed (Ghz)	2.66+
System memory per node (GB)	16
GPUs	Kepler K20, K40, K80
# of GPUs per CPU socket	1-4
GPU memory preference (GB)	6
GPU to CPU connection	PCIe 3.0 16x or higher
Server storage	2 TB
Network configuration	Infiniband QDR or better

Scale to multiple nodes with same single node configuration

CHARMM

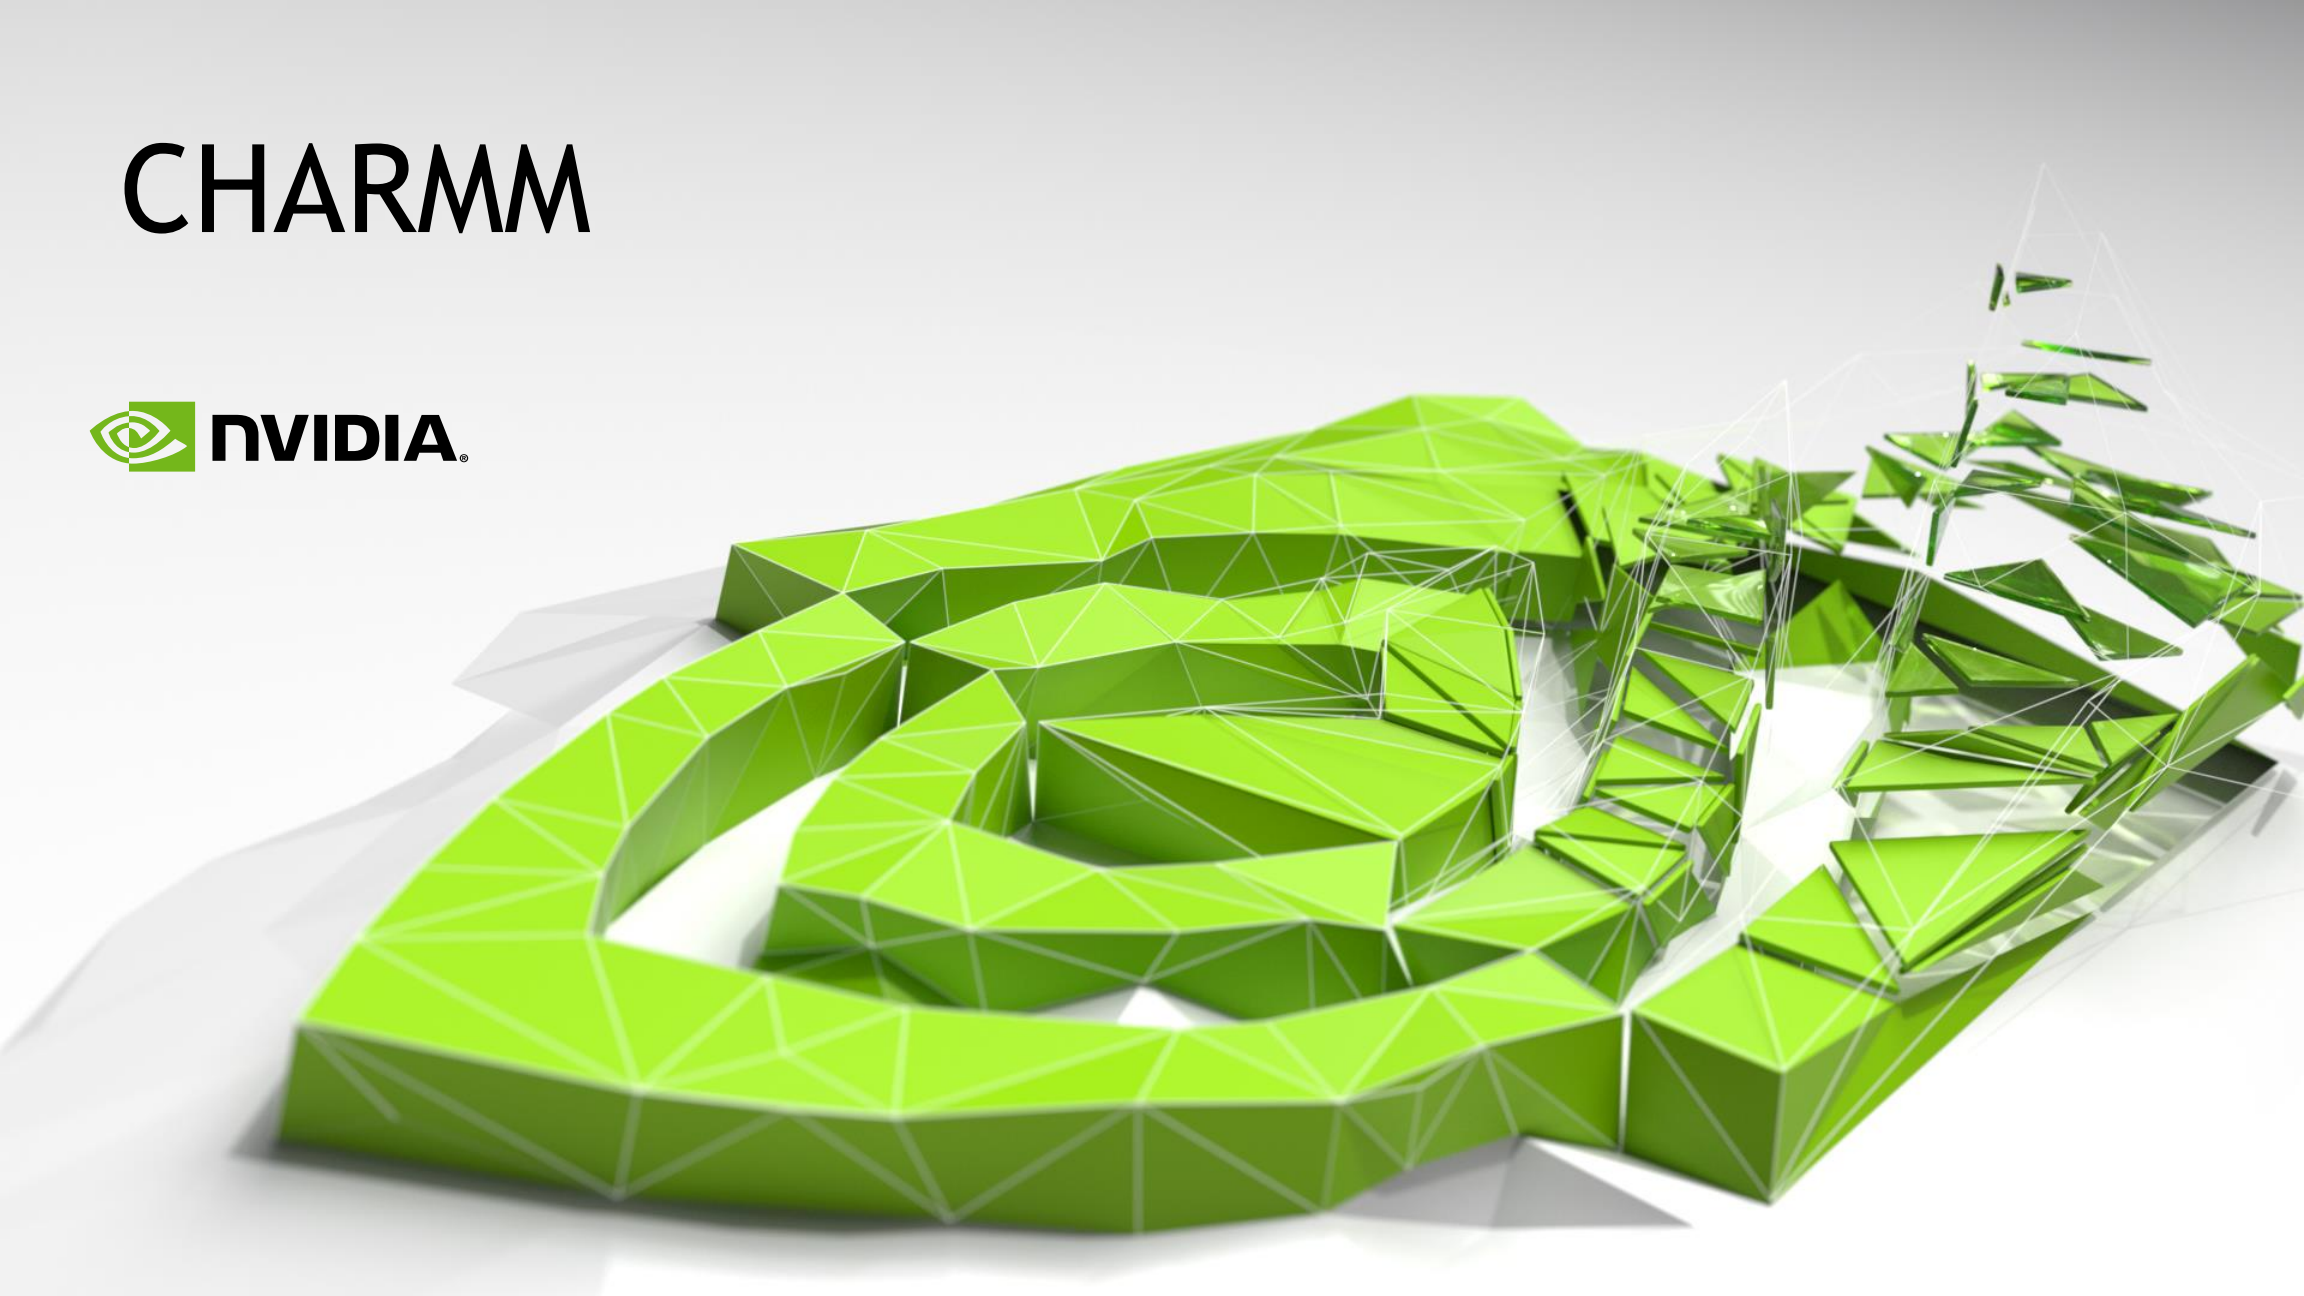

NREL
NATIONAL RENEWABLE ENERGY LABORATORY

NVIDIA

Heterogeneous CPU+GPU MD Engine in CHARMM


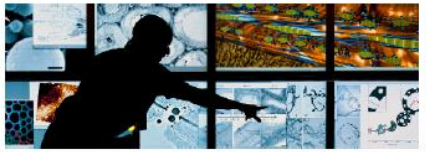

CHARMM meeting 2013, Zurich

Antti-Pekka Hynninen

antti.pekka.hynninen@nrel.gov

Courtesy of Antti-Pekka Hynninen @ NREL

DOMDEC_GPU MD step cycle for PME

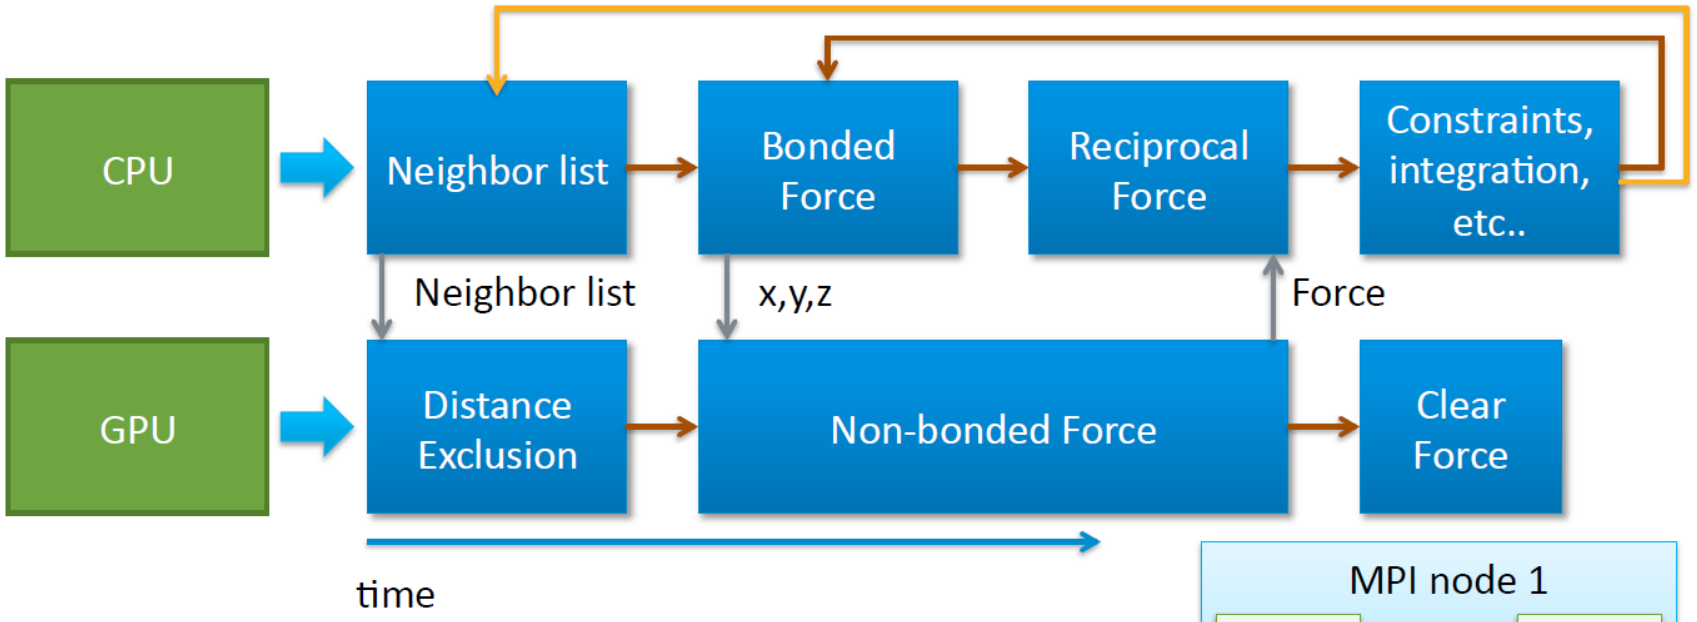

- **GPU calculates non-bonded force**
- **CPU calculates everything else**
- **Fully integrated with DOMDEC**
  - MPI task = CPU + GPU -pair
  - Built-in multinode GPU support

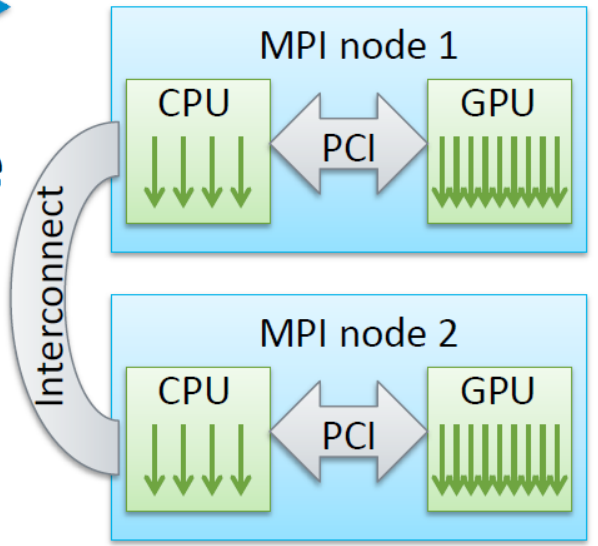

Courtesy of Antti-Pekka Hynninen @ NREL

Performance for DHFR on Tesla C2075*

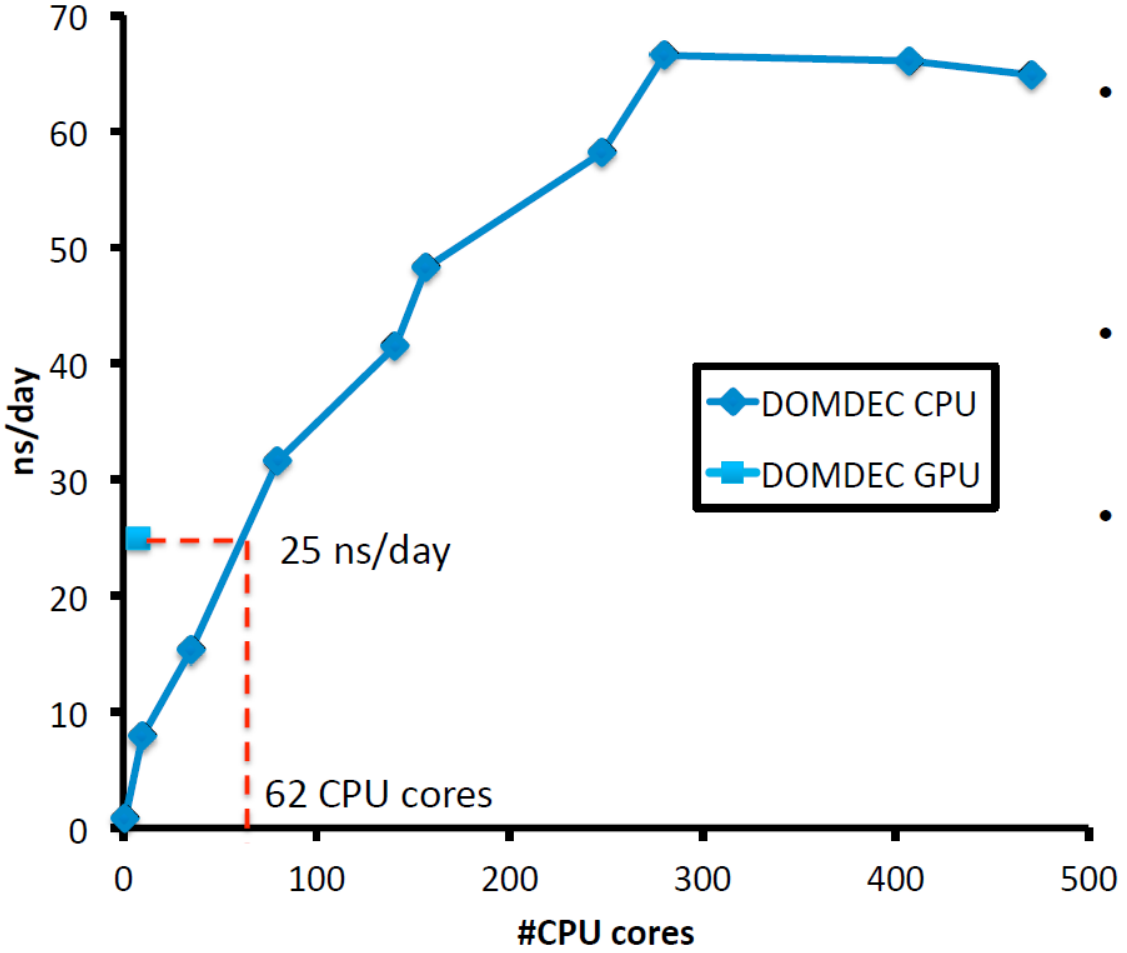

- Single node run using DOMDEC GPU corresponds to 62 CPU cores = 7-8 CPU nodes
- Great for the user, only single node taken for the run!
- Lots of room for improvement: Amber, Gromacs 4.6, and OpenMM run at about 50ns/day on a similar system

Courtesy of Antti-Pekka Hynninen @ NREL

*23558 atoms. Profiling run on 8 core Intel Xeon X5667 + NVIDIA Tesla C2075

Conclusions & Future Work

Conclusions

- Heterogeneous CPU + GPU MD engine implemented for CHARMM
  - Version 0.1
- Reasonable initial performance
- Natively multi-node capable
  - Scalable GPU computing
- OpenMP threads in CPU code
  - Important for current and future CPU architectures

Acknowledgements:
Michael Crowley, NREL
Charles Brooks, University of Michigan

Future Work

- Optimize (both CPU and GPU)
- Test & optimize multi-node simulation runs
- Load balancing PME calculation
- Looking for beta testers!
  - Need a relatively modern NVIDIA GPU computing node (compute capability 2.0 and above)

NVIDIA

Courtesy of Antti-Pekka Hynninen @ NREL

Greener Science with NVIDIA

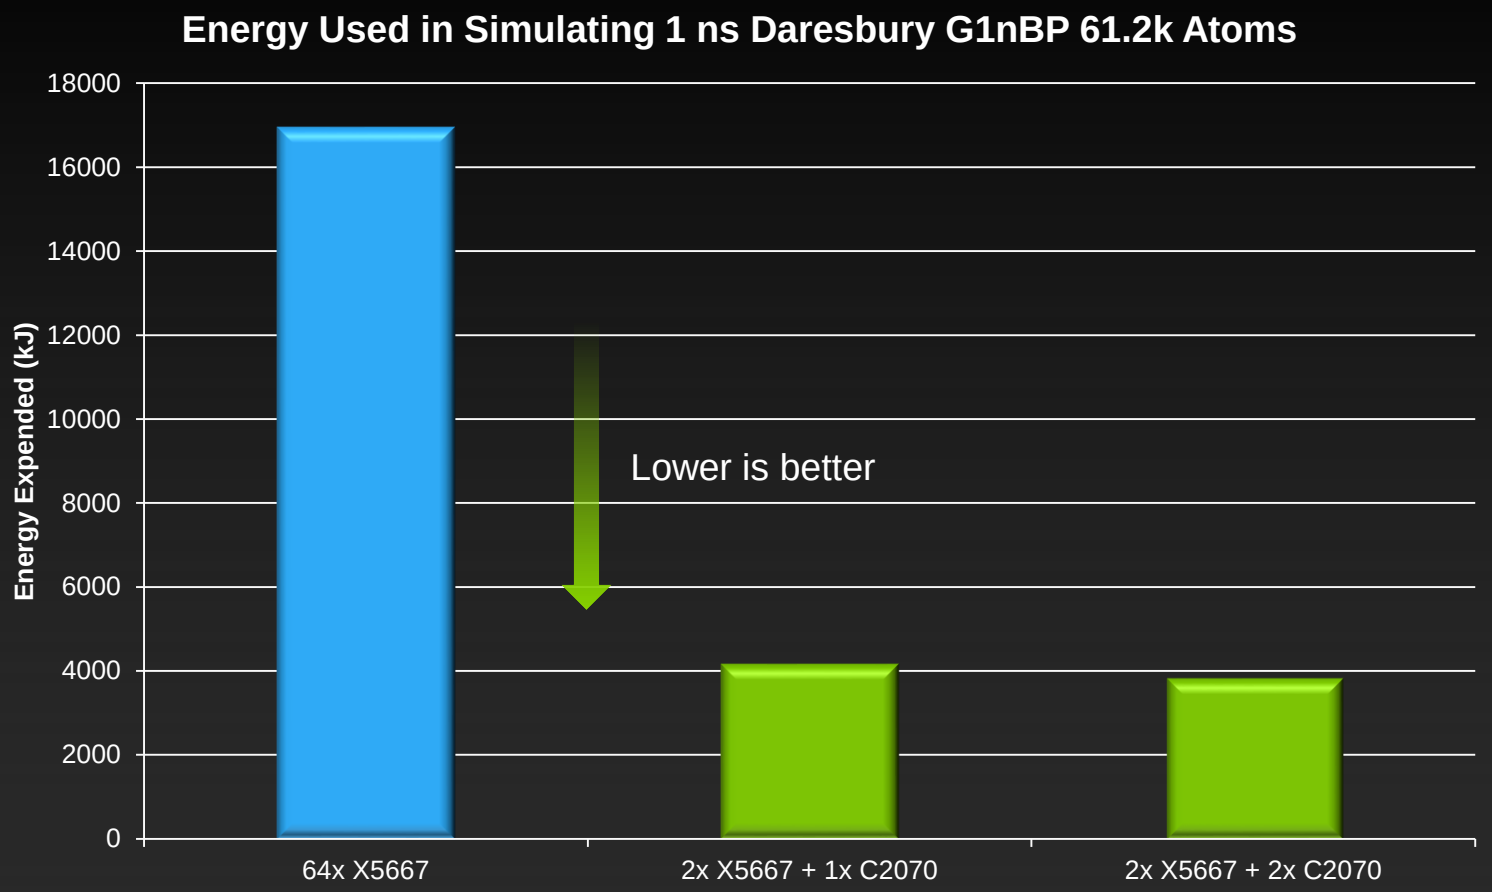

Using GPUs will decrease energy use by 75%

Running CHARMM release C37b1

The blue nodes contains 64 X5667 CPUs (95W, 4 Cores per CPU).

The green nodes contain 2 X5667 CPUs and 1 or 2 NVIDIA C2070 GPUs (238W each).

Note: Typical CPU and GPU node pricing used. Pricing may vary depending on node configuration. Contact your preferred HW vendor for actual pricing.

Energy Expended
= Power x Time

CHARMM c40a2

May 2016

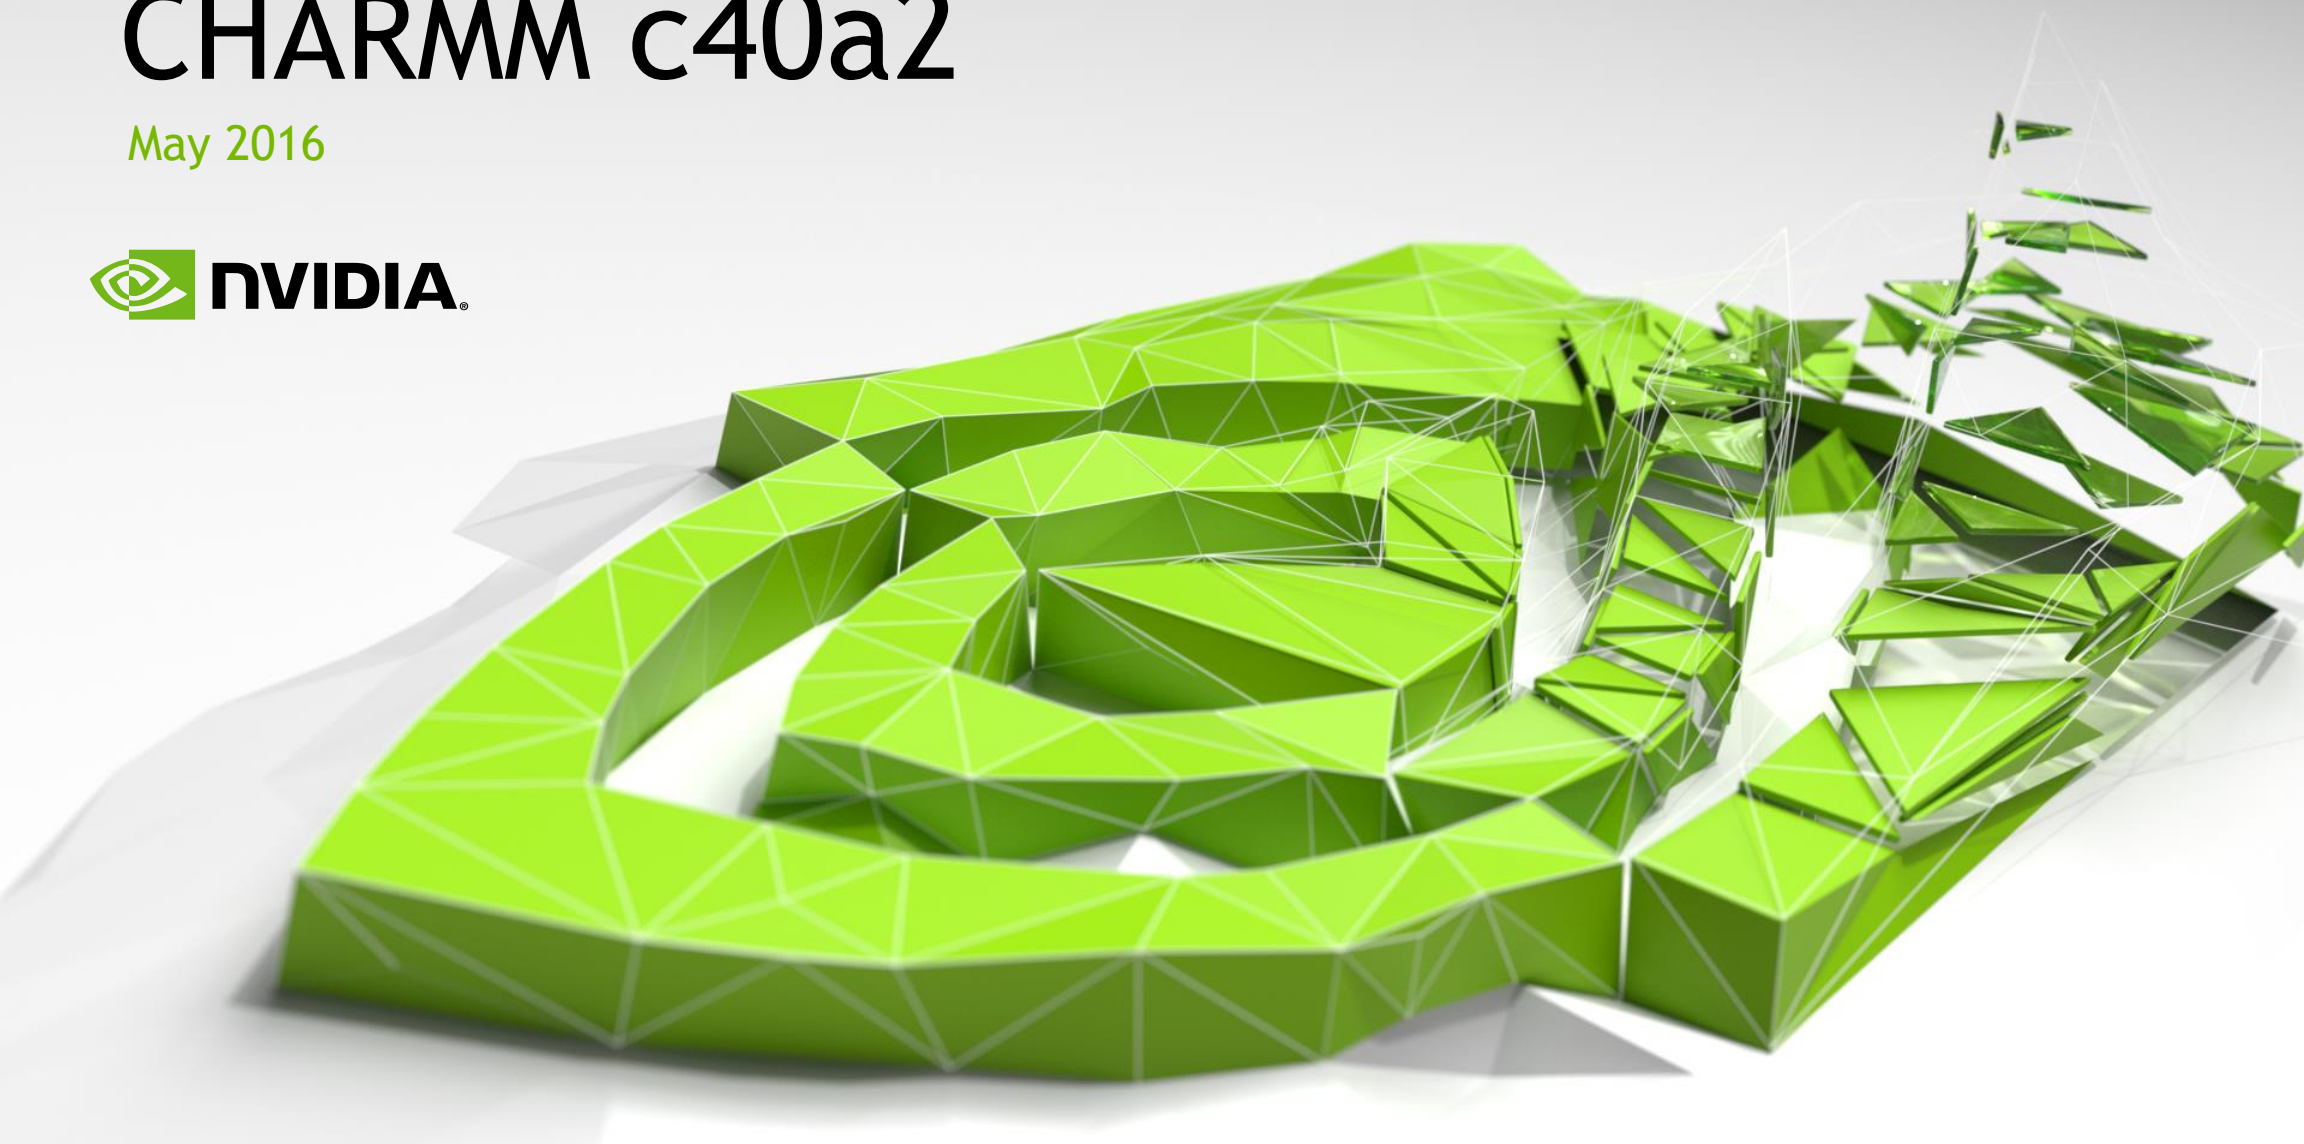

465K System on K80s

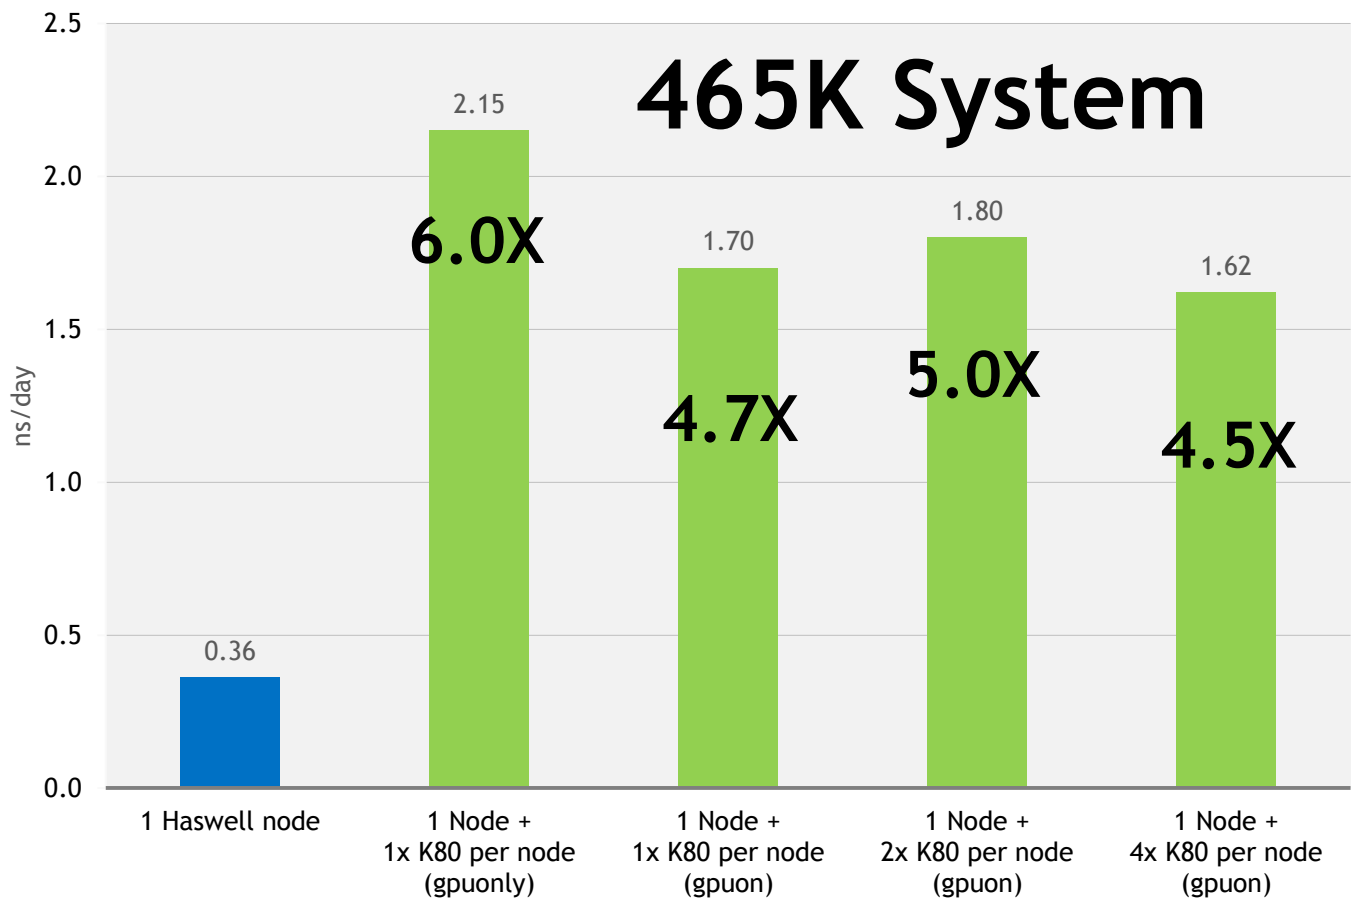

Running CHARMM version c40a2

The blue node contains Dual Intel Xeon Intel Xeon (R) ES-2698@2.30 MHZ (Haswell) CPUs

The green nodes contain Dual Intel Xeon Intel Xeon (R) ES-2698@2.30 MHZ (Haswell) CPUs + Tesla K80 (autoboost) GPUs

"Gpuonly" means all the forces are calculated in GPU
"Gpuon" means only non-bonded forces are calculated in GPU

534K System on K80s

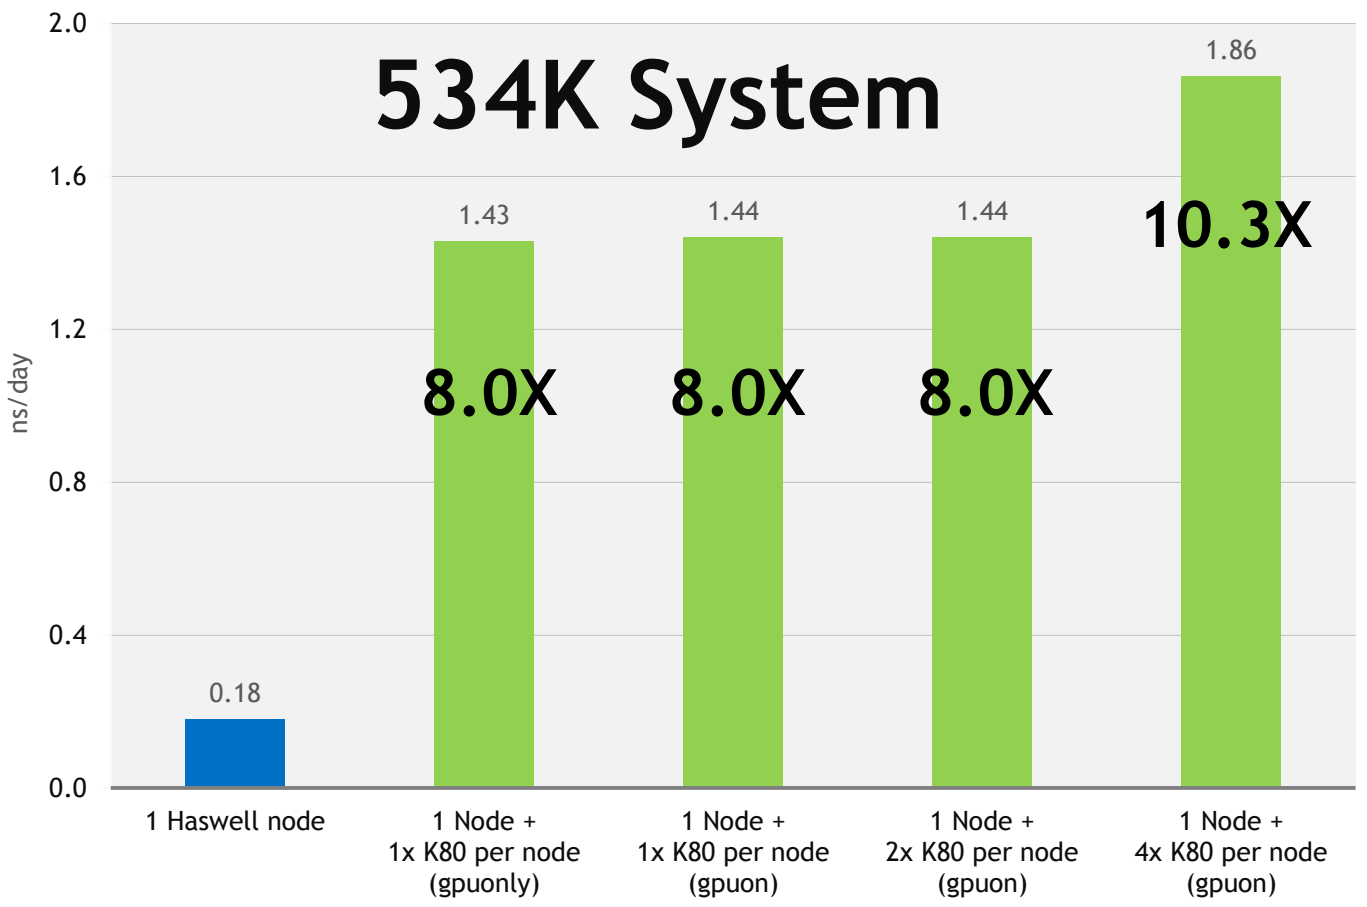

Running CHARMM version c40a2

The blue node contains Dual Intel Xeon Intel Xeon (R) ES-2698@2.30 MHZ (Haswell) CPUs

The green nodes contain Dual Intel Xeon Intel Xeon (R) ES-2698@2.30 MHZ (Haswell) CPUs + Tesla K80 (autoboost) GPUs

"Gpuonly" means all the forces are calculated in GPU
"Gpuon" means only non-bonded forces are calculated in GPU

GROMACS 5.1

October 2015

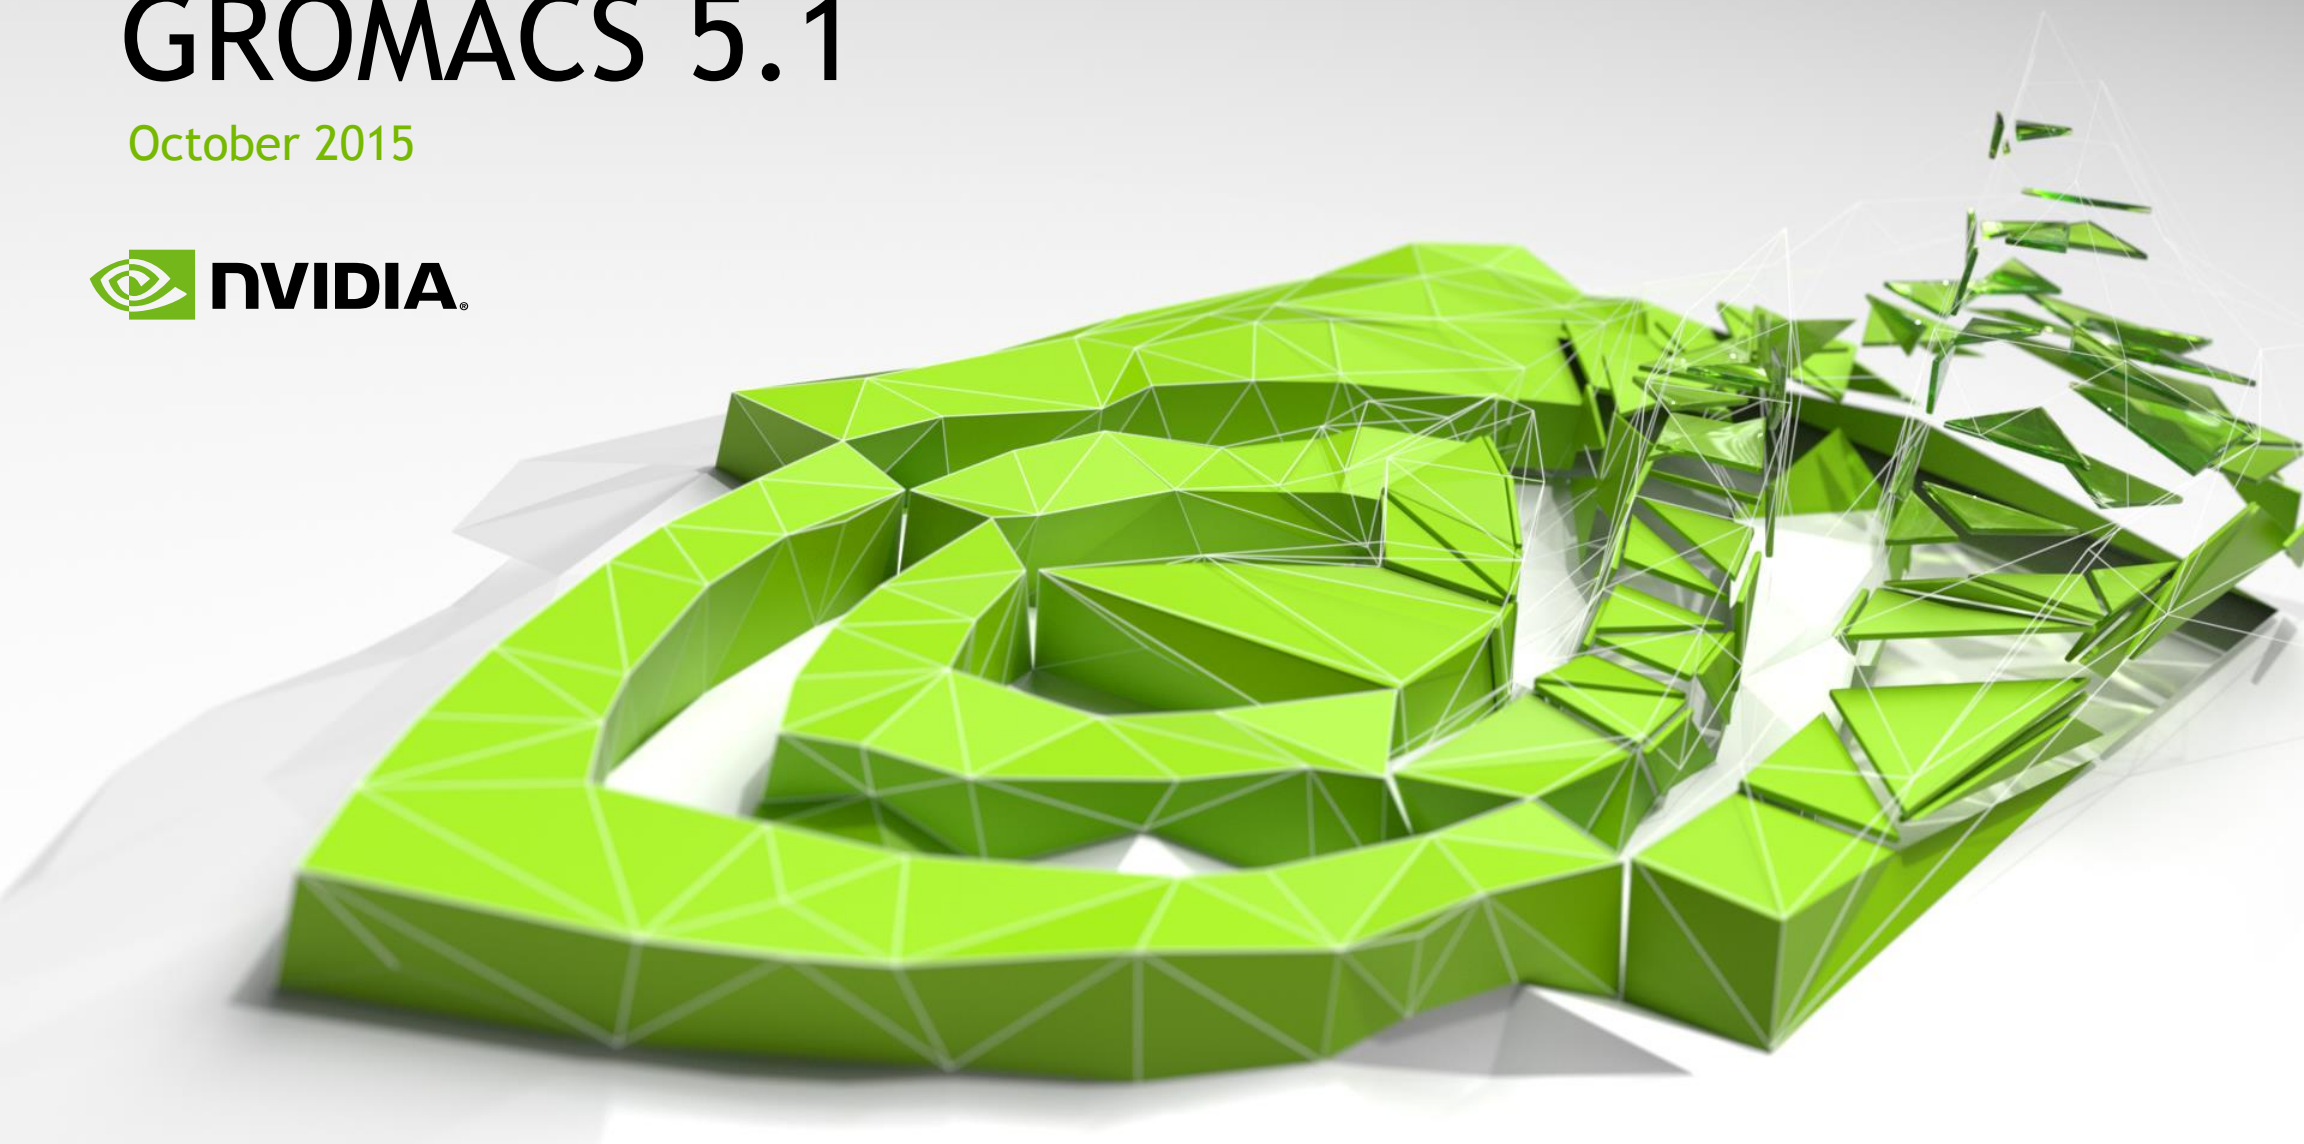

Erik Lindahl (GROMACS developer) video

384K Waters on K40s and K80s

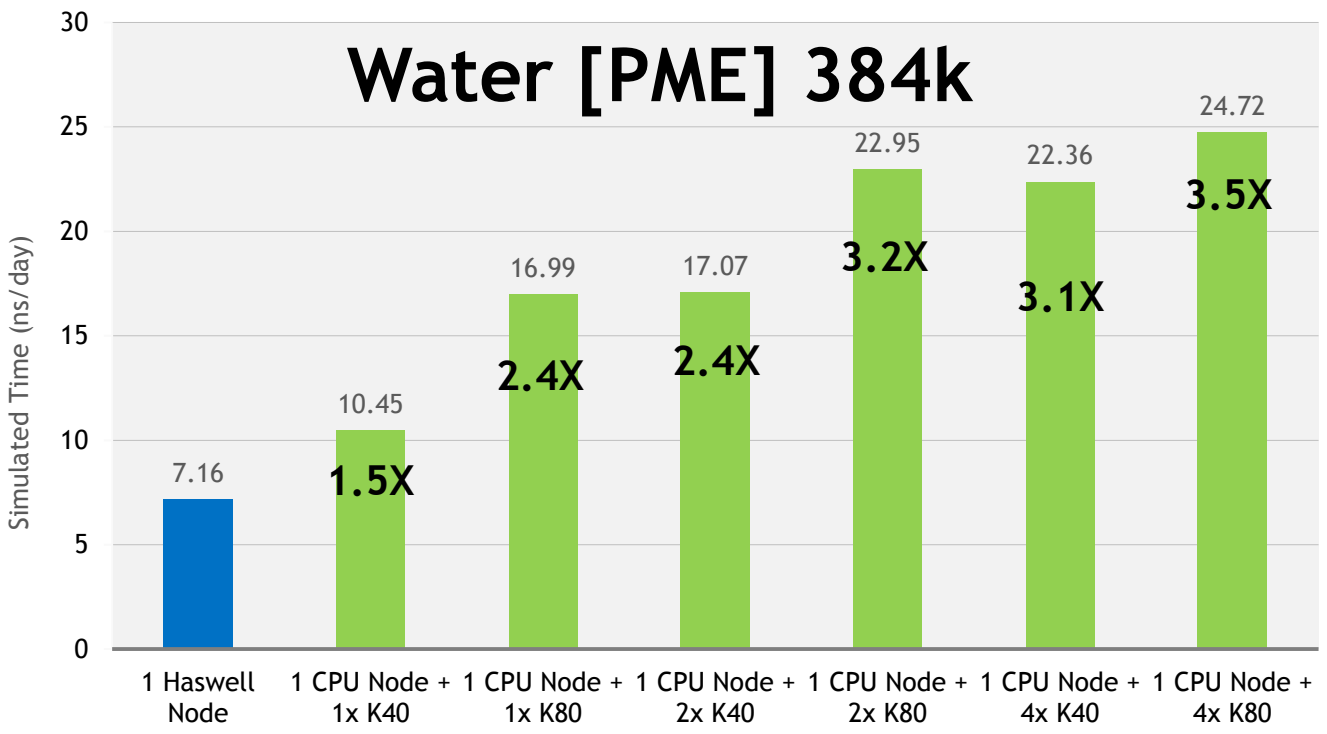

Running GROMACS version 5.1

The blue node contains Dual Intel E5-2698 v3@2.3GHz CPUs

The green nodes contain Dual Intel E5-2698 v3@2.3GHz CPUs + either NVIDIA Tesla K40@875Mhz or Tesla K80@562Mhz (autoboost) GPUs

384K Waters on Titan X

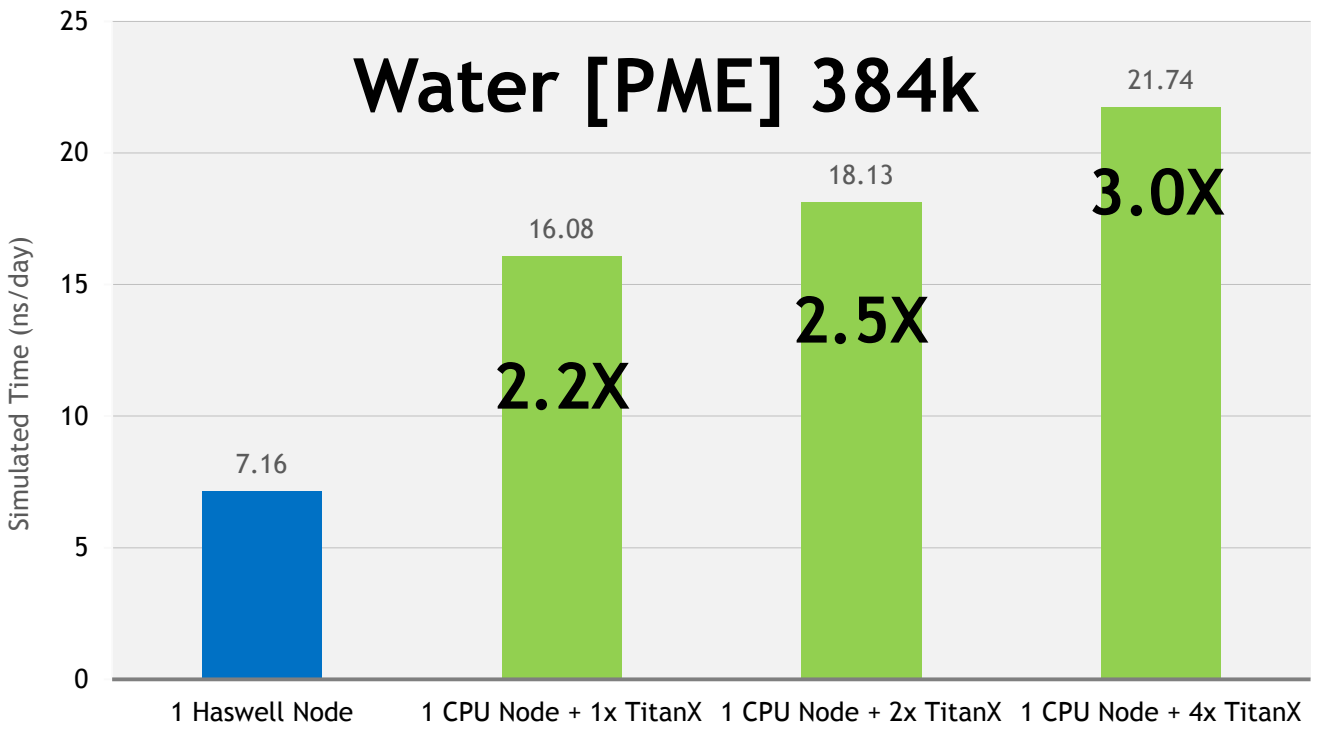

Running GROMACS version 5.1

The blue node contains Dual Intel E5-2698 v3@2.3GHz CPUs

The green nodes contain Dual Intel E5-2698 v3@2.3GHz CPUs + GeForce GTX TitanX@1000Mhz GPUs

768K Waters on K40s and K80s

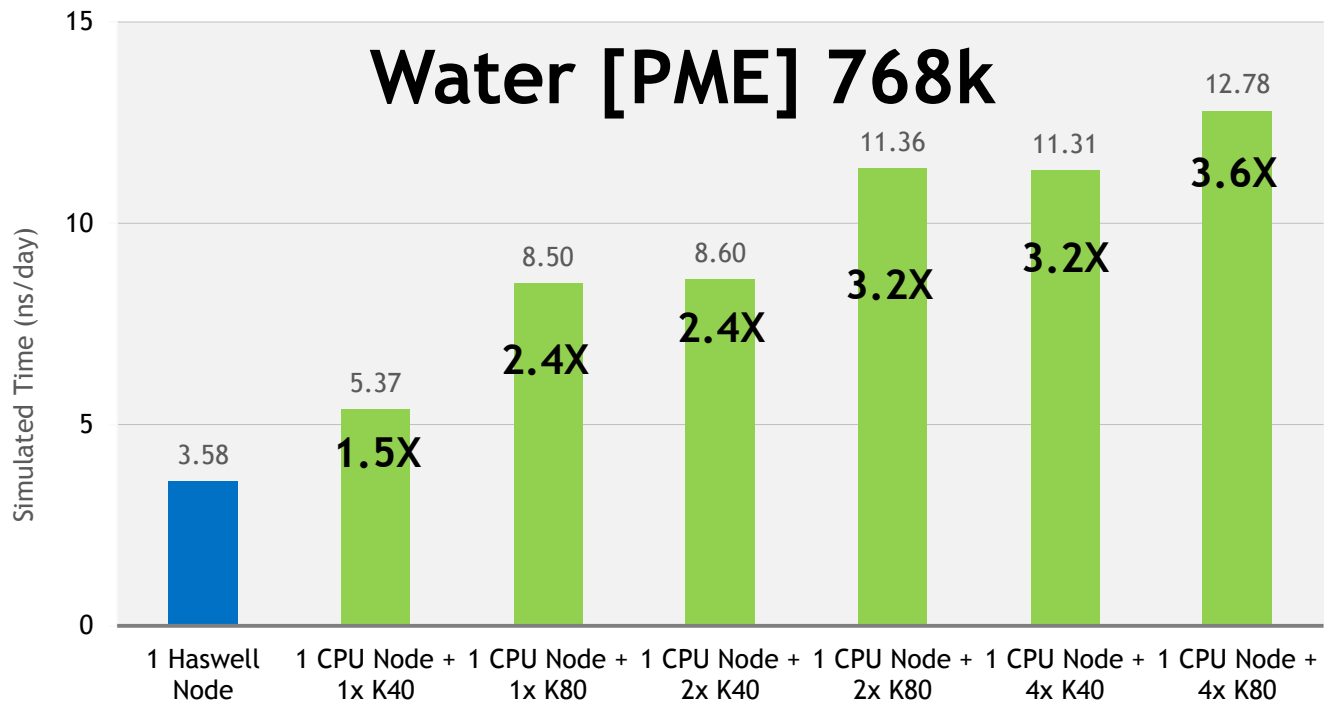

Running GROMACS version 5.1

The blue node contains Dual Intel E5-2698 v3@2.3GHz CPUs

The green nodes contain Dual Intel E5-2698 v3@2.3GHz CPUs + either NVIDIA Tesla K40@875Mhz or Tesla K80@562Mhz (autoboost) GPUs

768K Waters on Titan X

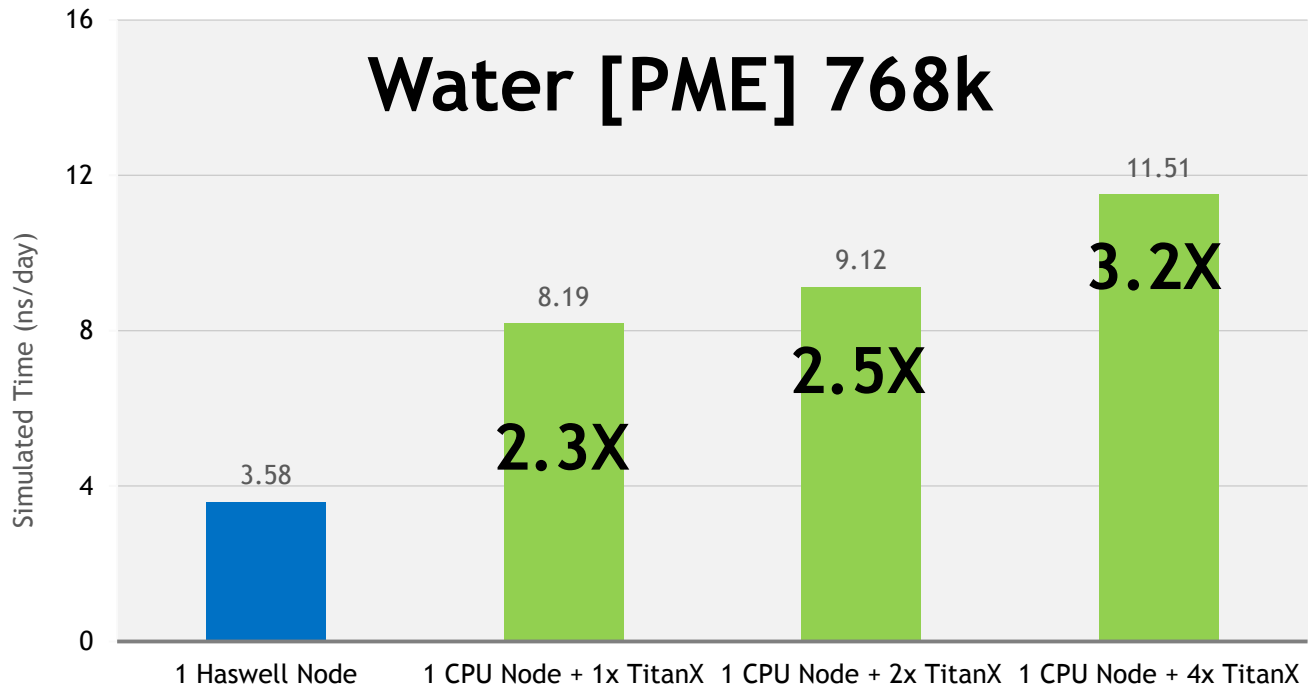

Running GROMACS version 5.1

The blue node contains Dual Intel E5-2698 v3@2.3GHz CPUs

The green nodes contain Dual Intel E5-2698 v3@2.3GHz CPUs + GeForce GTX TitanX@1000Mhz GPUs

1.5M Waters on K40s and K80s

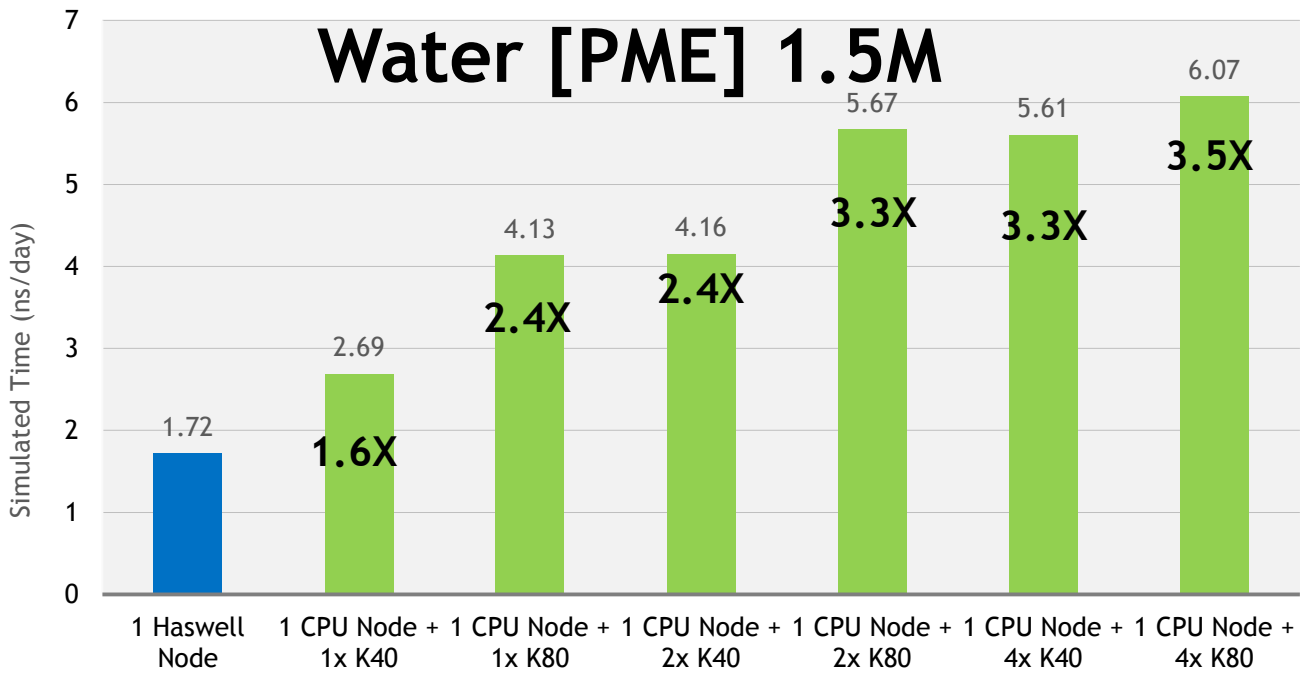

Running GROMACS version 5.1

The blue node contains Dual Intel E5-2698 v3@2.3GHz CPUs

The green nodes contain Dual Intel E5-2698 v3@2.3GHz CPUs + either NVIDIA Tesla K40@875Mhz or Tesla K80@562Mhz (autoboost) GPUs

1.5M Waters on Titan X

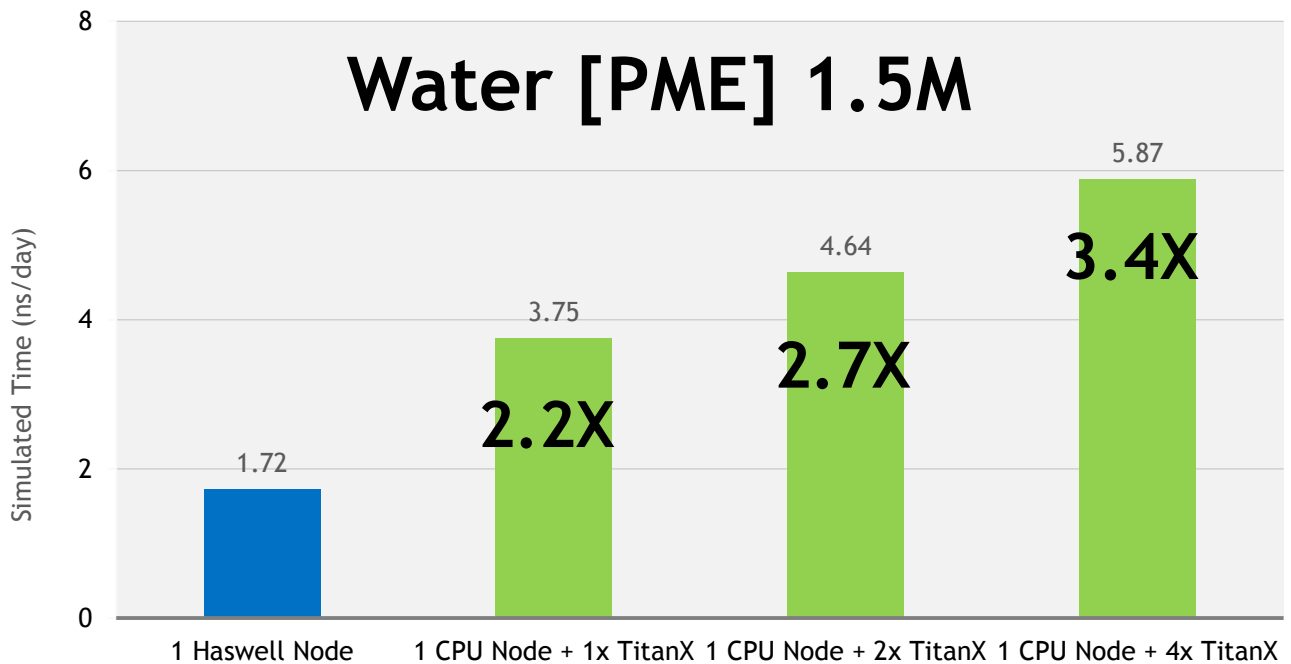

Running GROMACS version 5.1

The blue node contains Dual Intel E5-2698 v3@2.3GHz CPUs

The green nodes contain Dual Intel E5-2698 v3@2.3GHz CPUs + GeForce GTX TitanX@1000Mhz GPUs

3M Waters on K40s and K80s

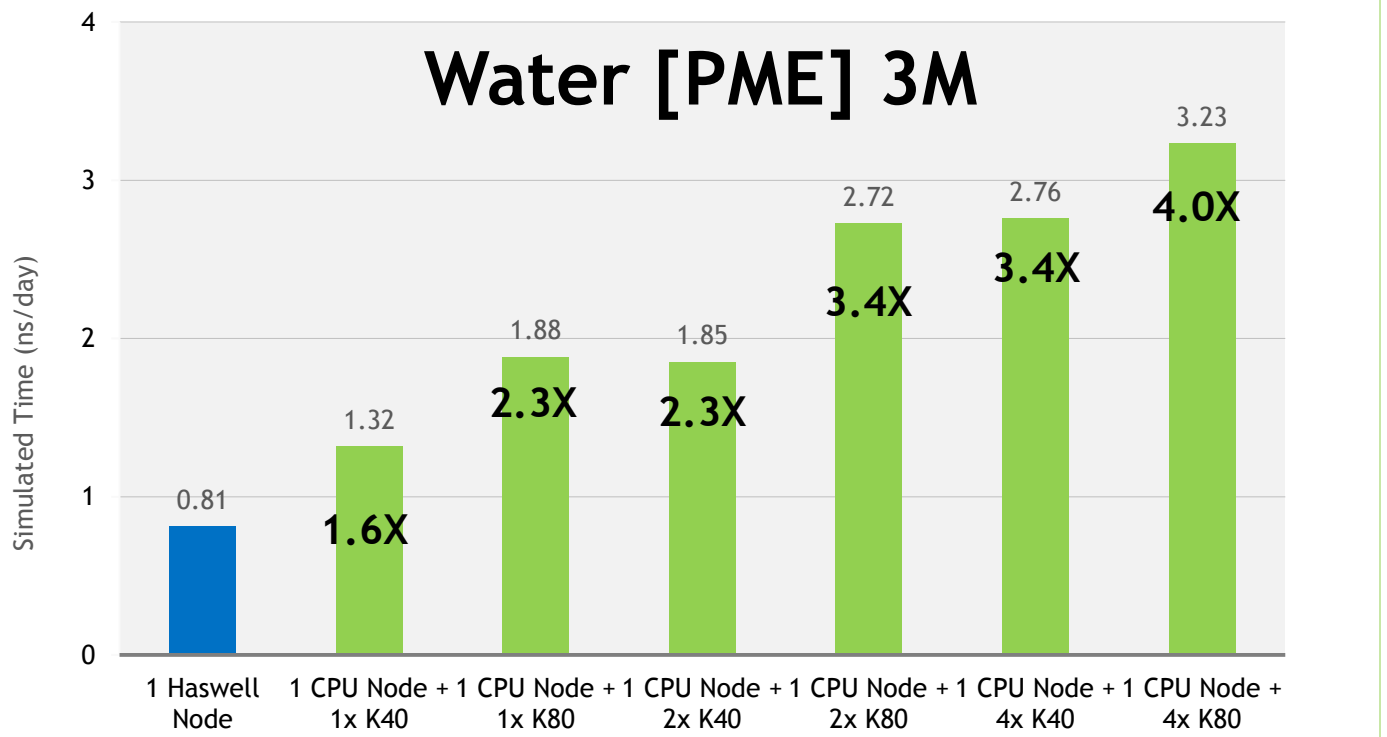

Running GROMACS version 5.1

The blue node contains Dual Intel E5-2698 v3@2.3GHz CPUs

The green nodes contain Dual Intel E5-2698 v3@2.3GHz CPUs + either NVIDIA Tesla K40@875Mhz or Tesla K80@562Mhz (autoboost) GPUs

3M Waters on Titan X

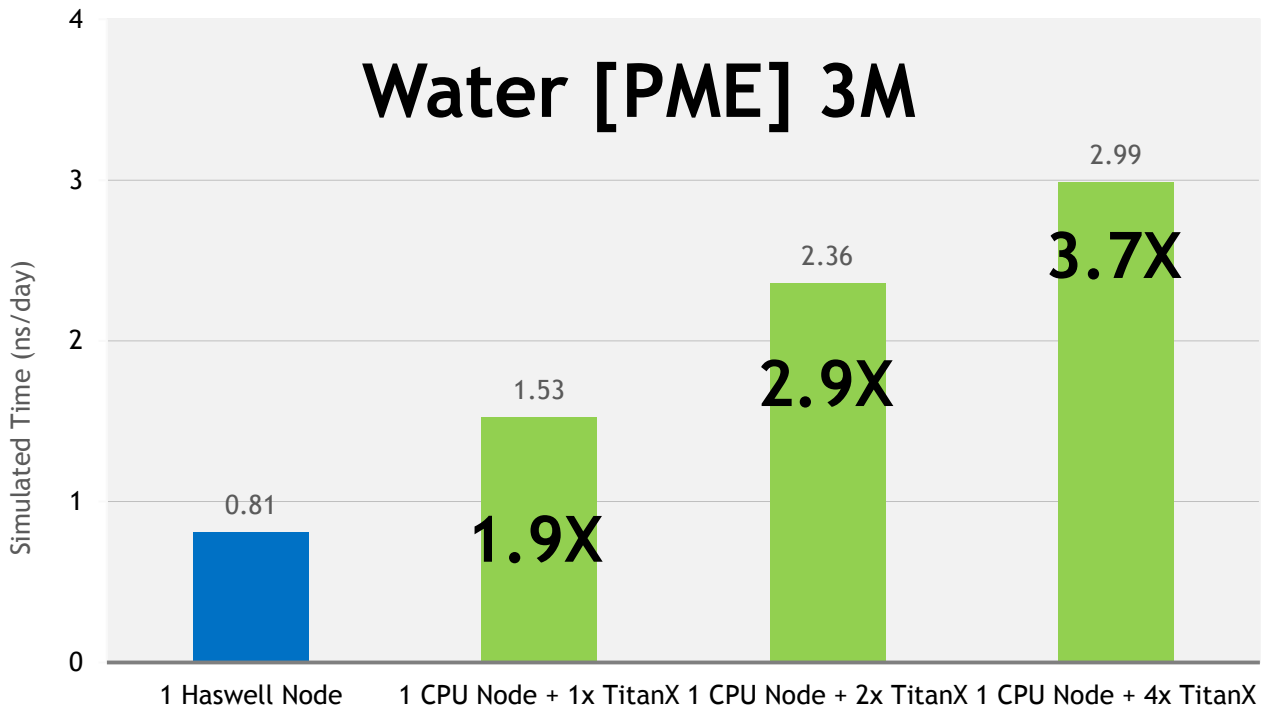

Running GROMACS version 5.1

The blue node contains Dual Intel E5-2698 v3@2.3GHz CPUs

The green nodes contain Dual Intel E5-2698 v3@2.3GHz CPUs + GeForce GTX TitanX@1000Mhz GPUs

GROMACS 5.0: Phi vs. Kepler K40 fastest GPU!

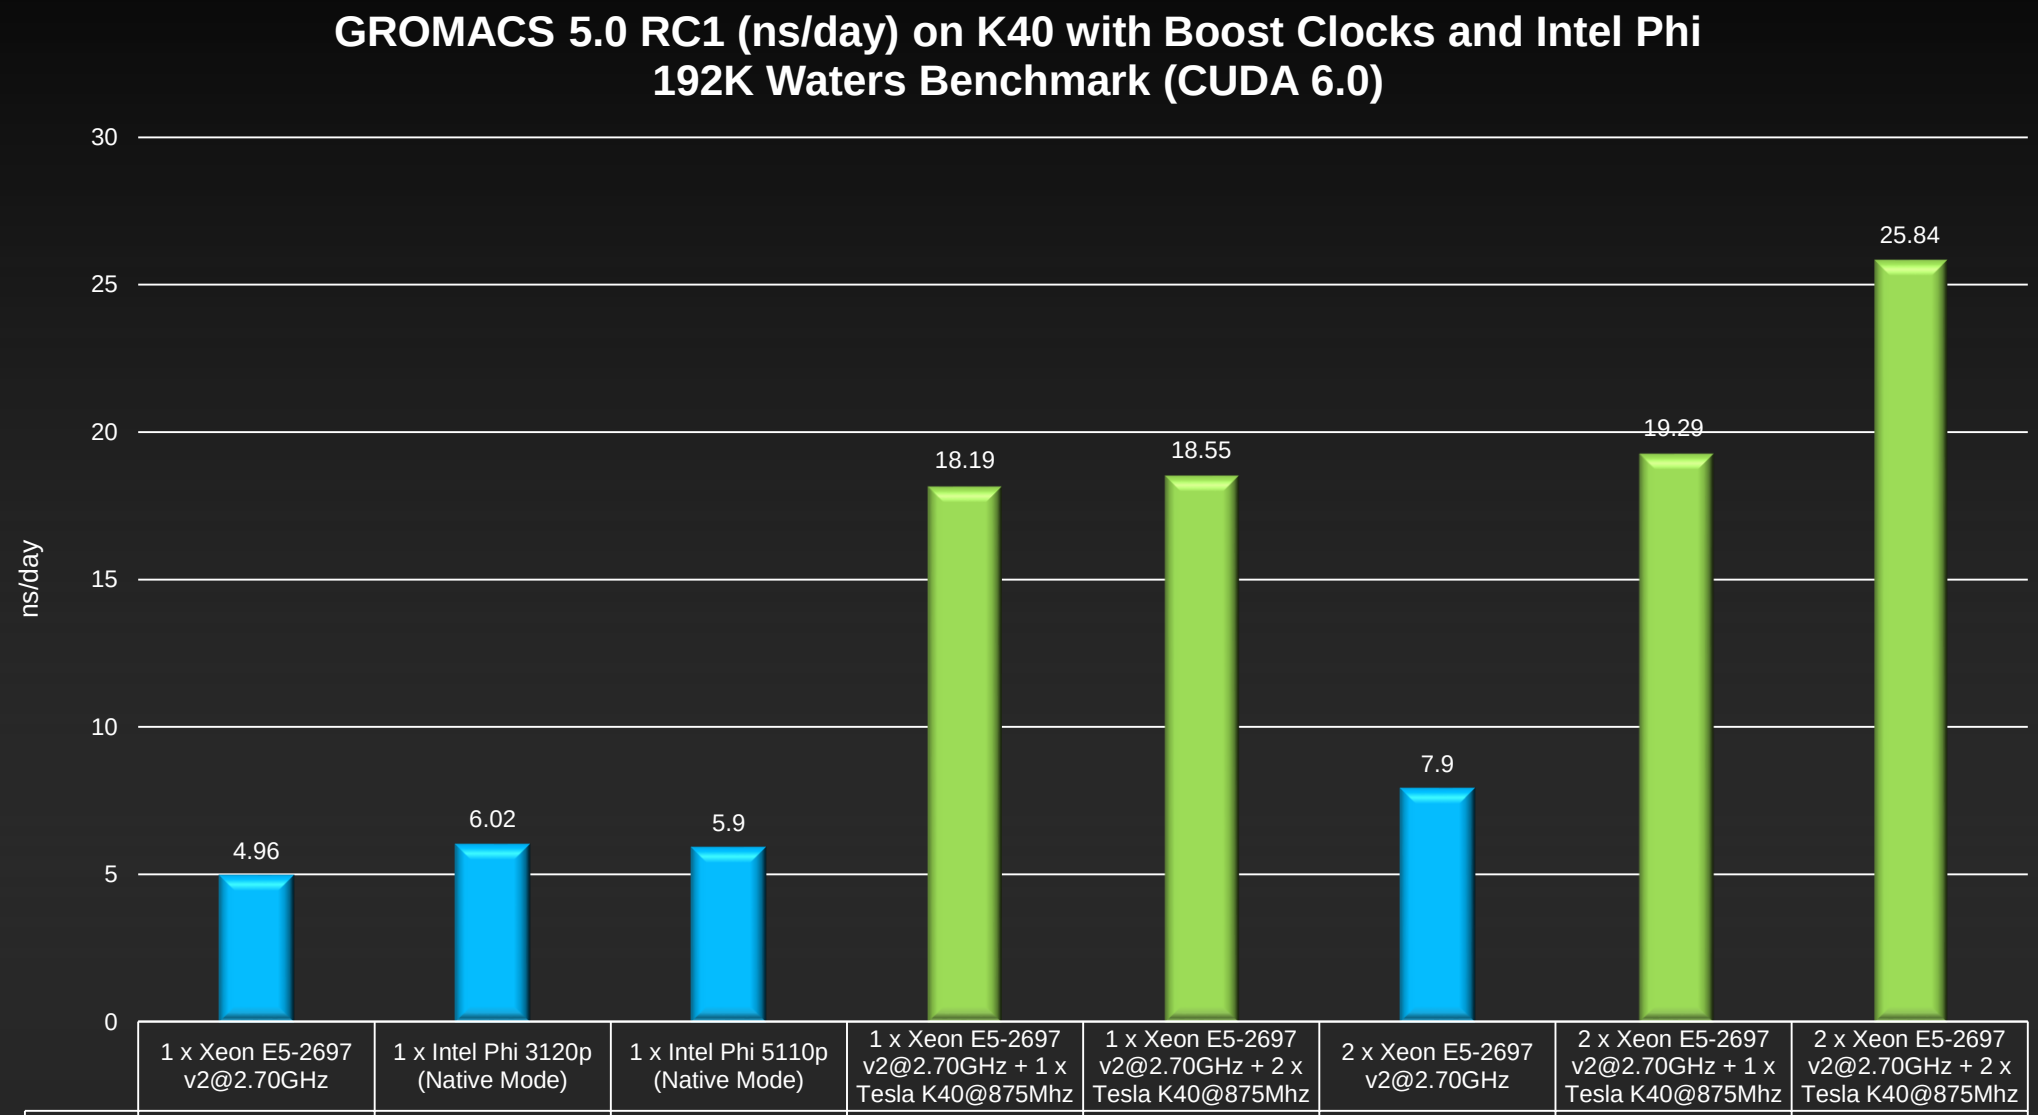

GROMACS 5.0 & Fastest Kepler GPUs yet!

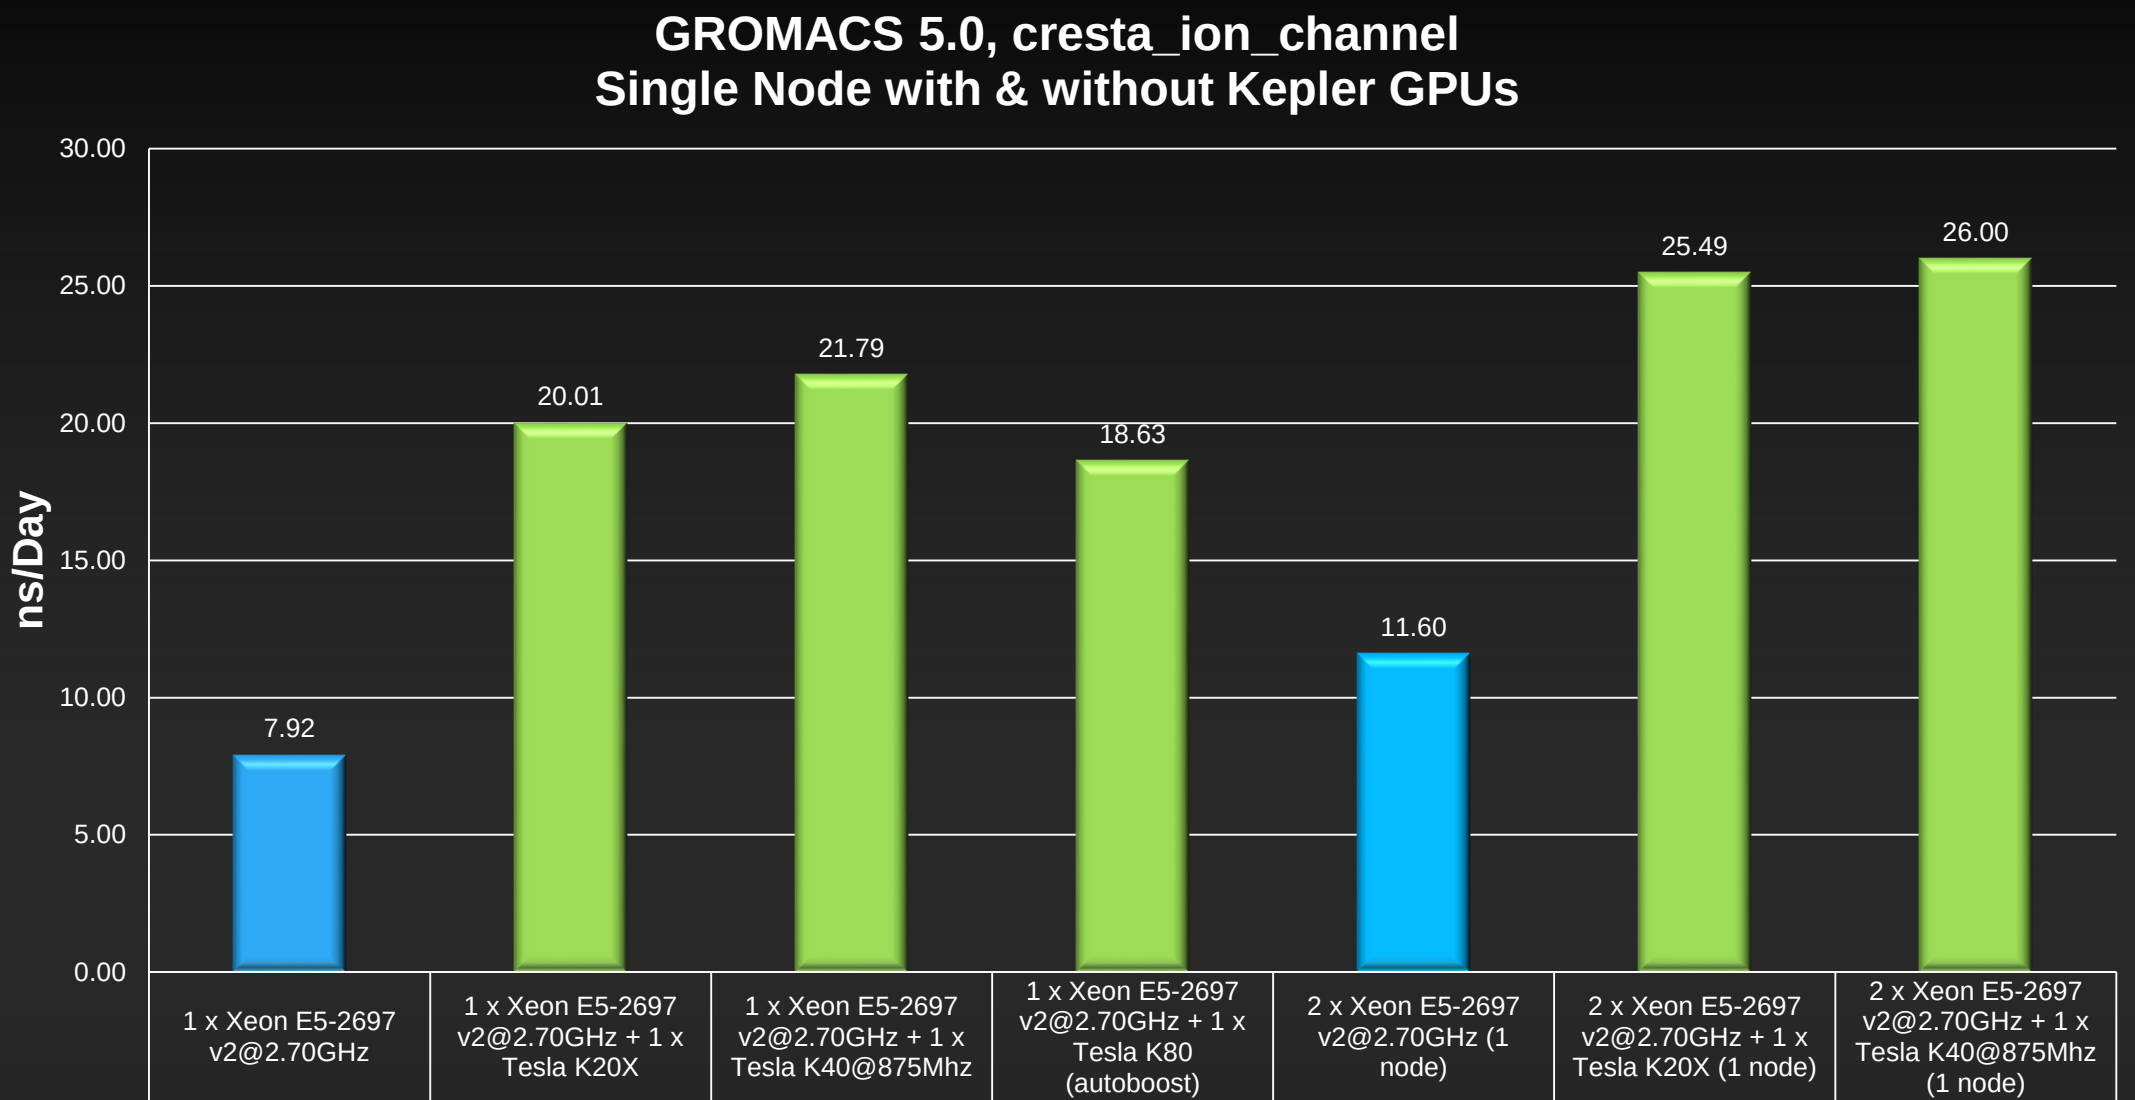

GROMACS 5.0 & Fastest Kepler GPUs yet!

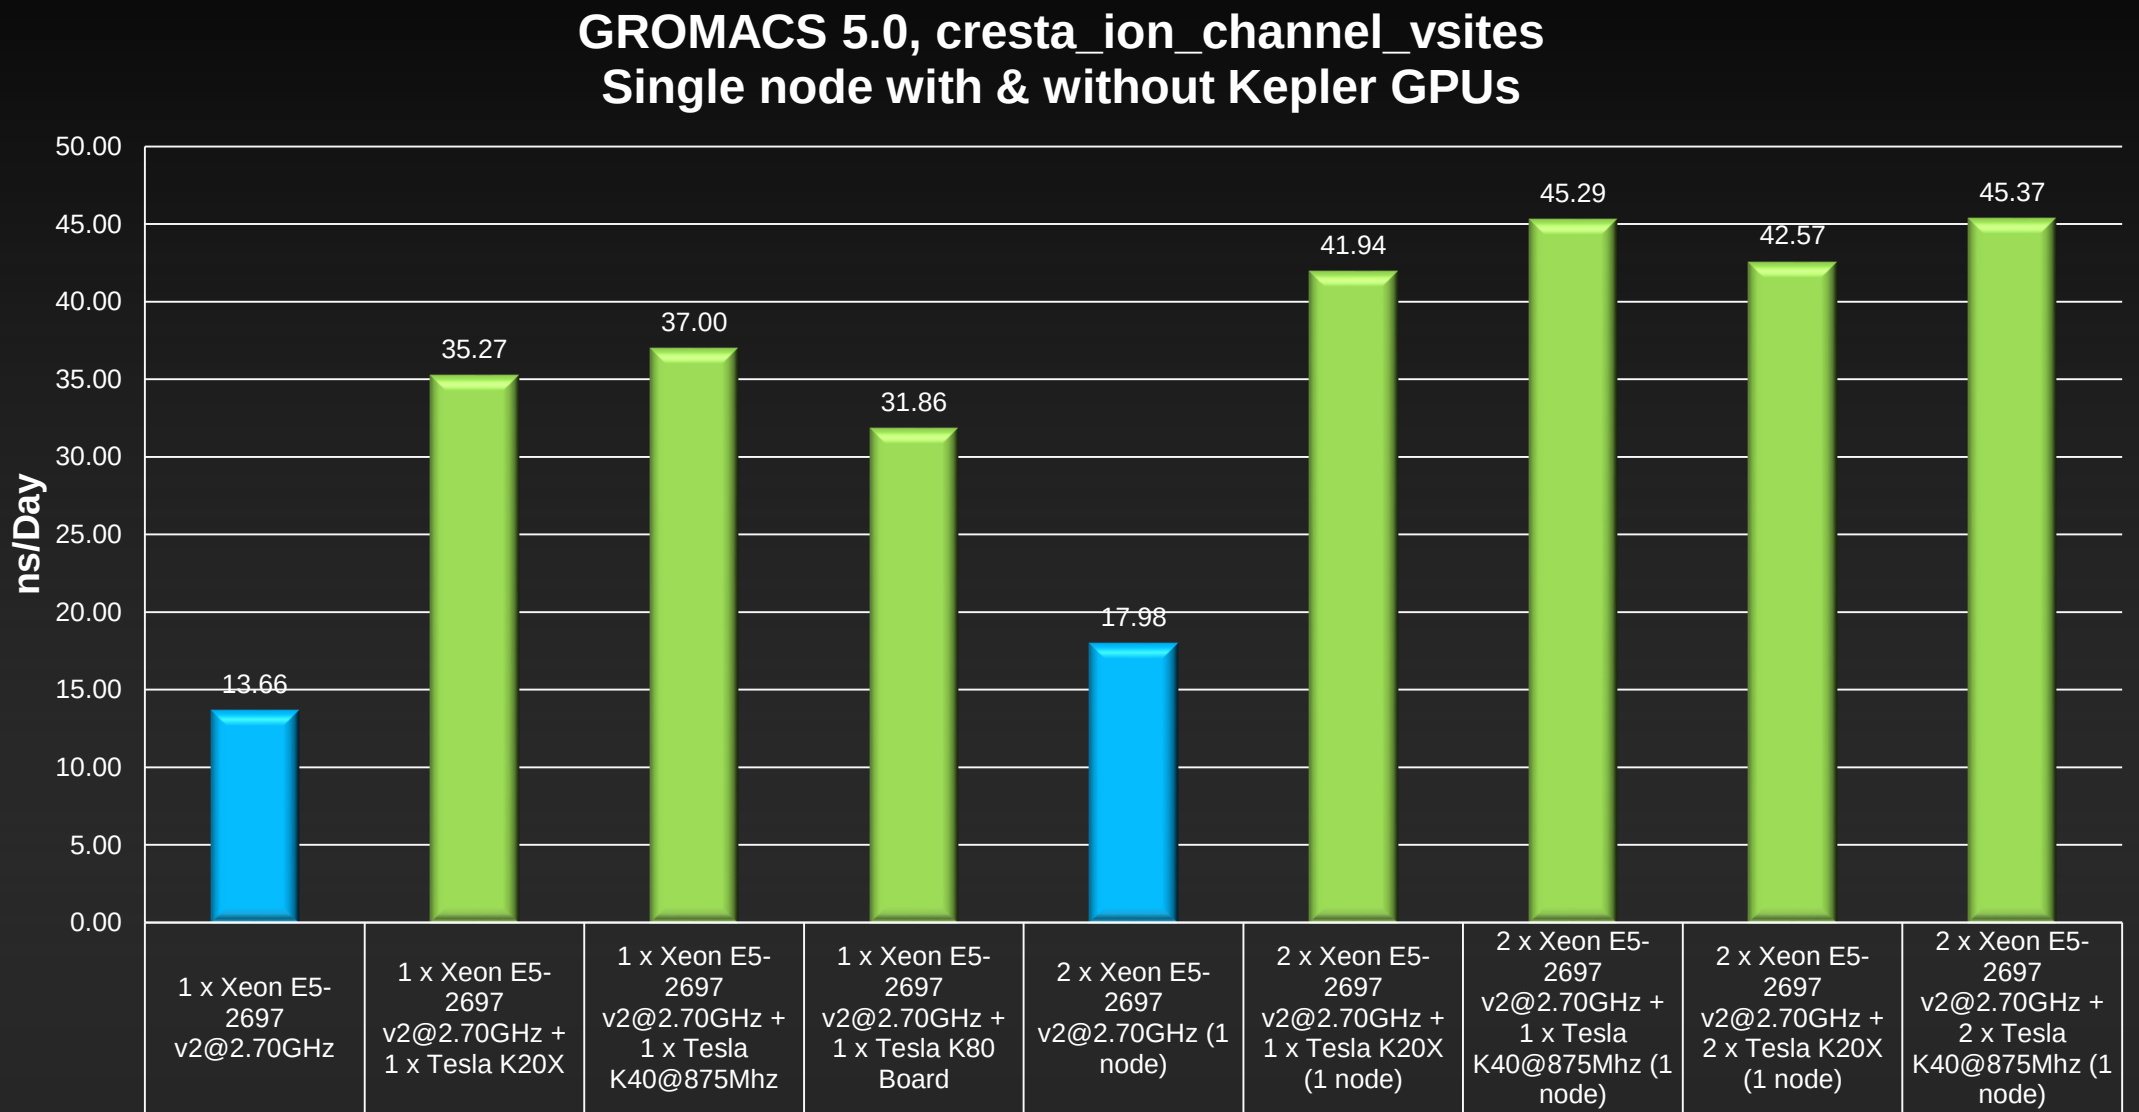

GROMACS 5.0 & Fastest Kepler GPUs yet!

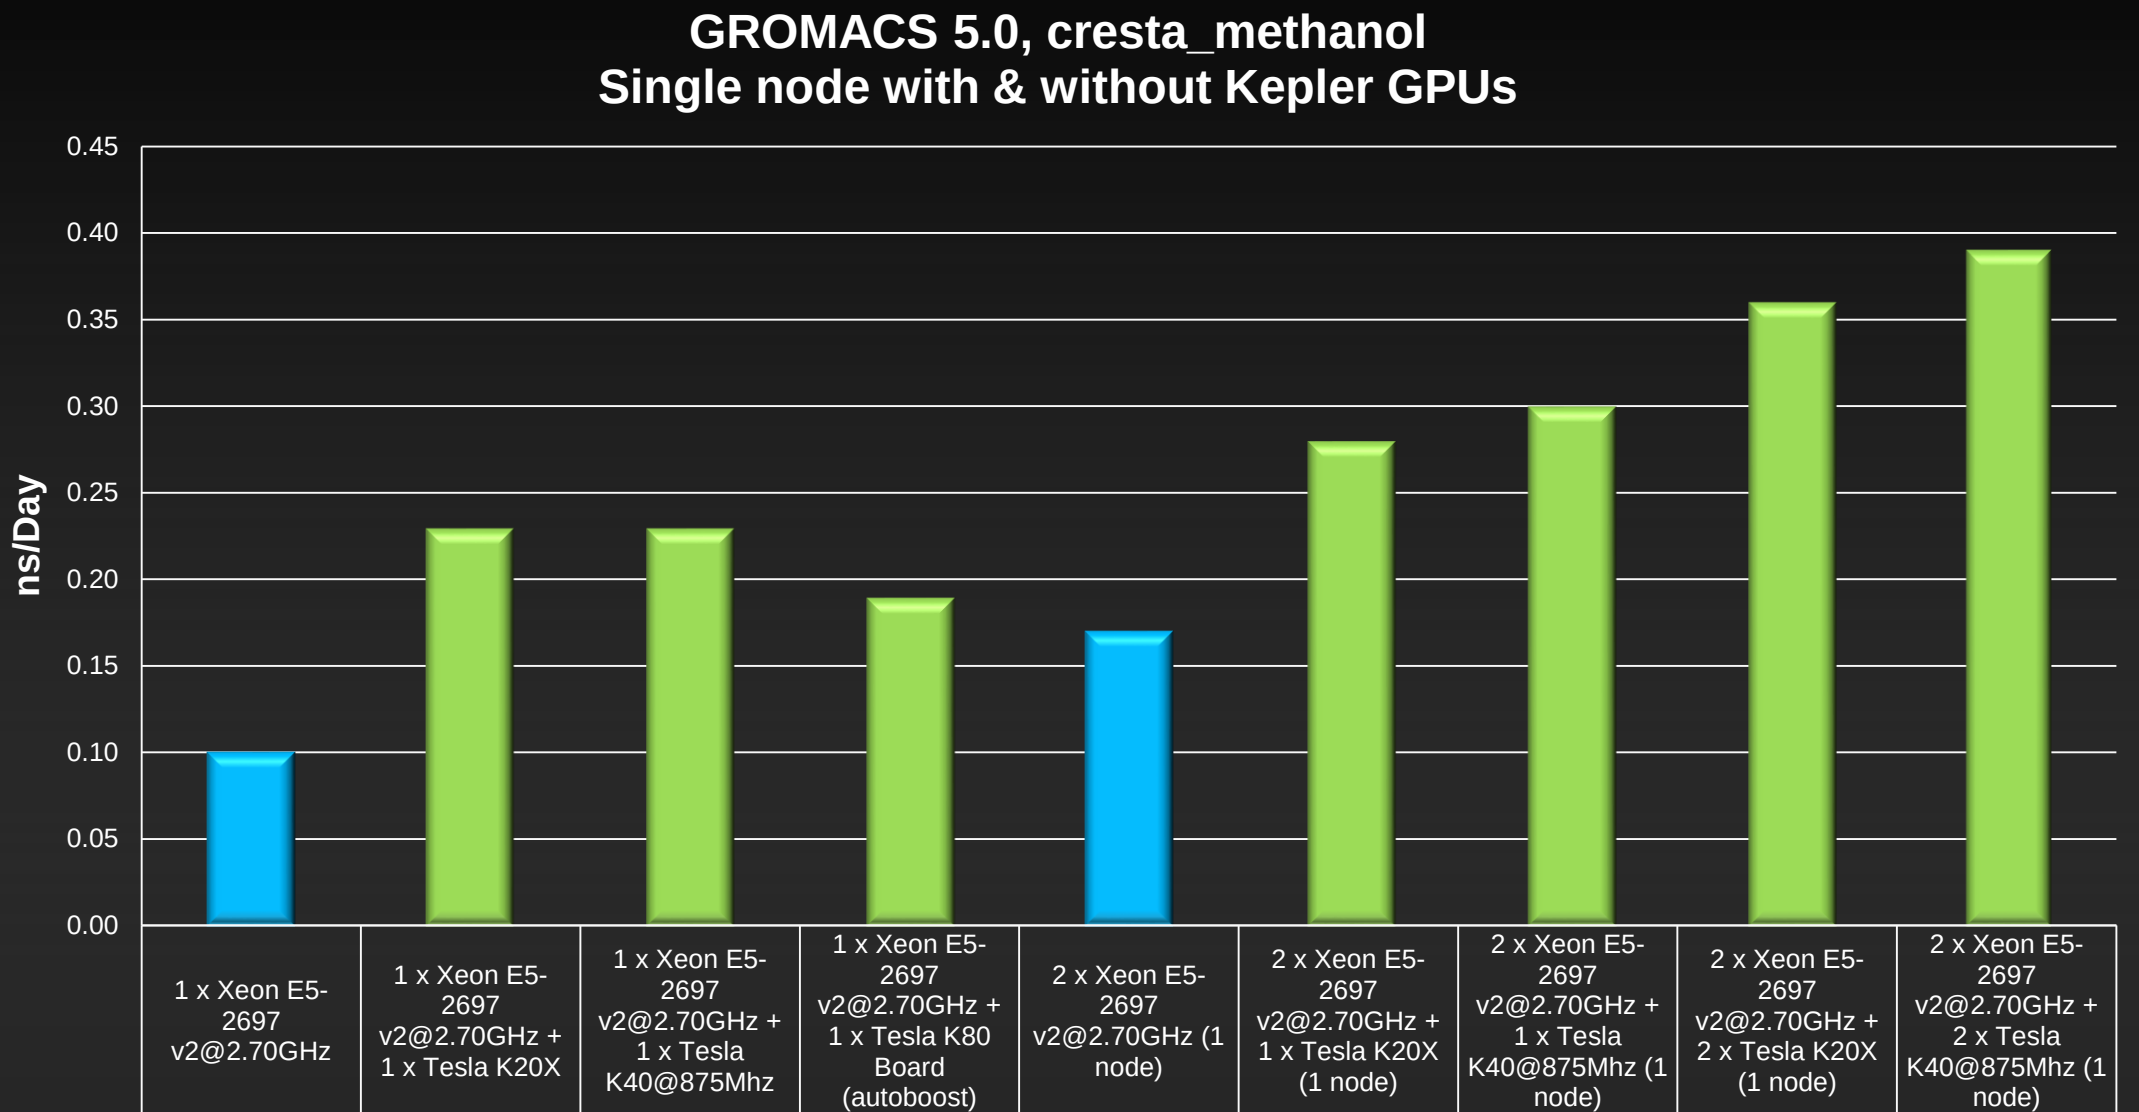

GROMACS 5.0 & Fastest Kepler GPUs yet!

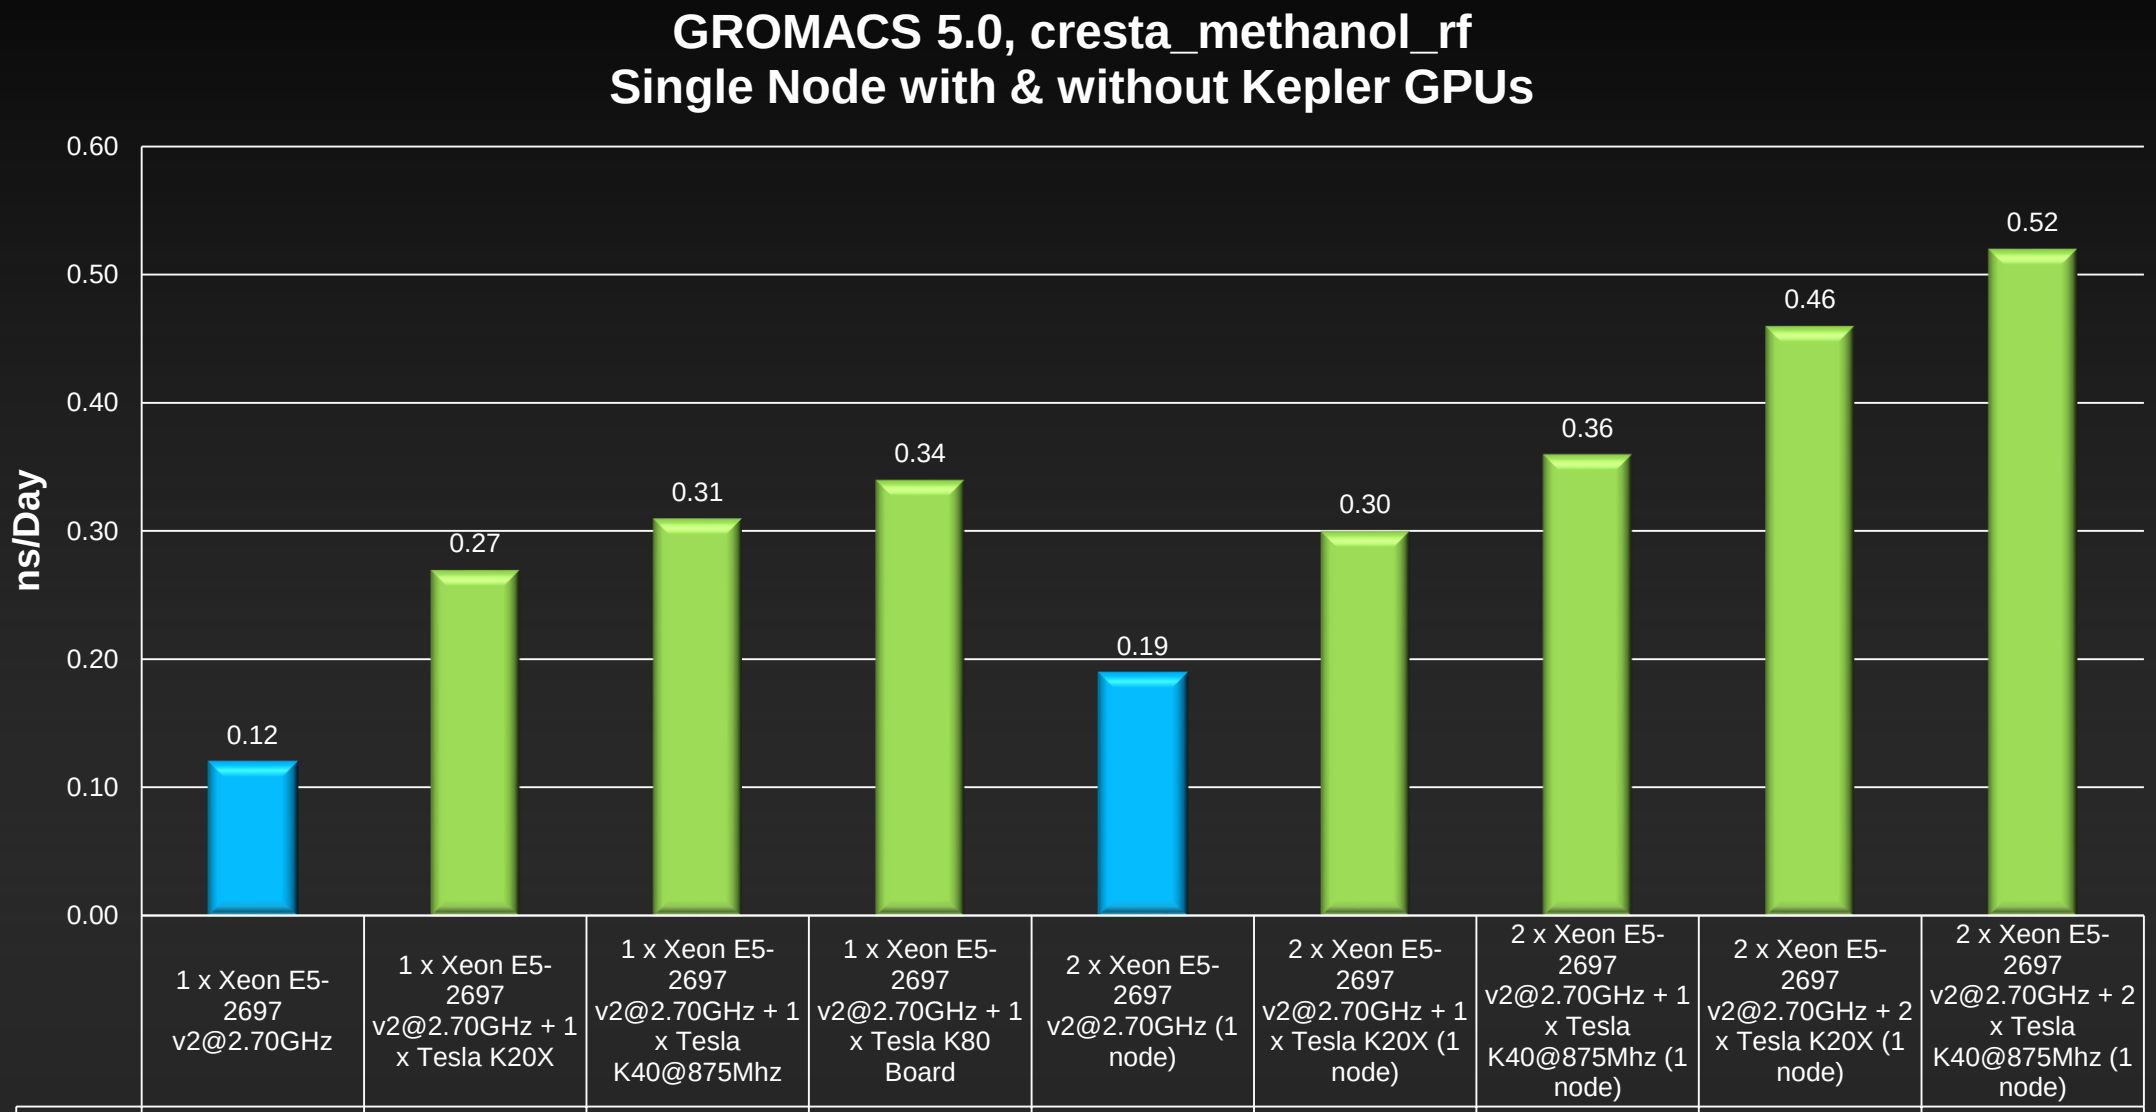

GROMACS 5.0 & Fastest Kepler GPUs yet!

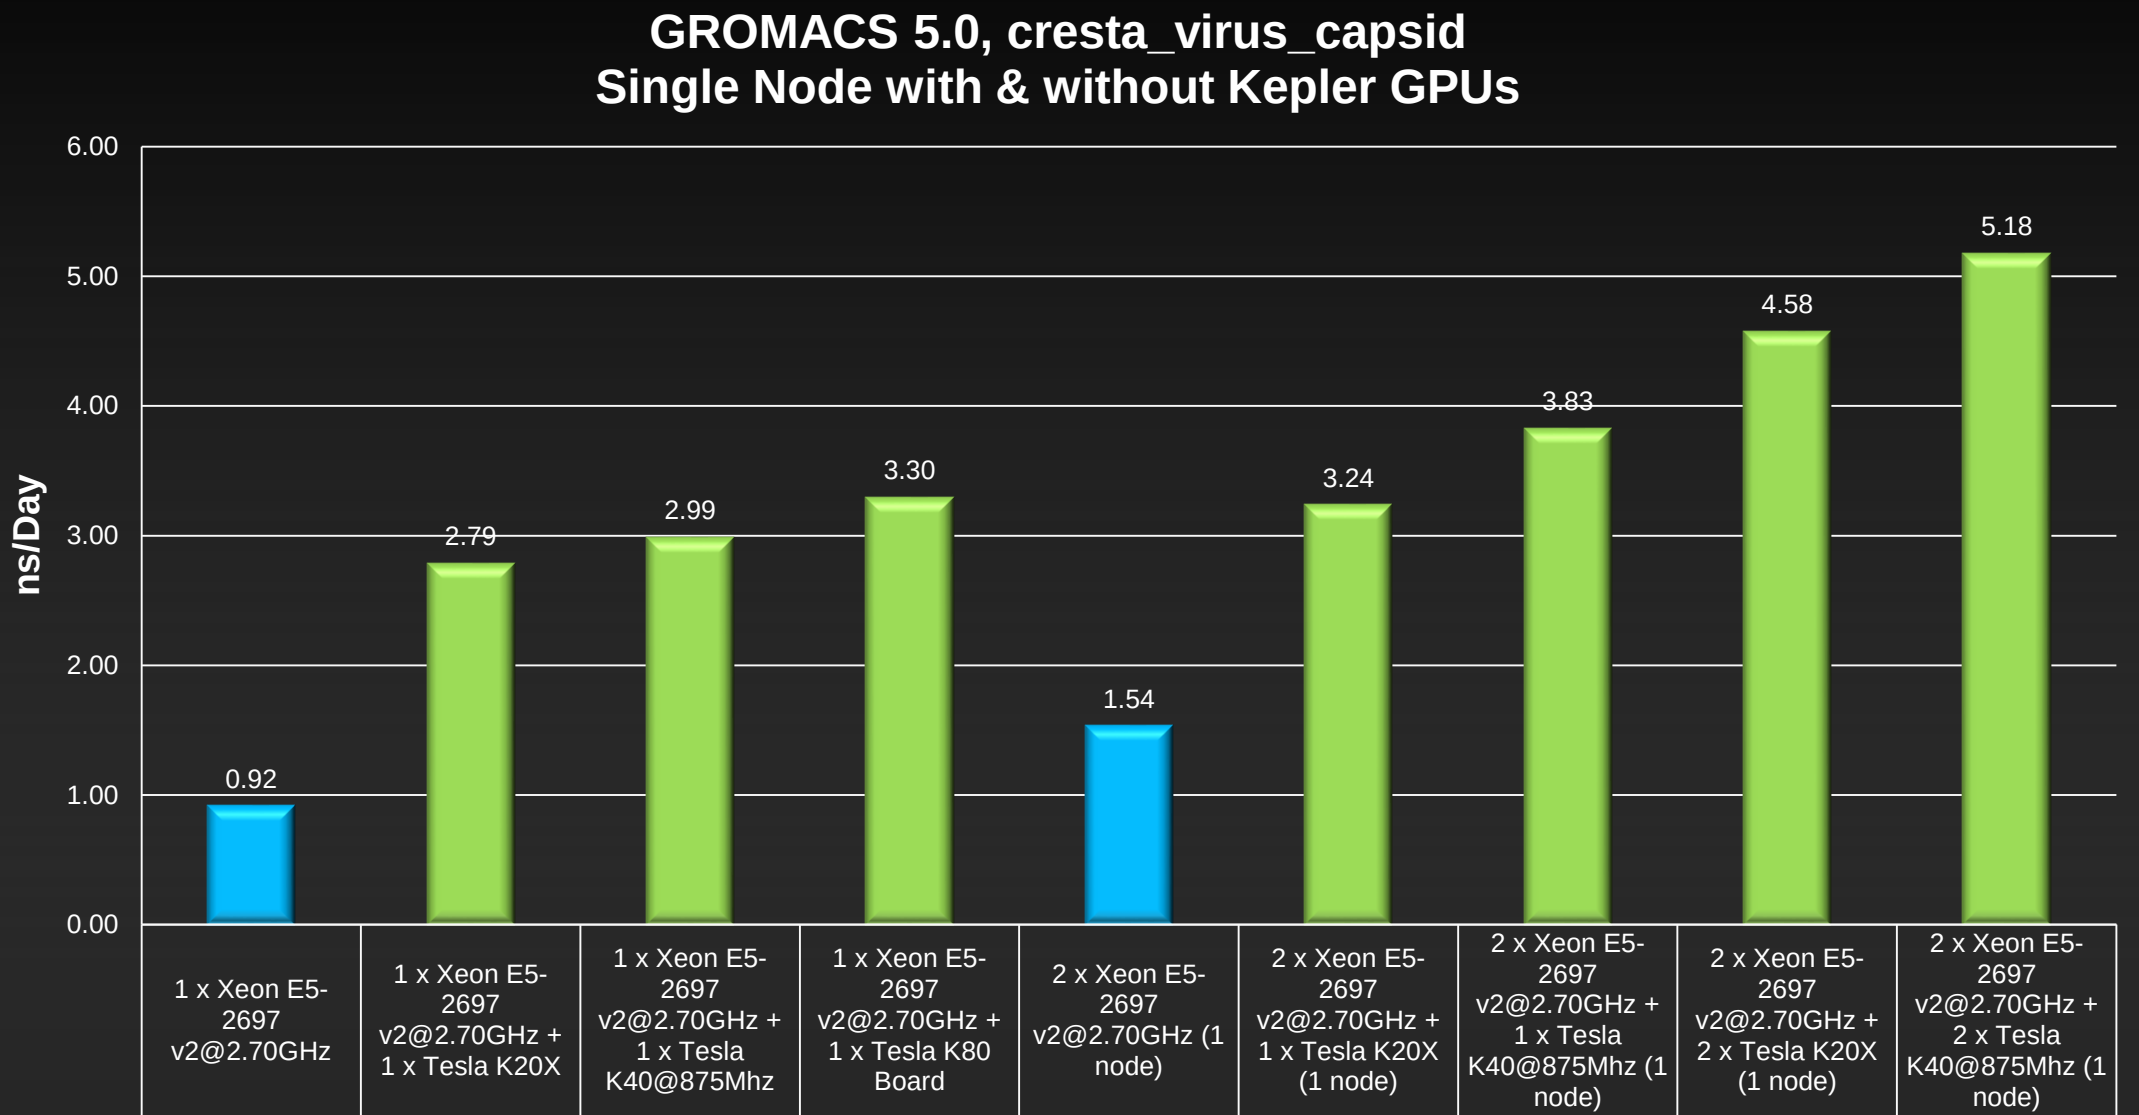

GROMACS 5.0 & Fastest Kepler GPUs yet!

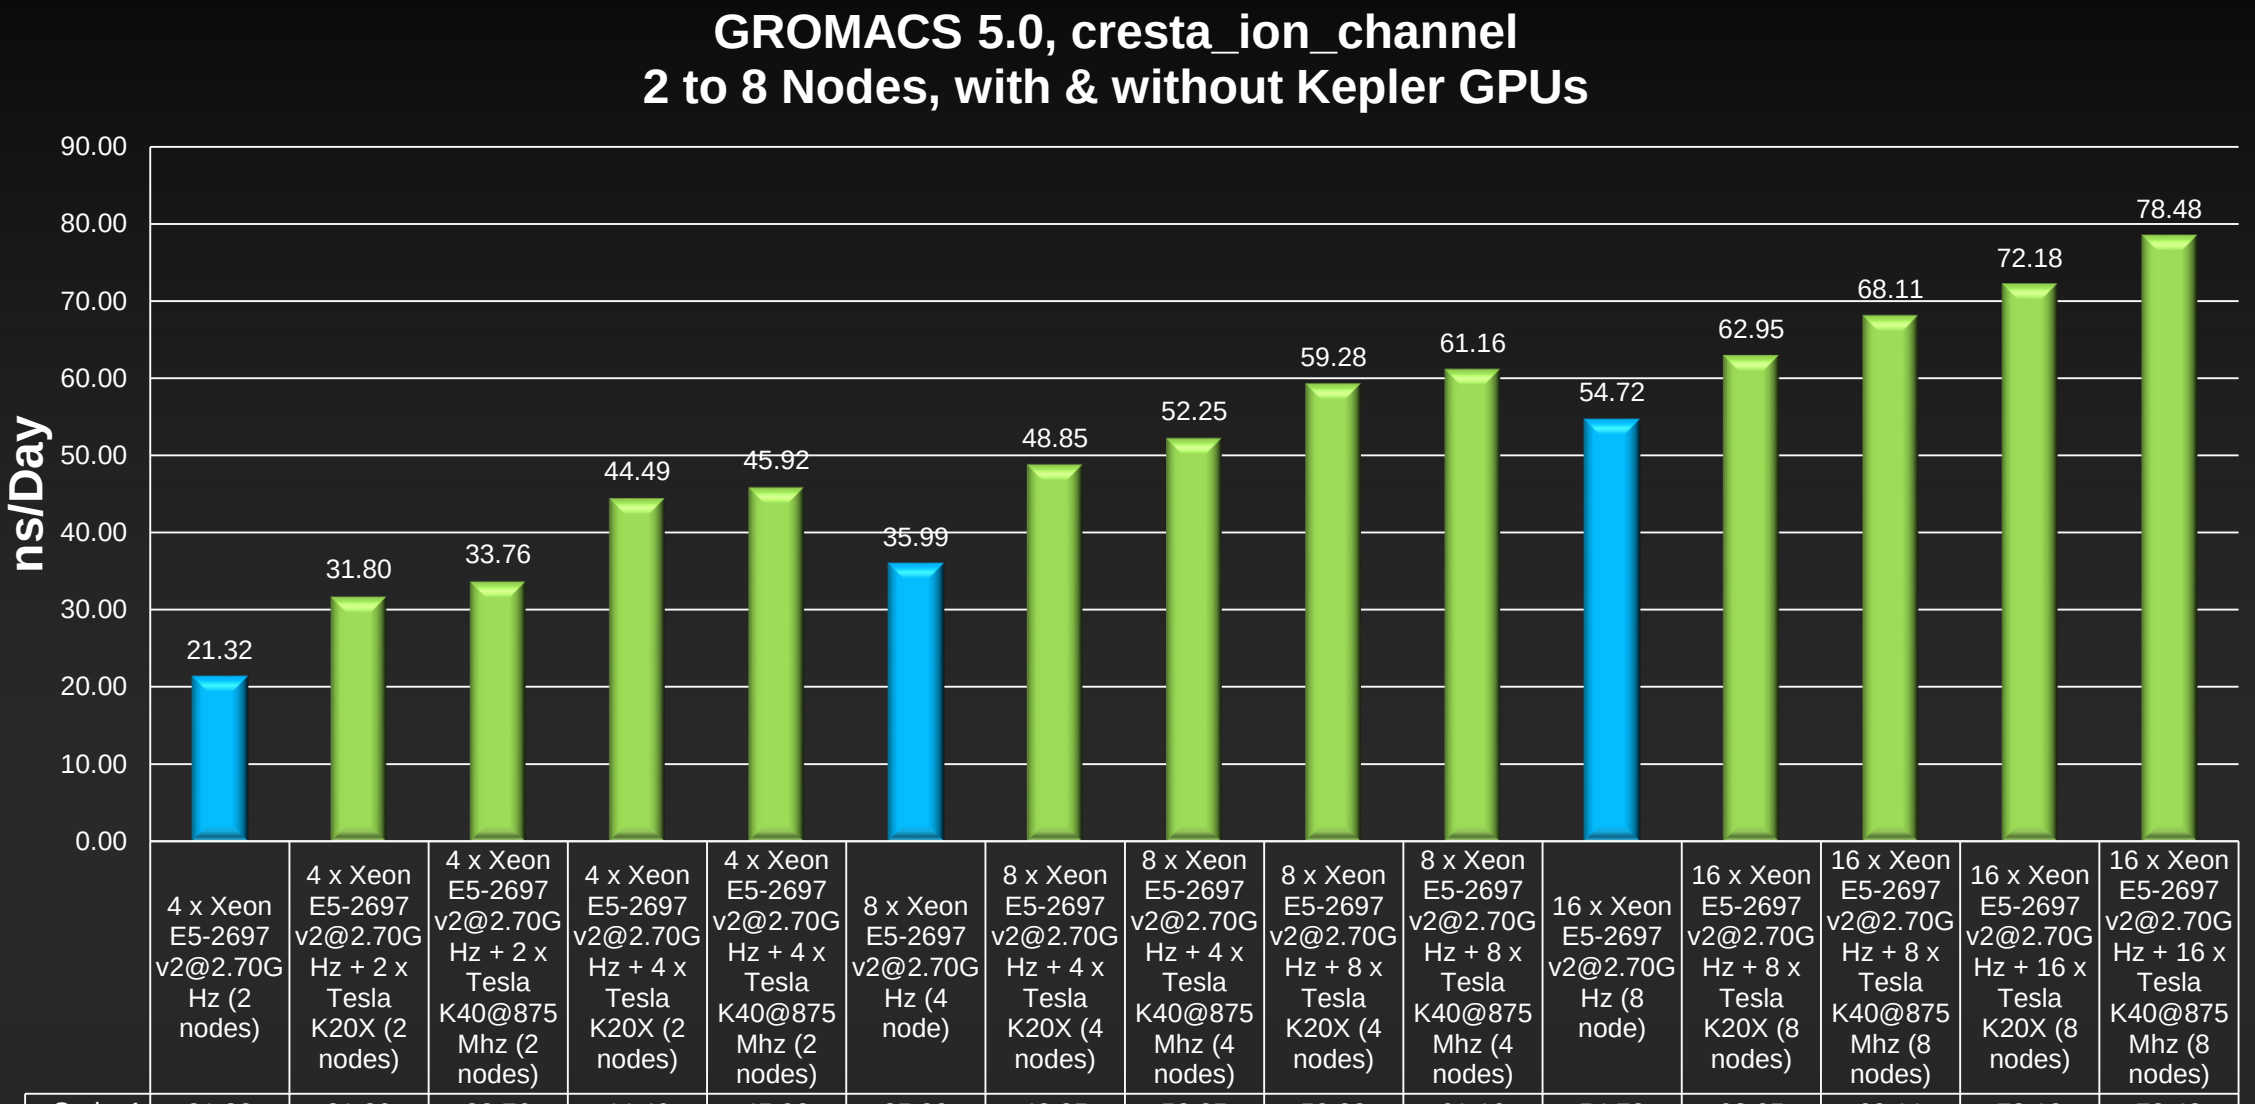

GROMACS 5.0 & Fastest Kepler GPUs yet!

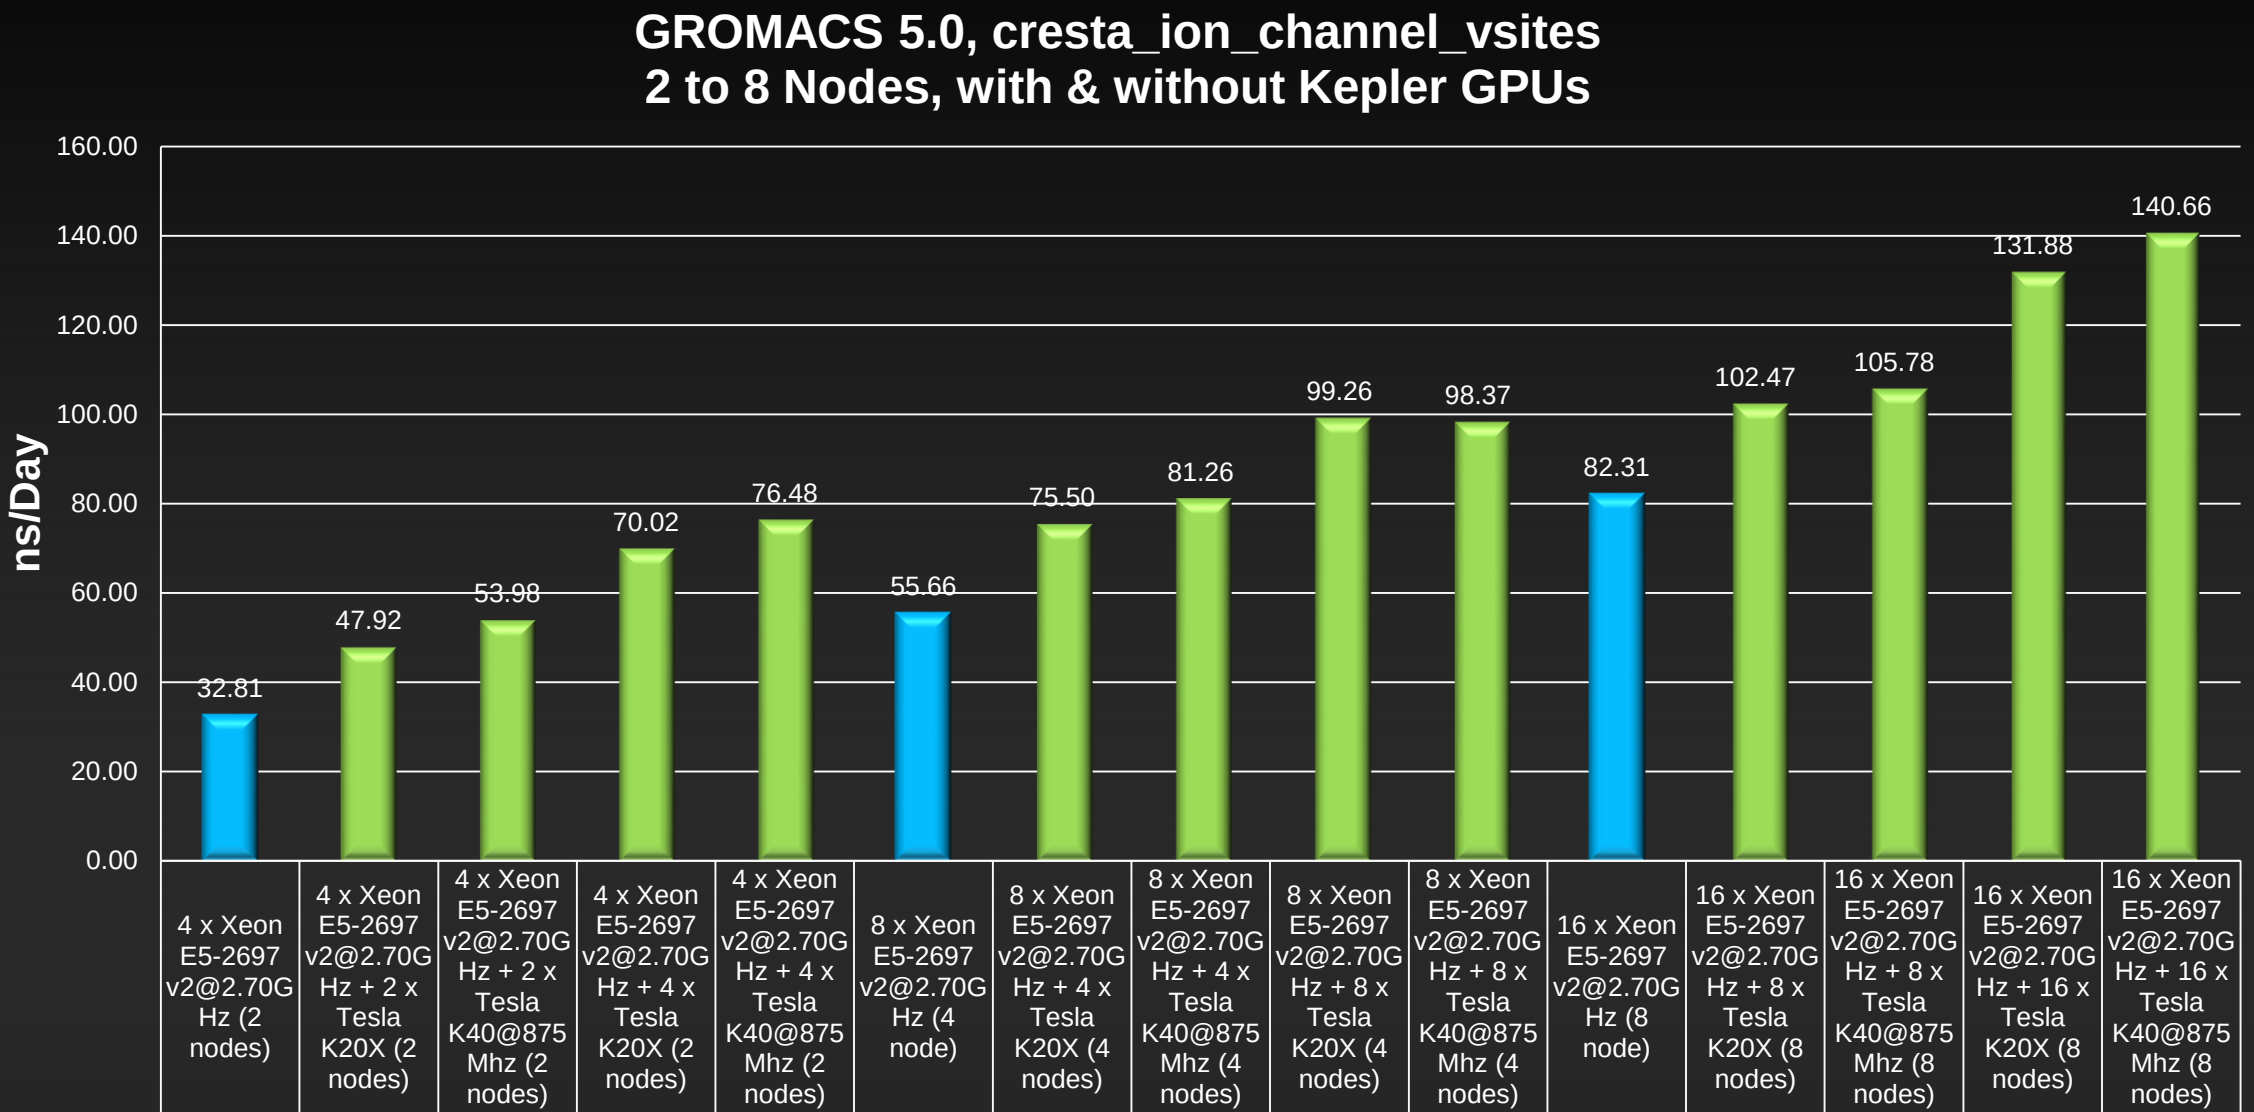

GROMACS 5.0 & Fastest Kepler GPUs yet!

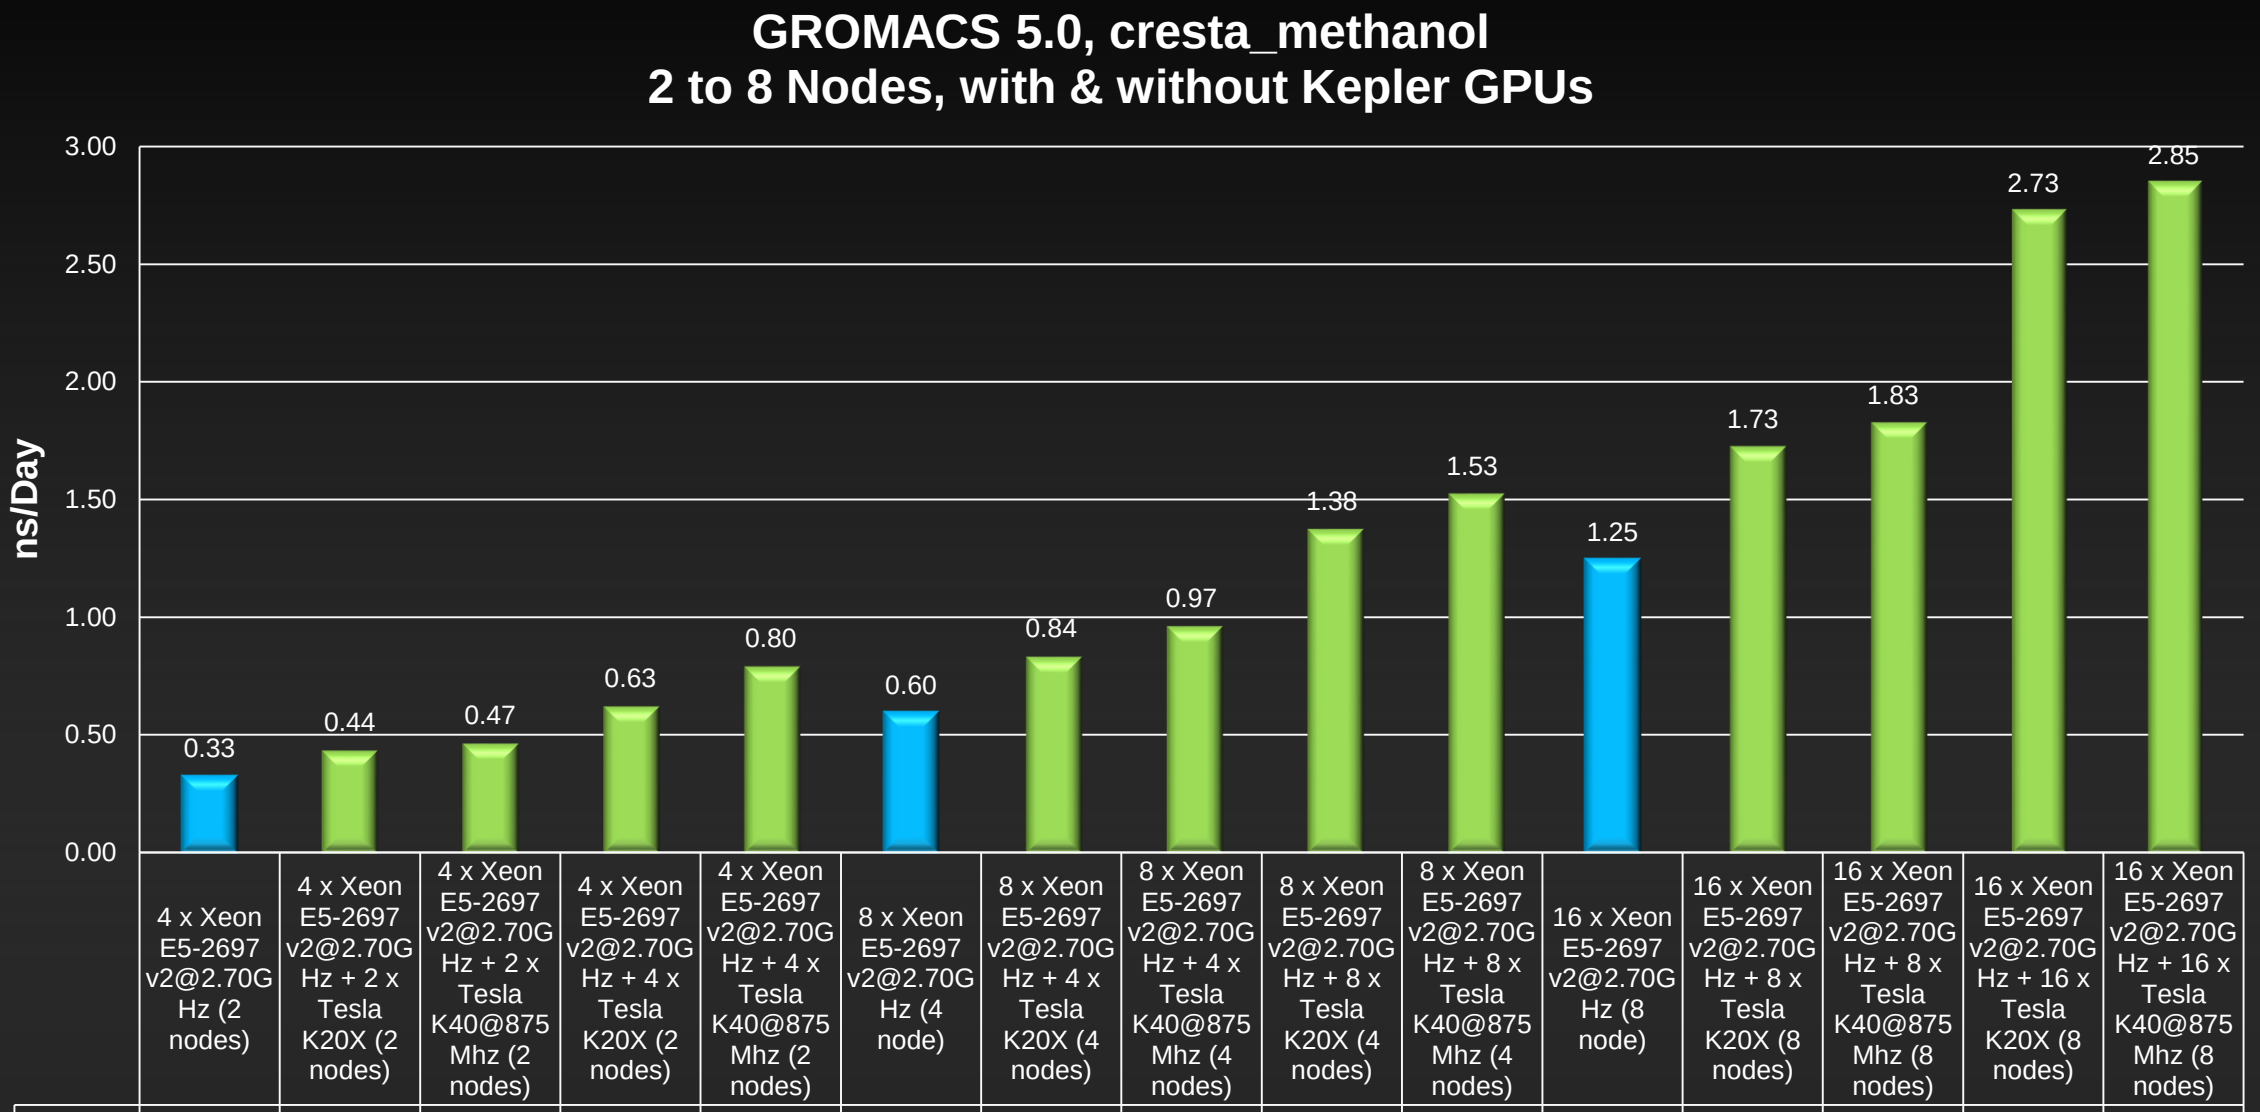

GROMACS 5.0 & Fastest Kepler GPUs yet!

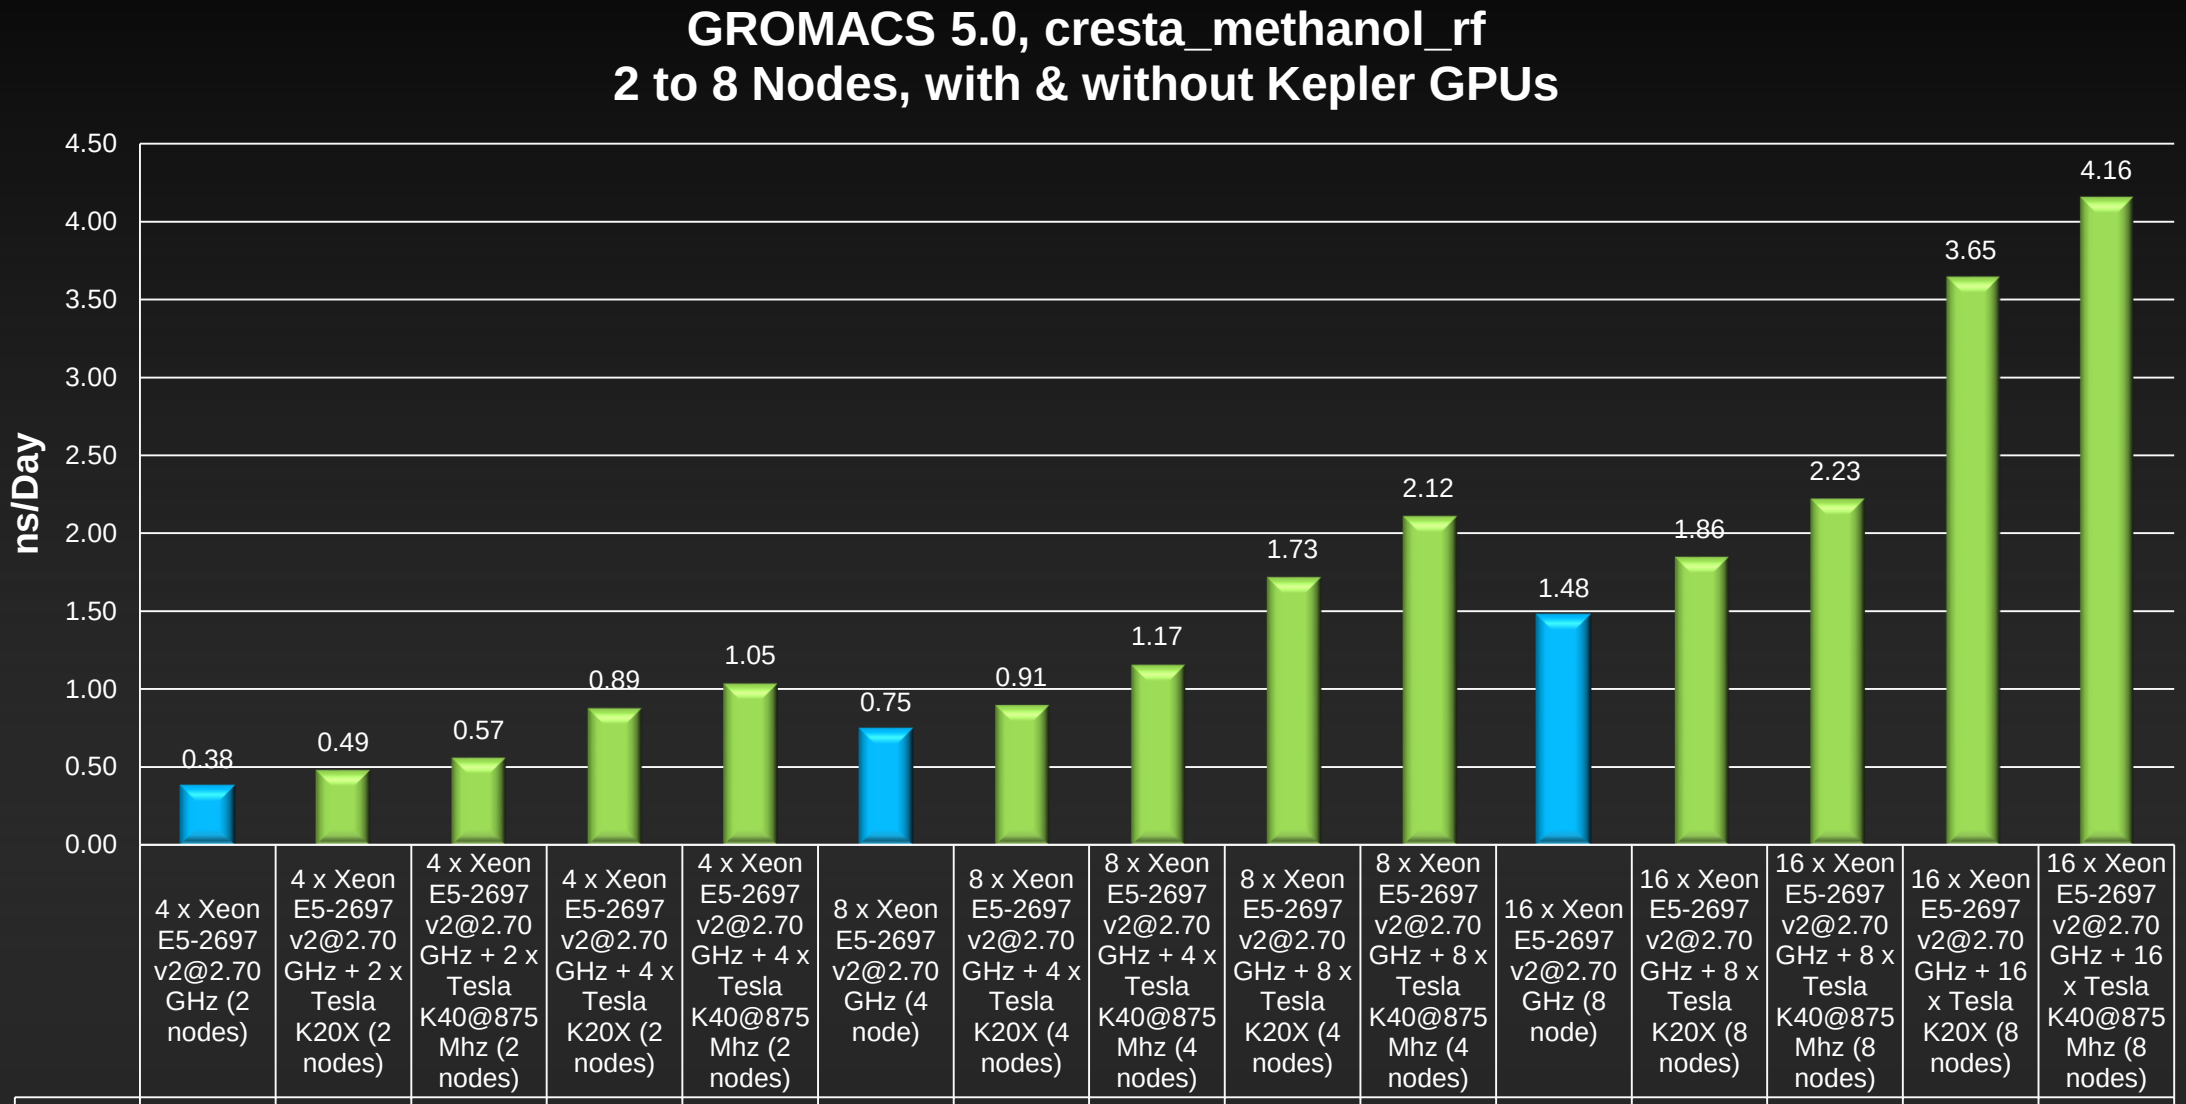

GROMACS 5.0 & Fastest Kepler GPUs yet!

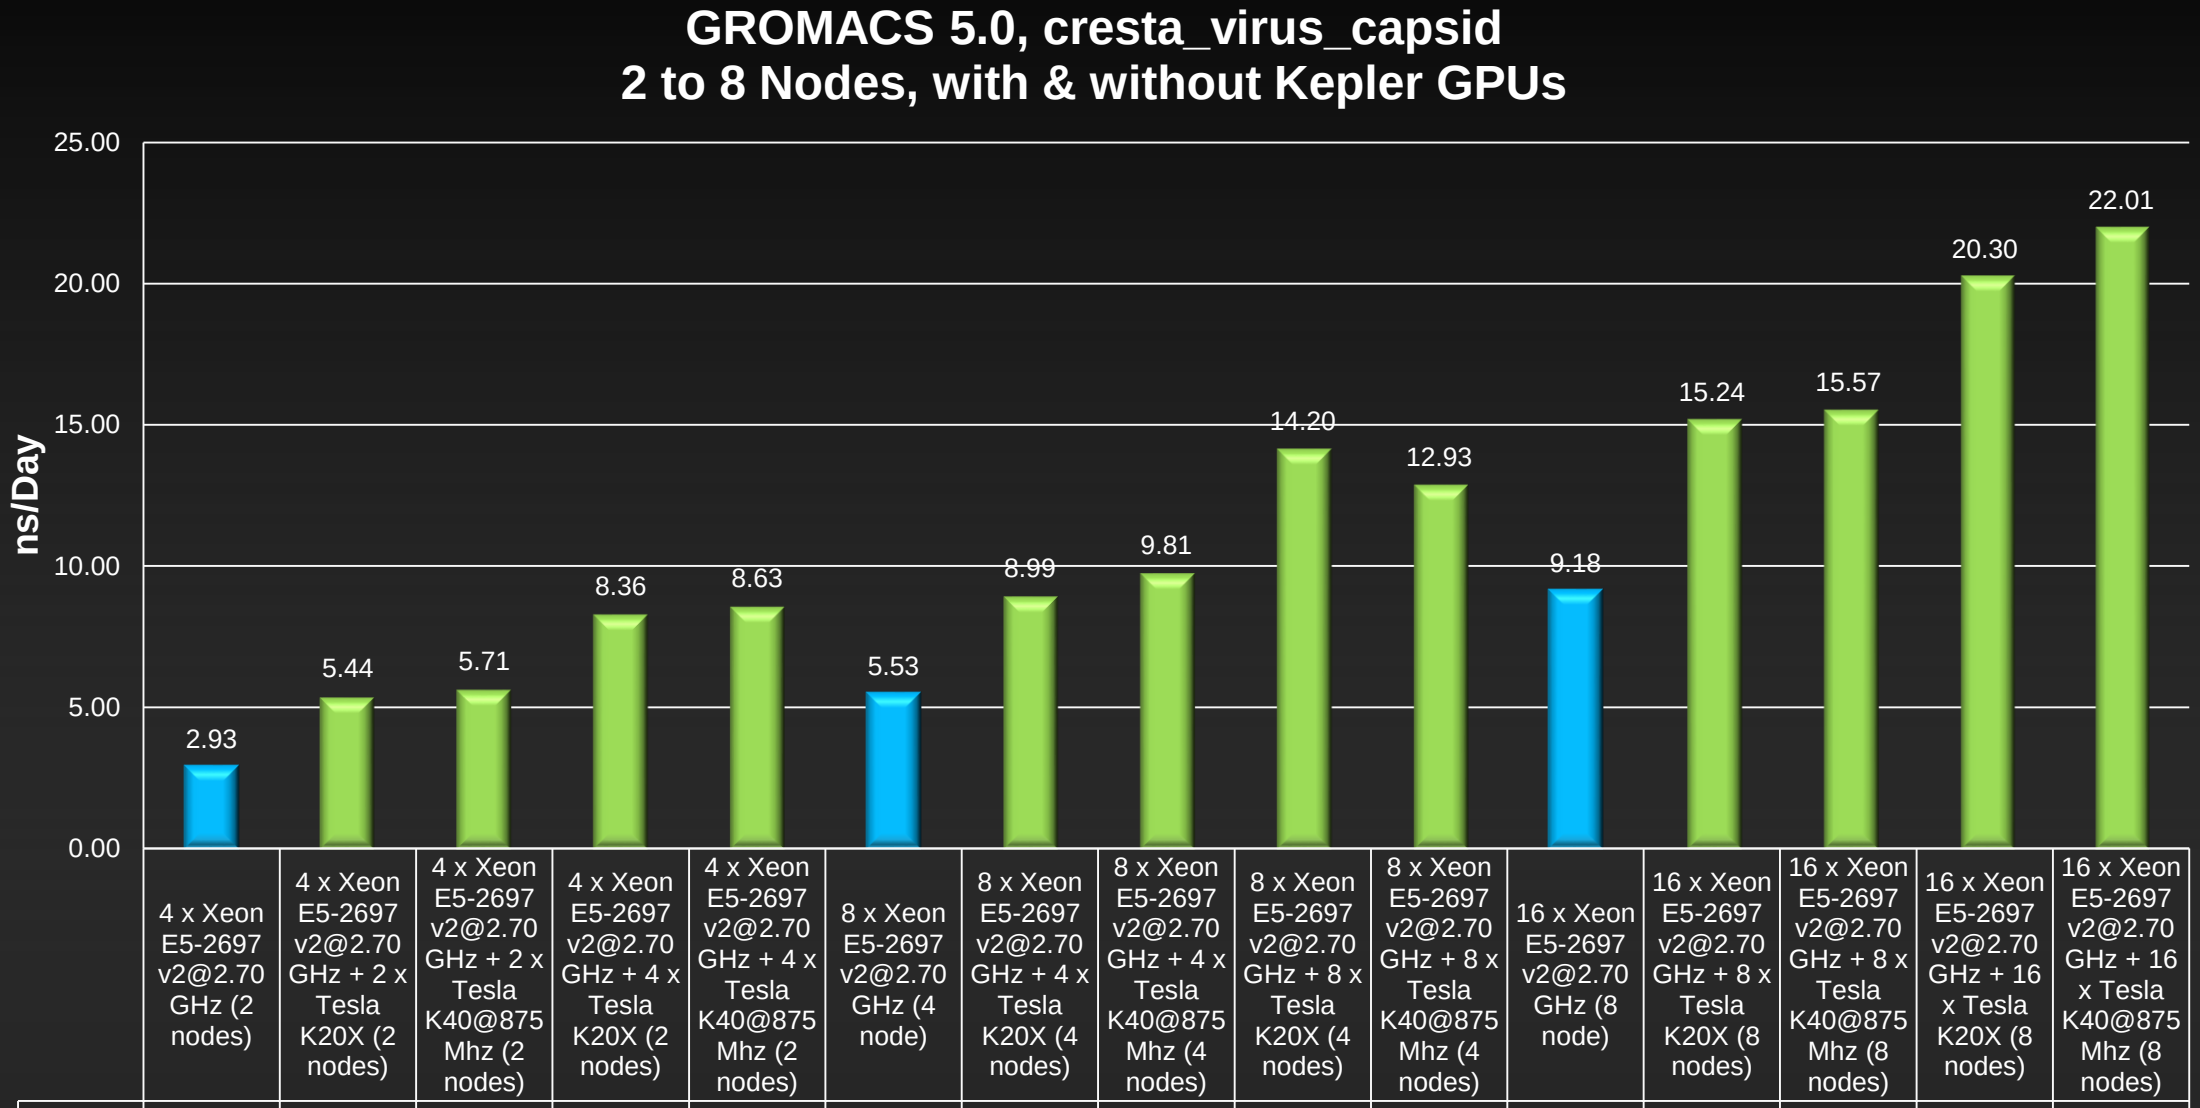

Slides - courtesy of GROMACS Dev Team

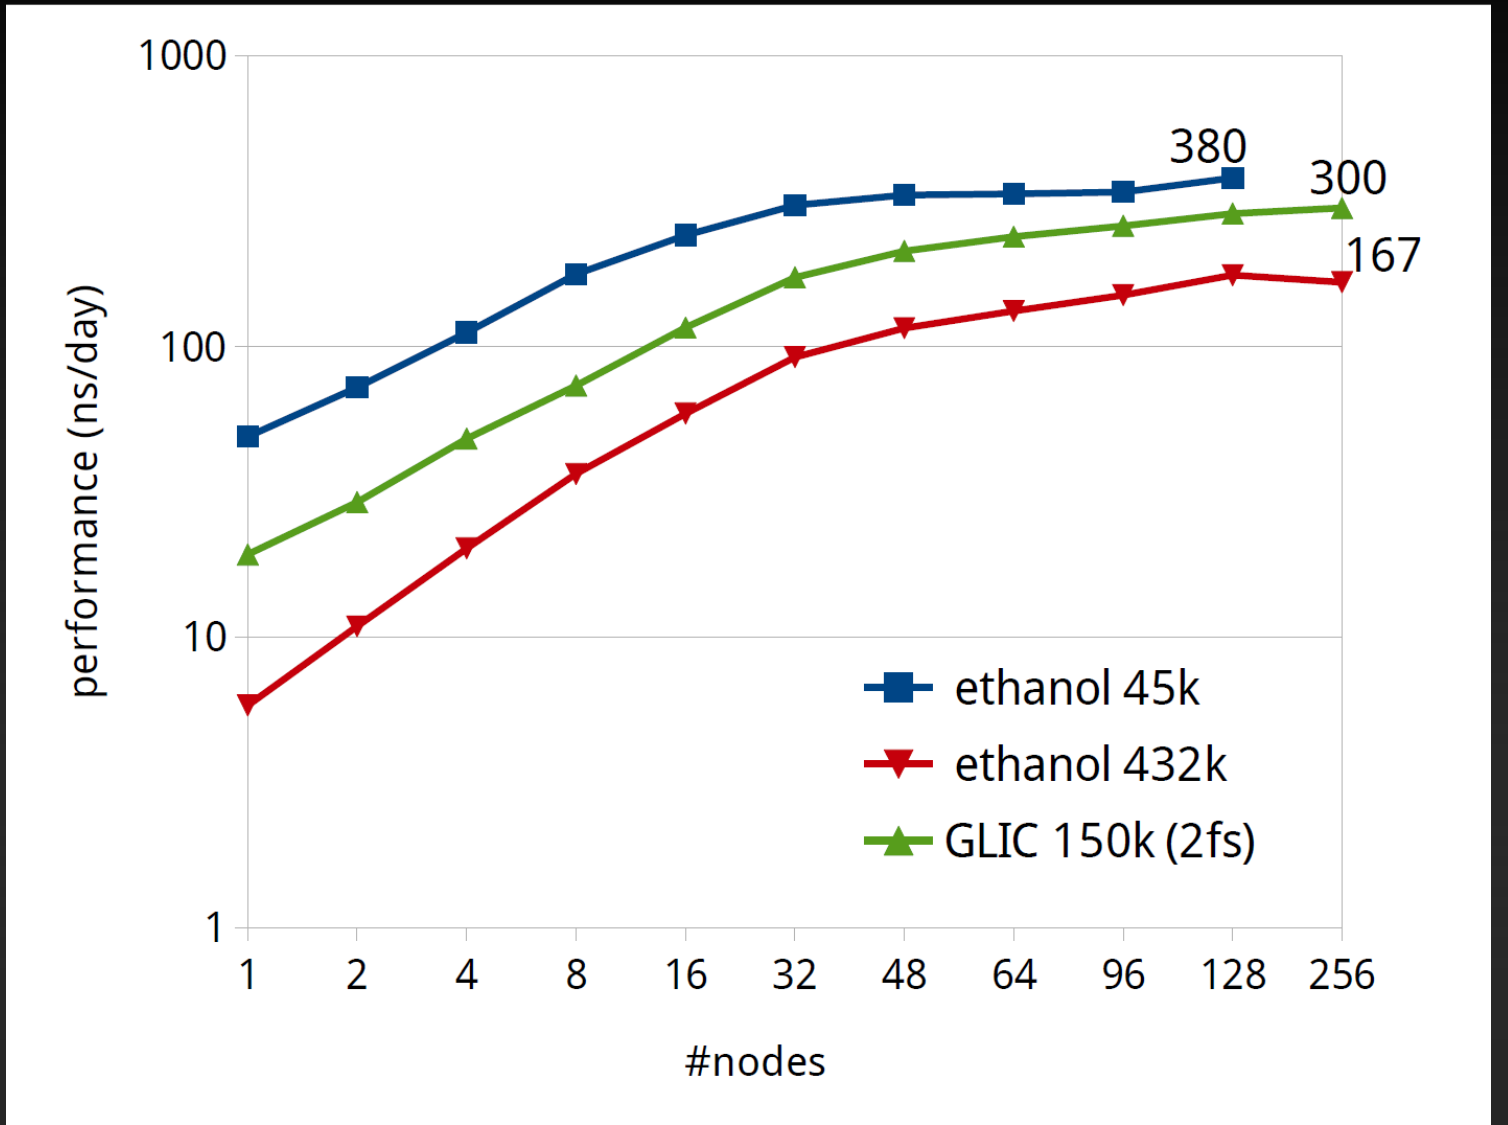

Slides - courtesy of GROMACS Dev Team

Kernel performance and scaling

Strong scaling regime:

This is where most of our efforts go!

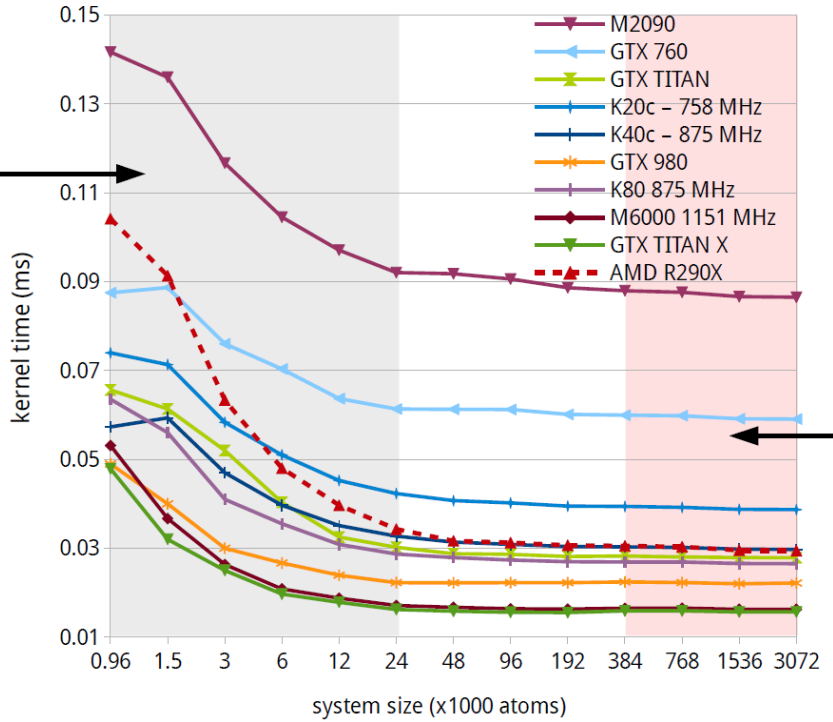

Benchmark "show-off" regime:

This is where the "free lunch" from new hardware comes in full effect

Slides - courtesy of GROMACS Dev Team

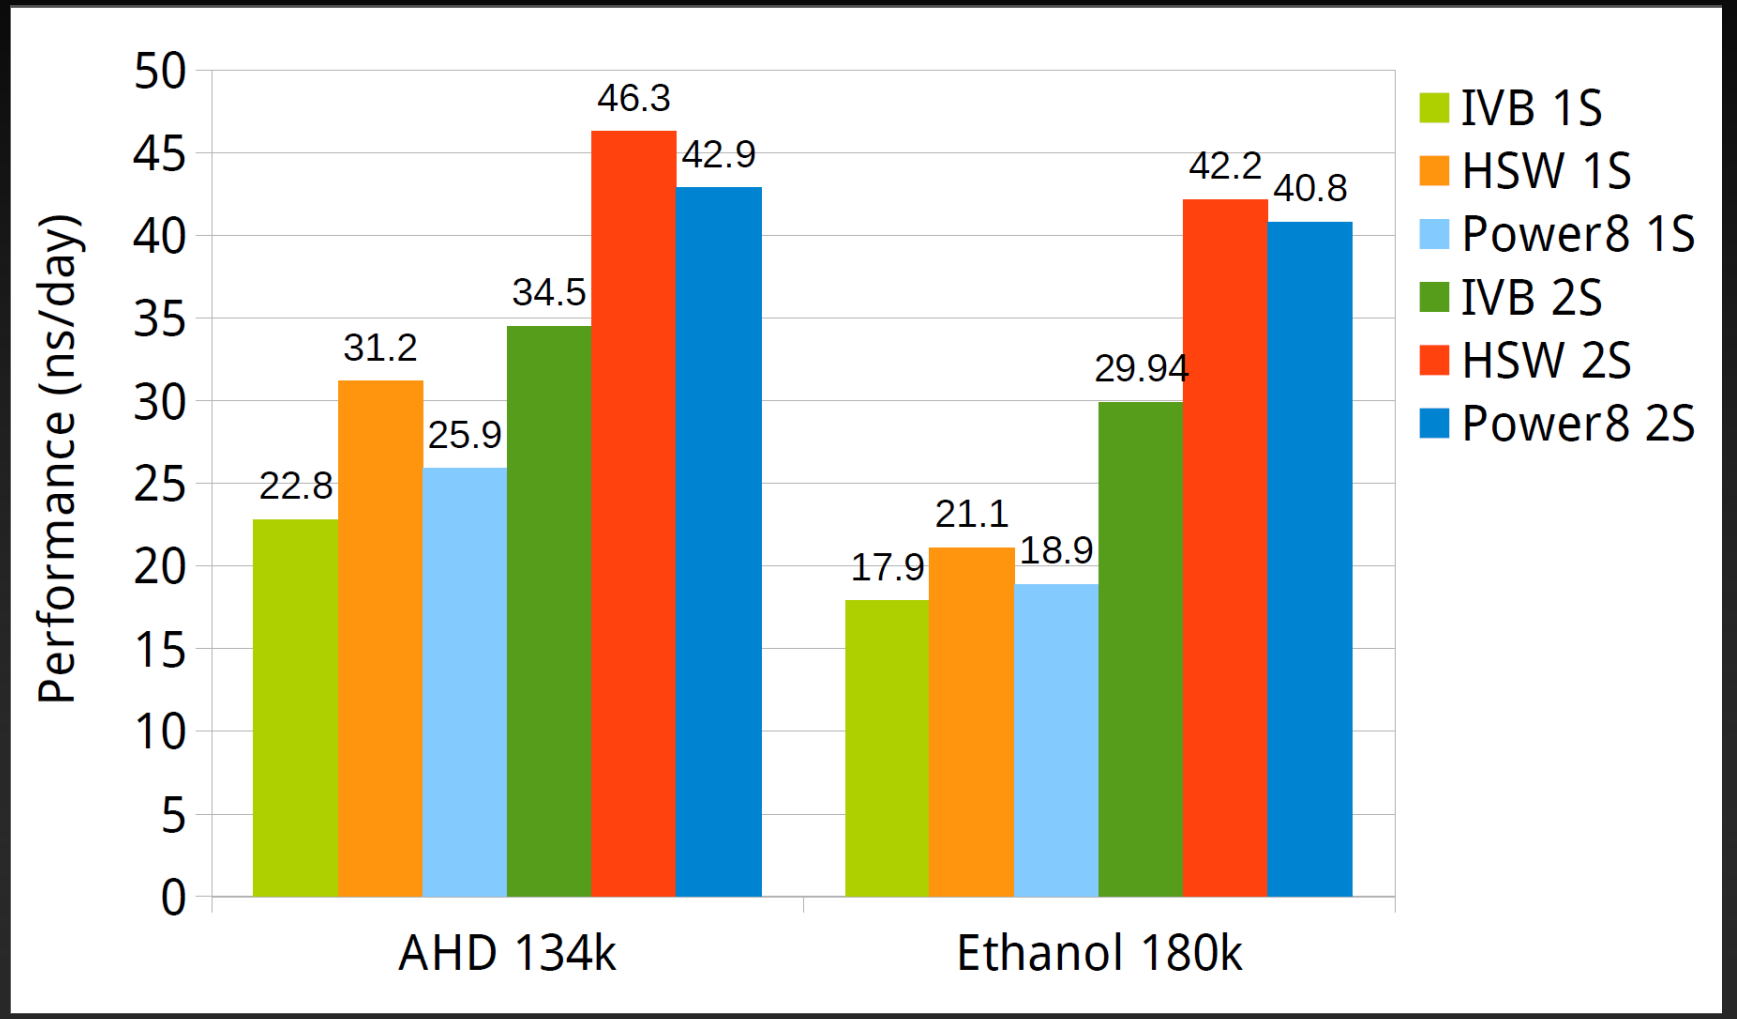

Slides - courtesy of GROMACS Dev Team

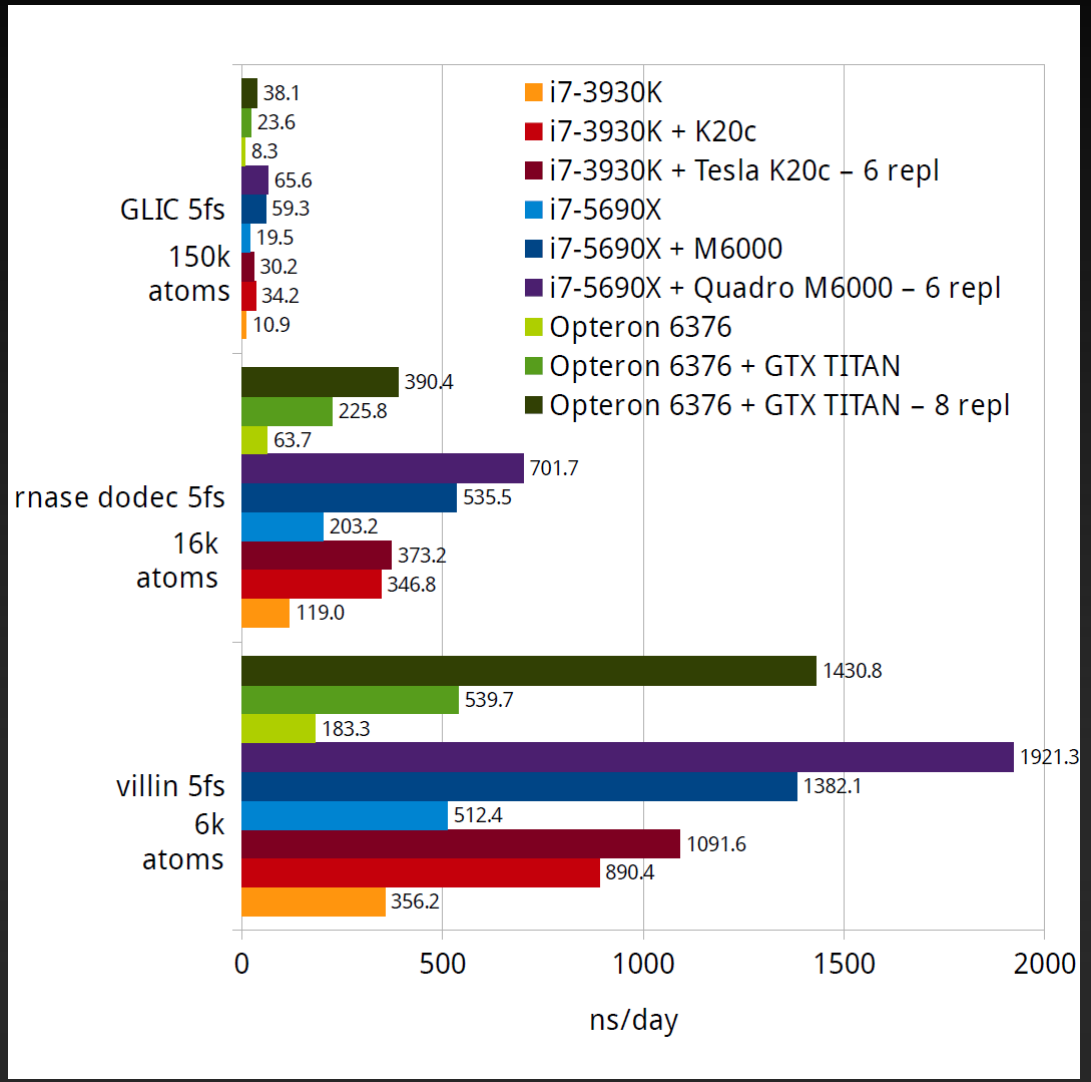

Greener Science

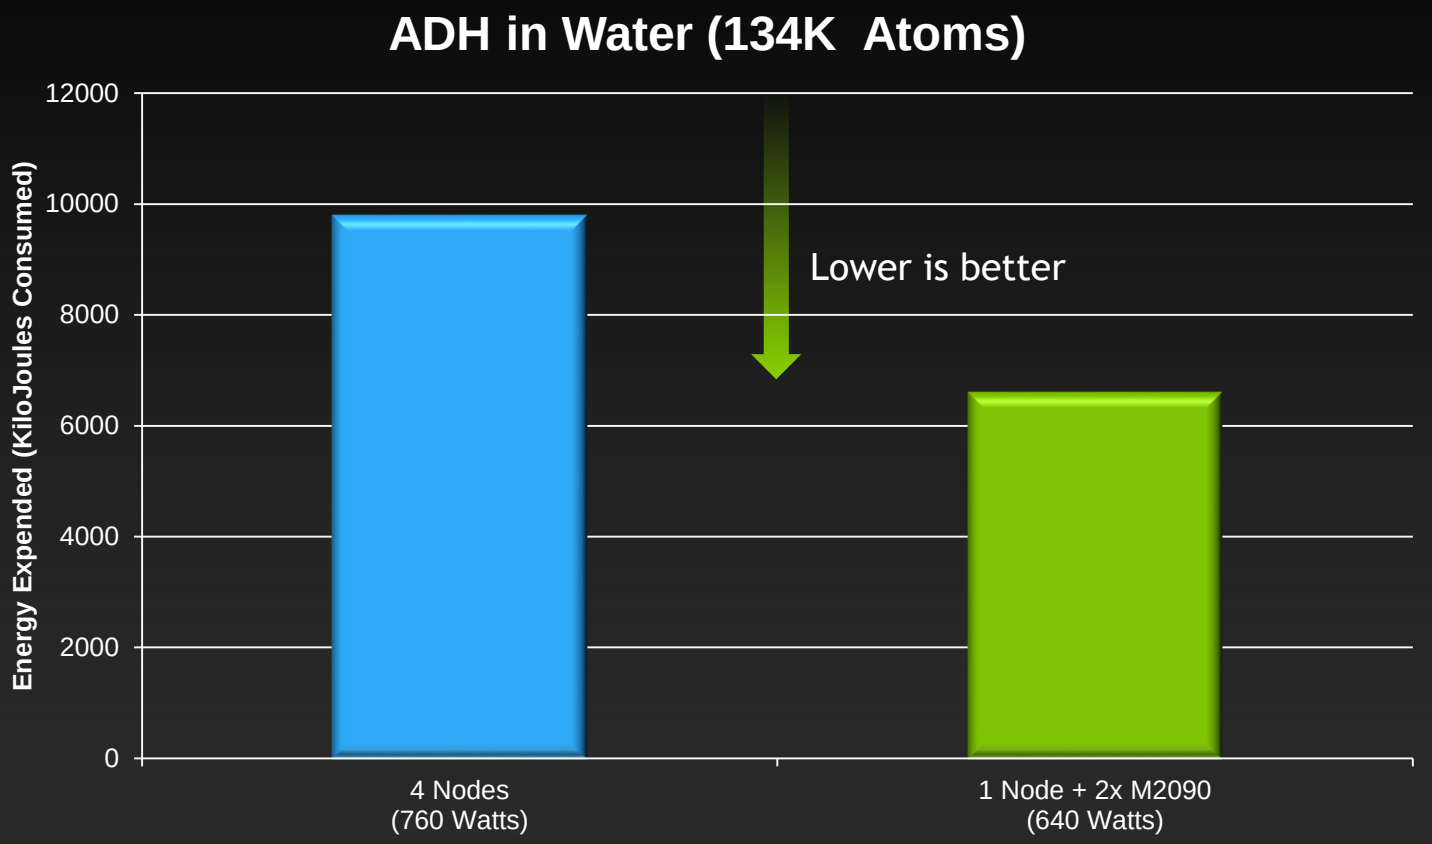

Running GROMACS 4.6 with CUDA 4.1

The blue nodes contain 2x Intel X5550 CPUs (95W TDP, 4 Cores per CPU)

The green node contains 2x Intel X5550 CPUs, 4 Cores per CPU) and 2x NVIDIA M2090s GPUs (225W TDP per GPU)

Energy Expended
= Power x Time

In simulating each nanosecond, the GPU-accelerated system uses 33% less energy

Recommended GPU Node Configuration for GROMACS Computational Chemistry

Workstation or Single Node Configuration	
# of CPU sockets	2
Cores per CPU socket	6+
CPU speed (Ghz)	2.66+
System memory per socket (GB)	32
GPUs	Kepler K20, K40, K80
# of GPUs per CPU socket	1x Kepler GPUs: need fast Sandy Bridge or Ivy Bridge, or high-end AMD Opterons
GPU memory preference (GB)	6
GPU to CPU connection	PCIe 3.0 or higher
Server storage	500 GB or higher
Network configuration	Gemini, InfiniBand

HOOMD-Blue

March 2016

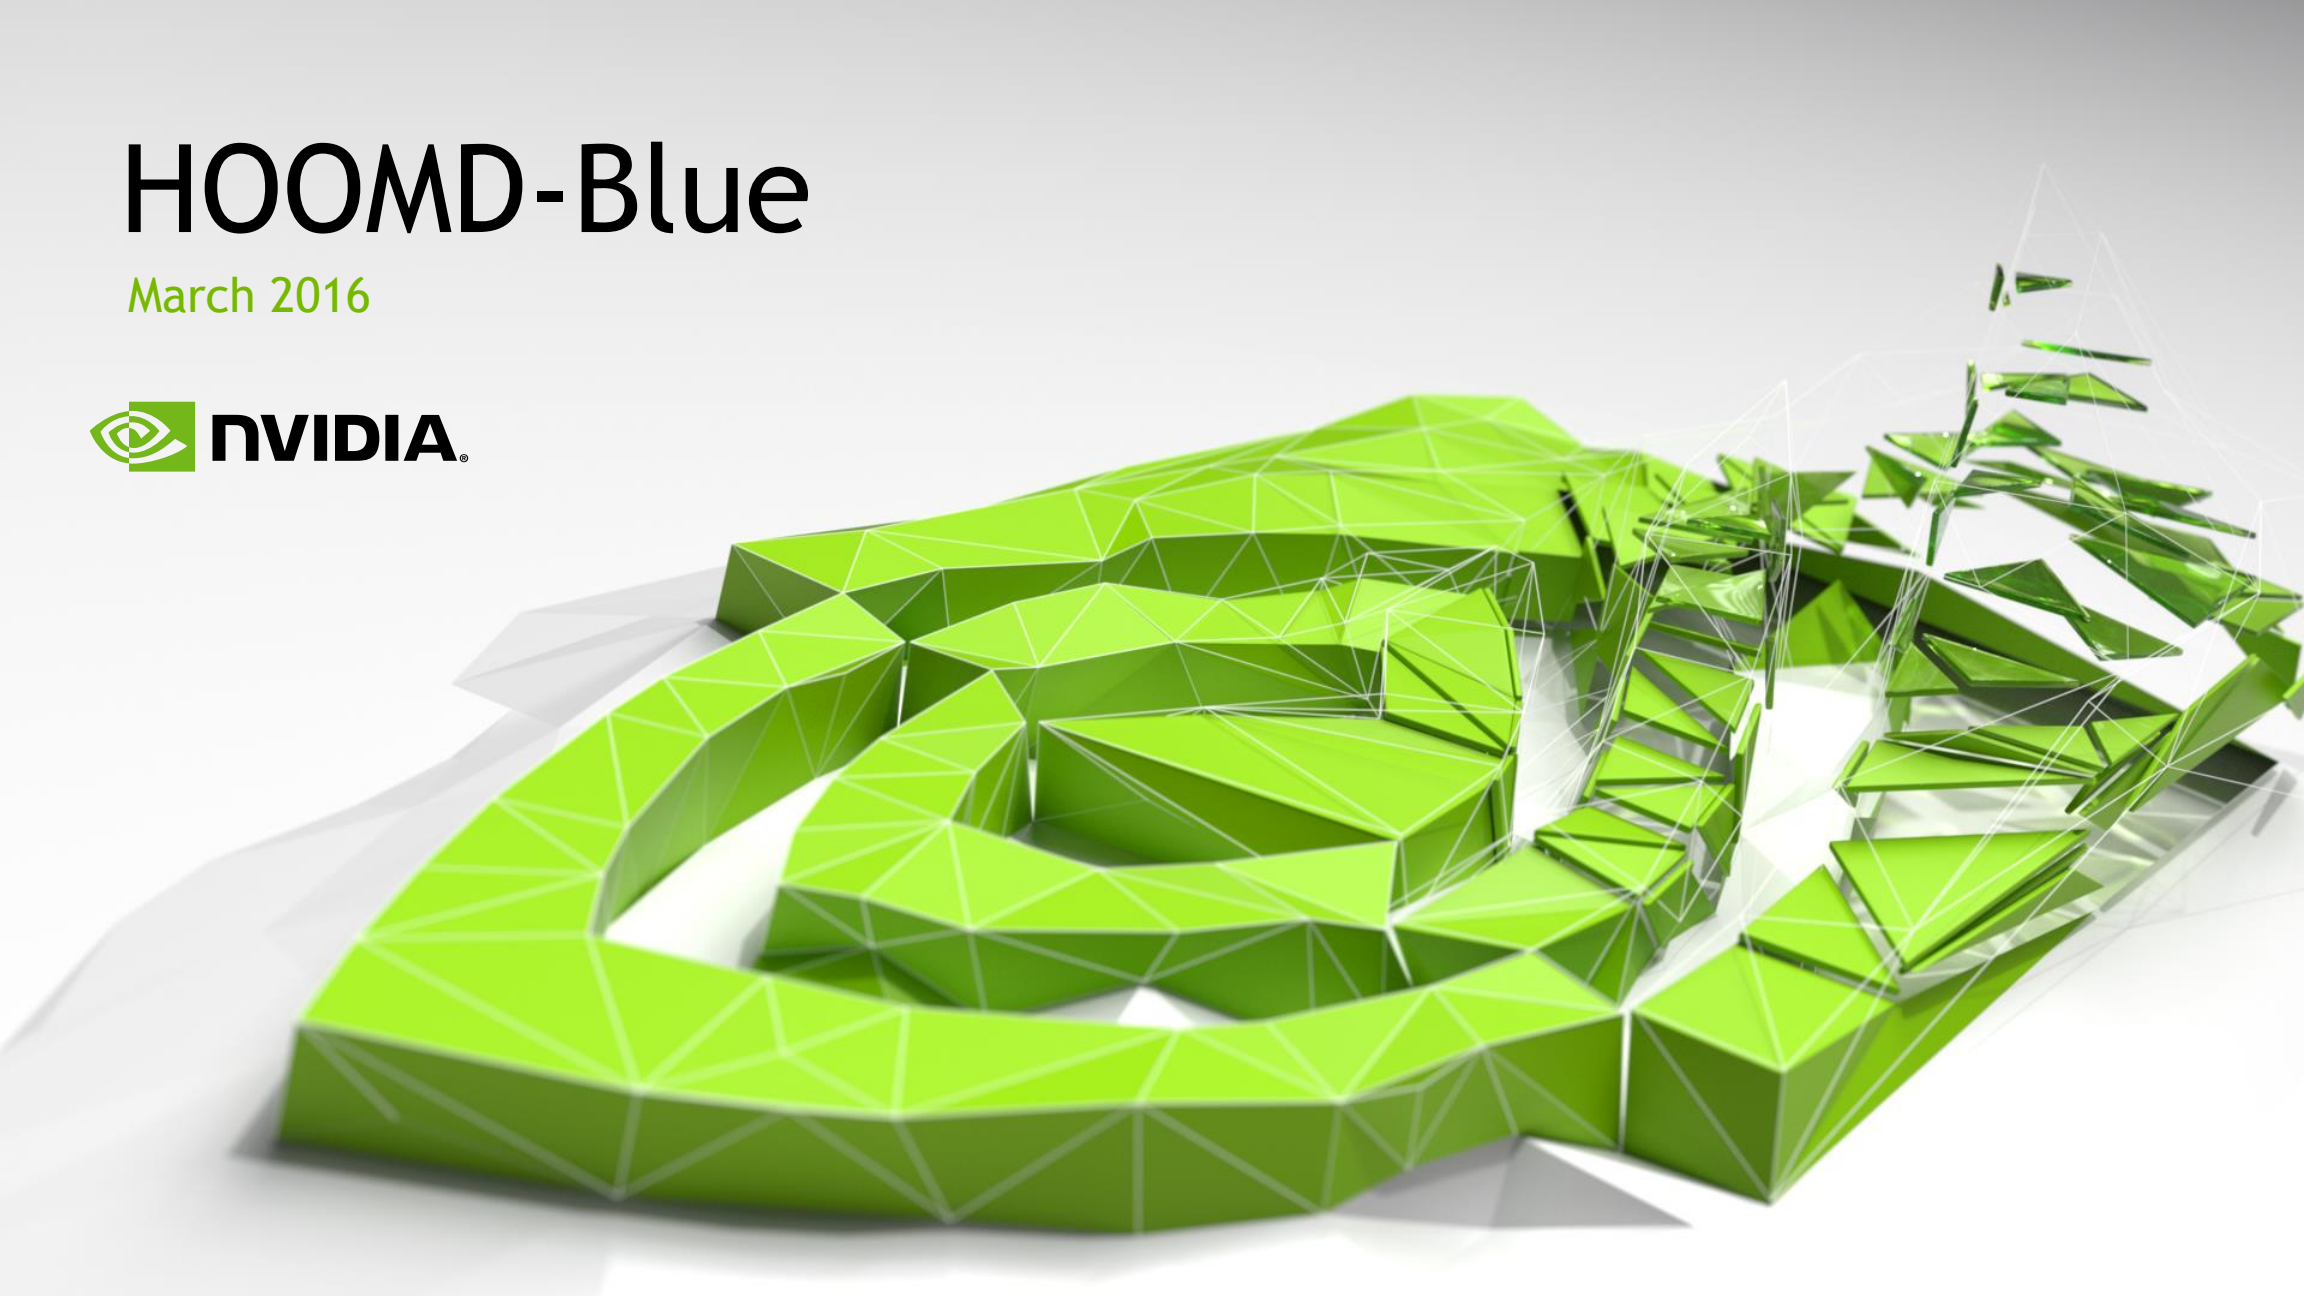

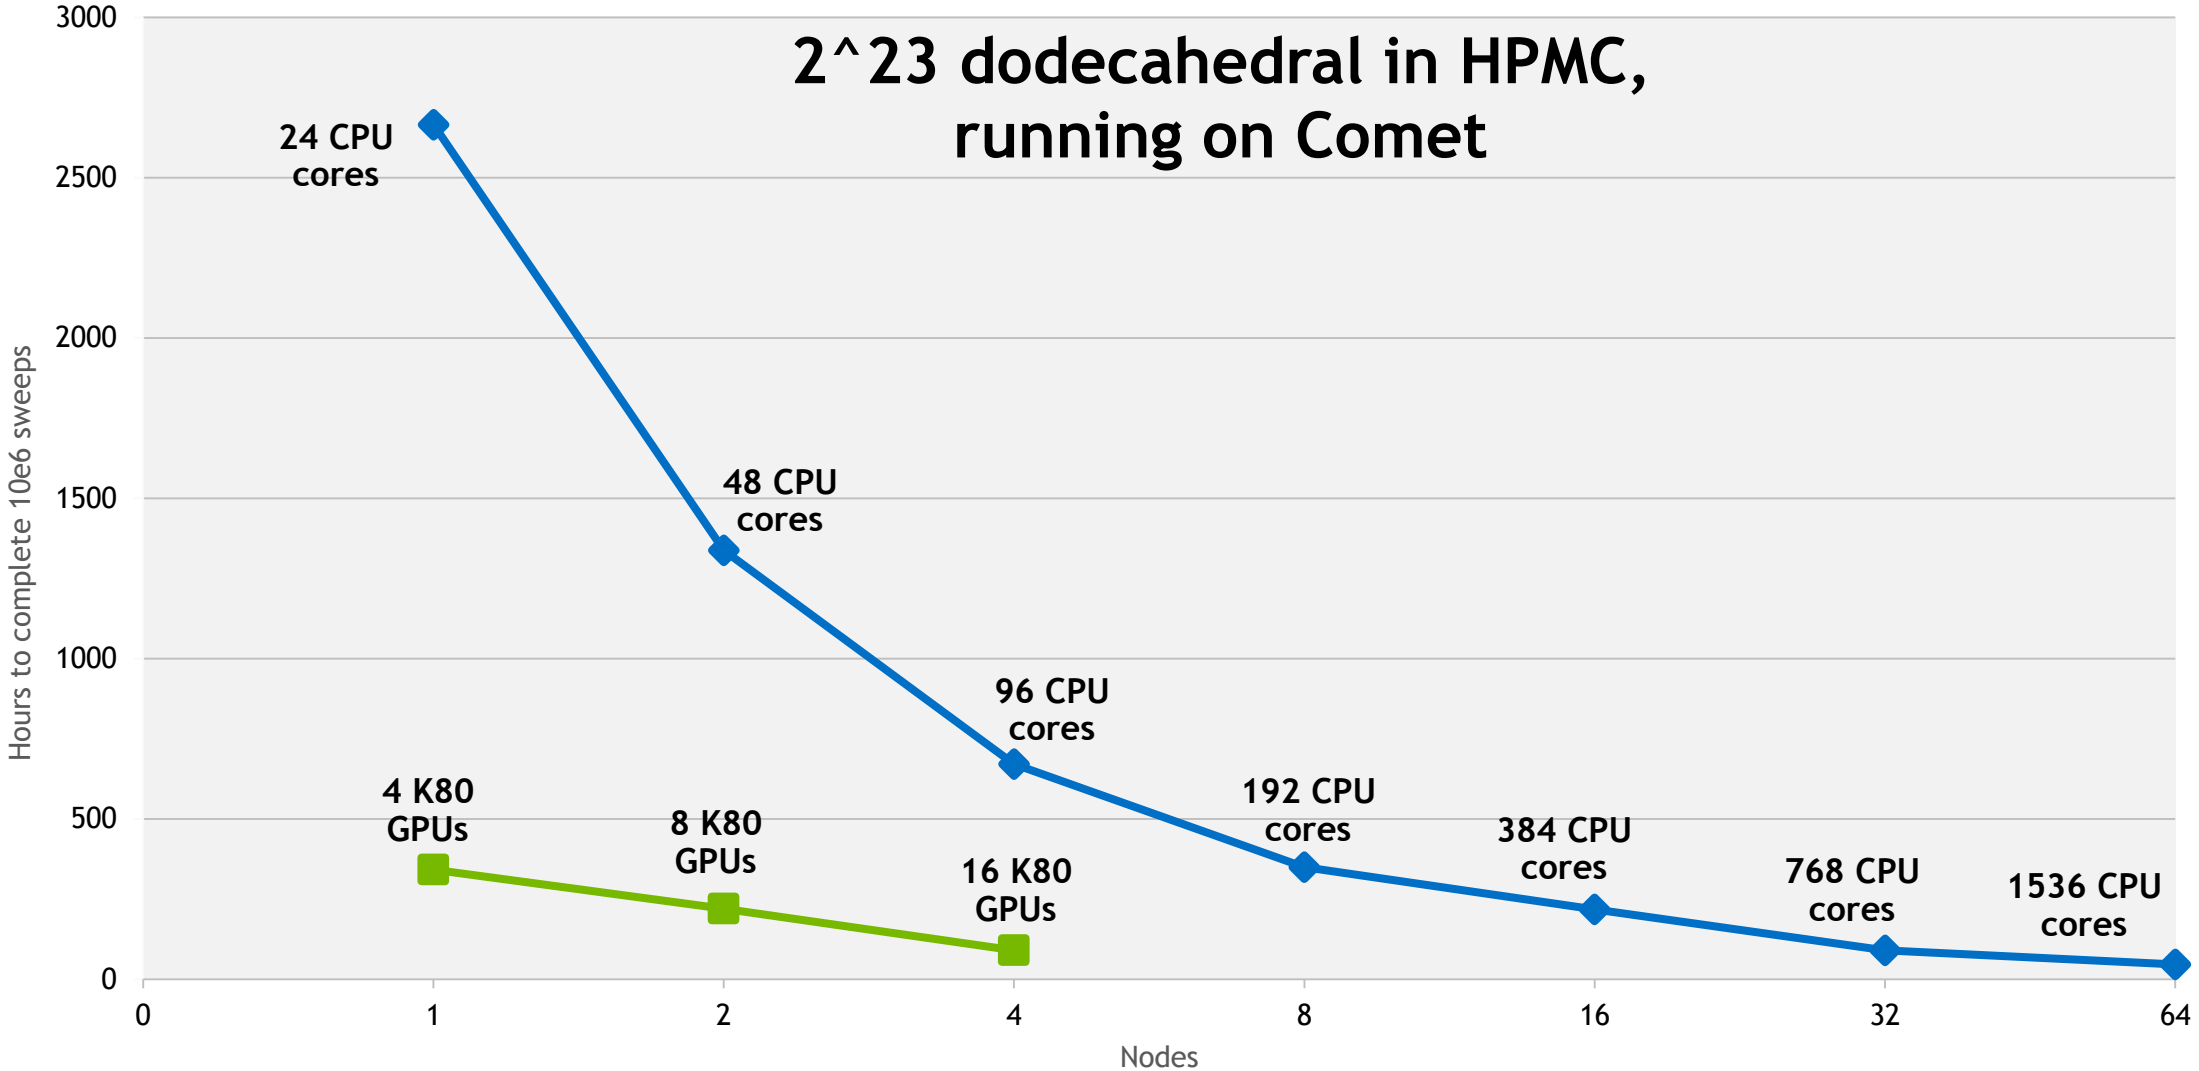

Blue nodes contain Dual Intel Xeon E5-2680 v3@2.50 GHz (Haswell) CPUs
Green nodes contain Tesla K80 (autoboost) GPUs

HOOMD-BLUE 1.0

October 2015

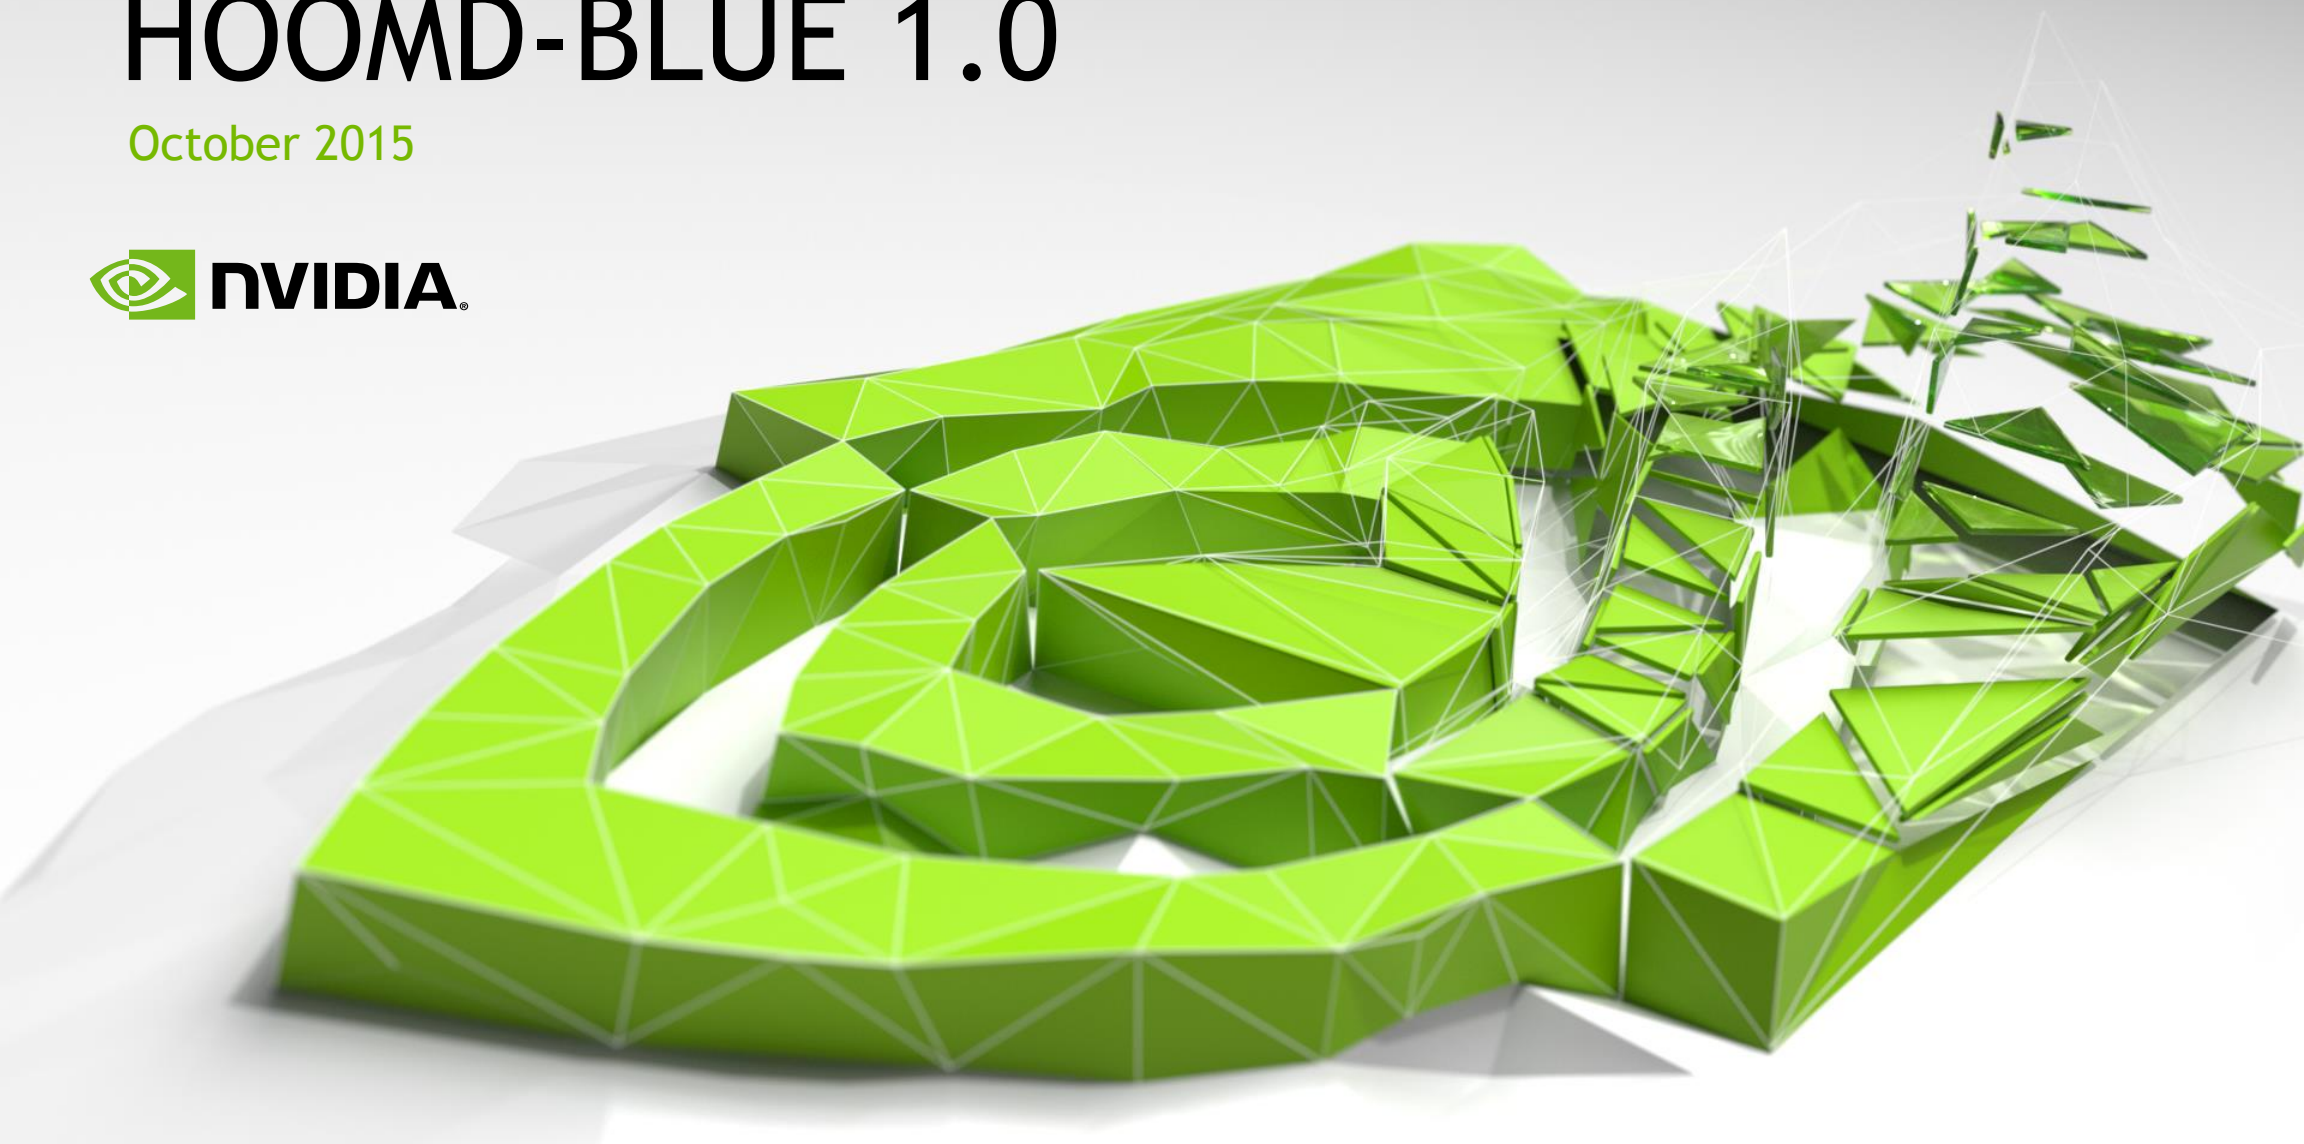

HOOMD-Blue 1.0, K40 & K80, Boost impact!

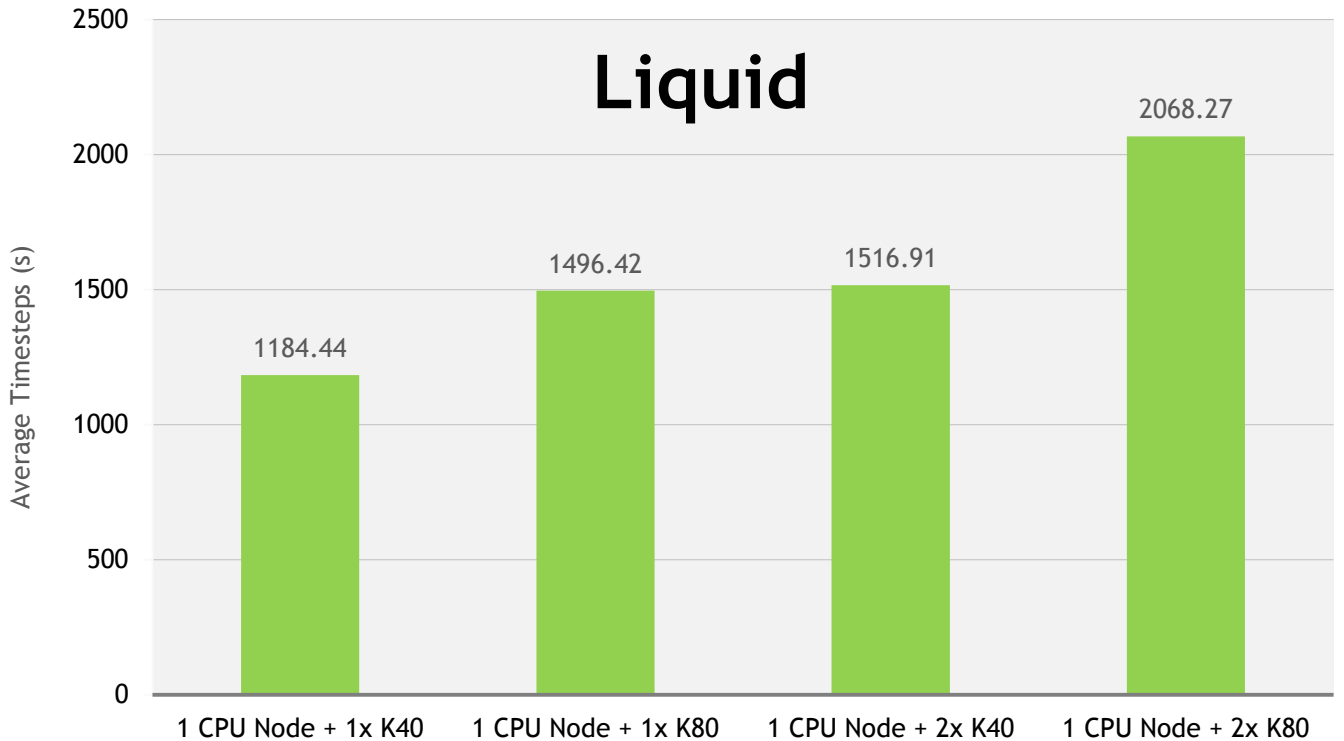

Running HOOMD-Blue version 1.0

The green nodes contain Dual Intel E5-2697 v2@2.70GHz CPUs + either NVIDIA Tesla K40@875Mhz or Tesla K80 (autoboost) GPUs

HOOMD-Blue 1.0, K40 & K80, Boost impact!

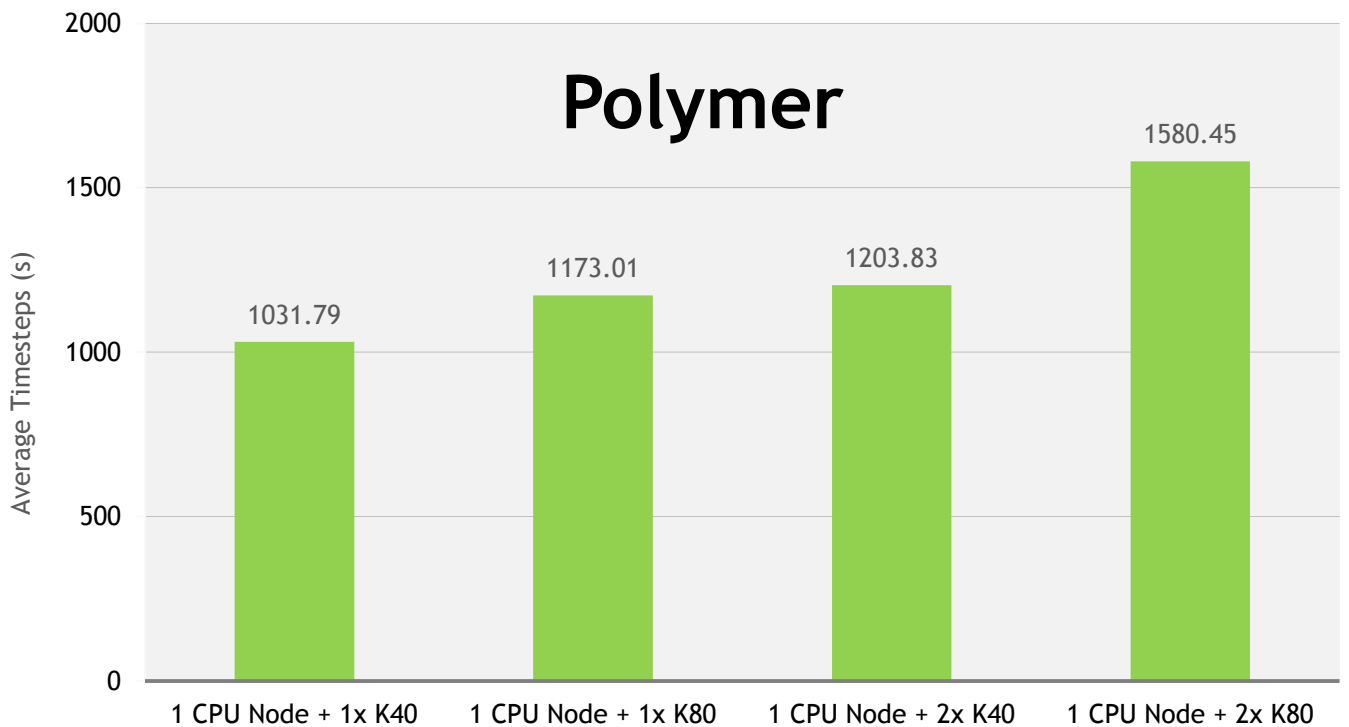

Running HOOMD-Blue version 1.0

The green nodes contain Dual Intel E5-2697 v2@2.70GHz CPUs + either NVIDIA Tesla K40@875Mhz or Tesla K80 (autoboost) GPUs

HOOMD-Blue 1.0, K40 & K80, Boost impact!

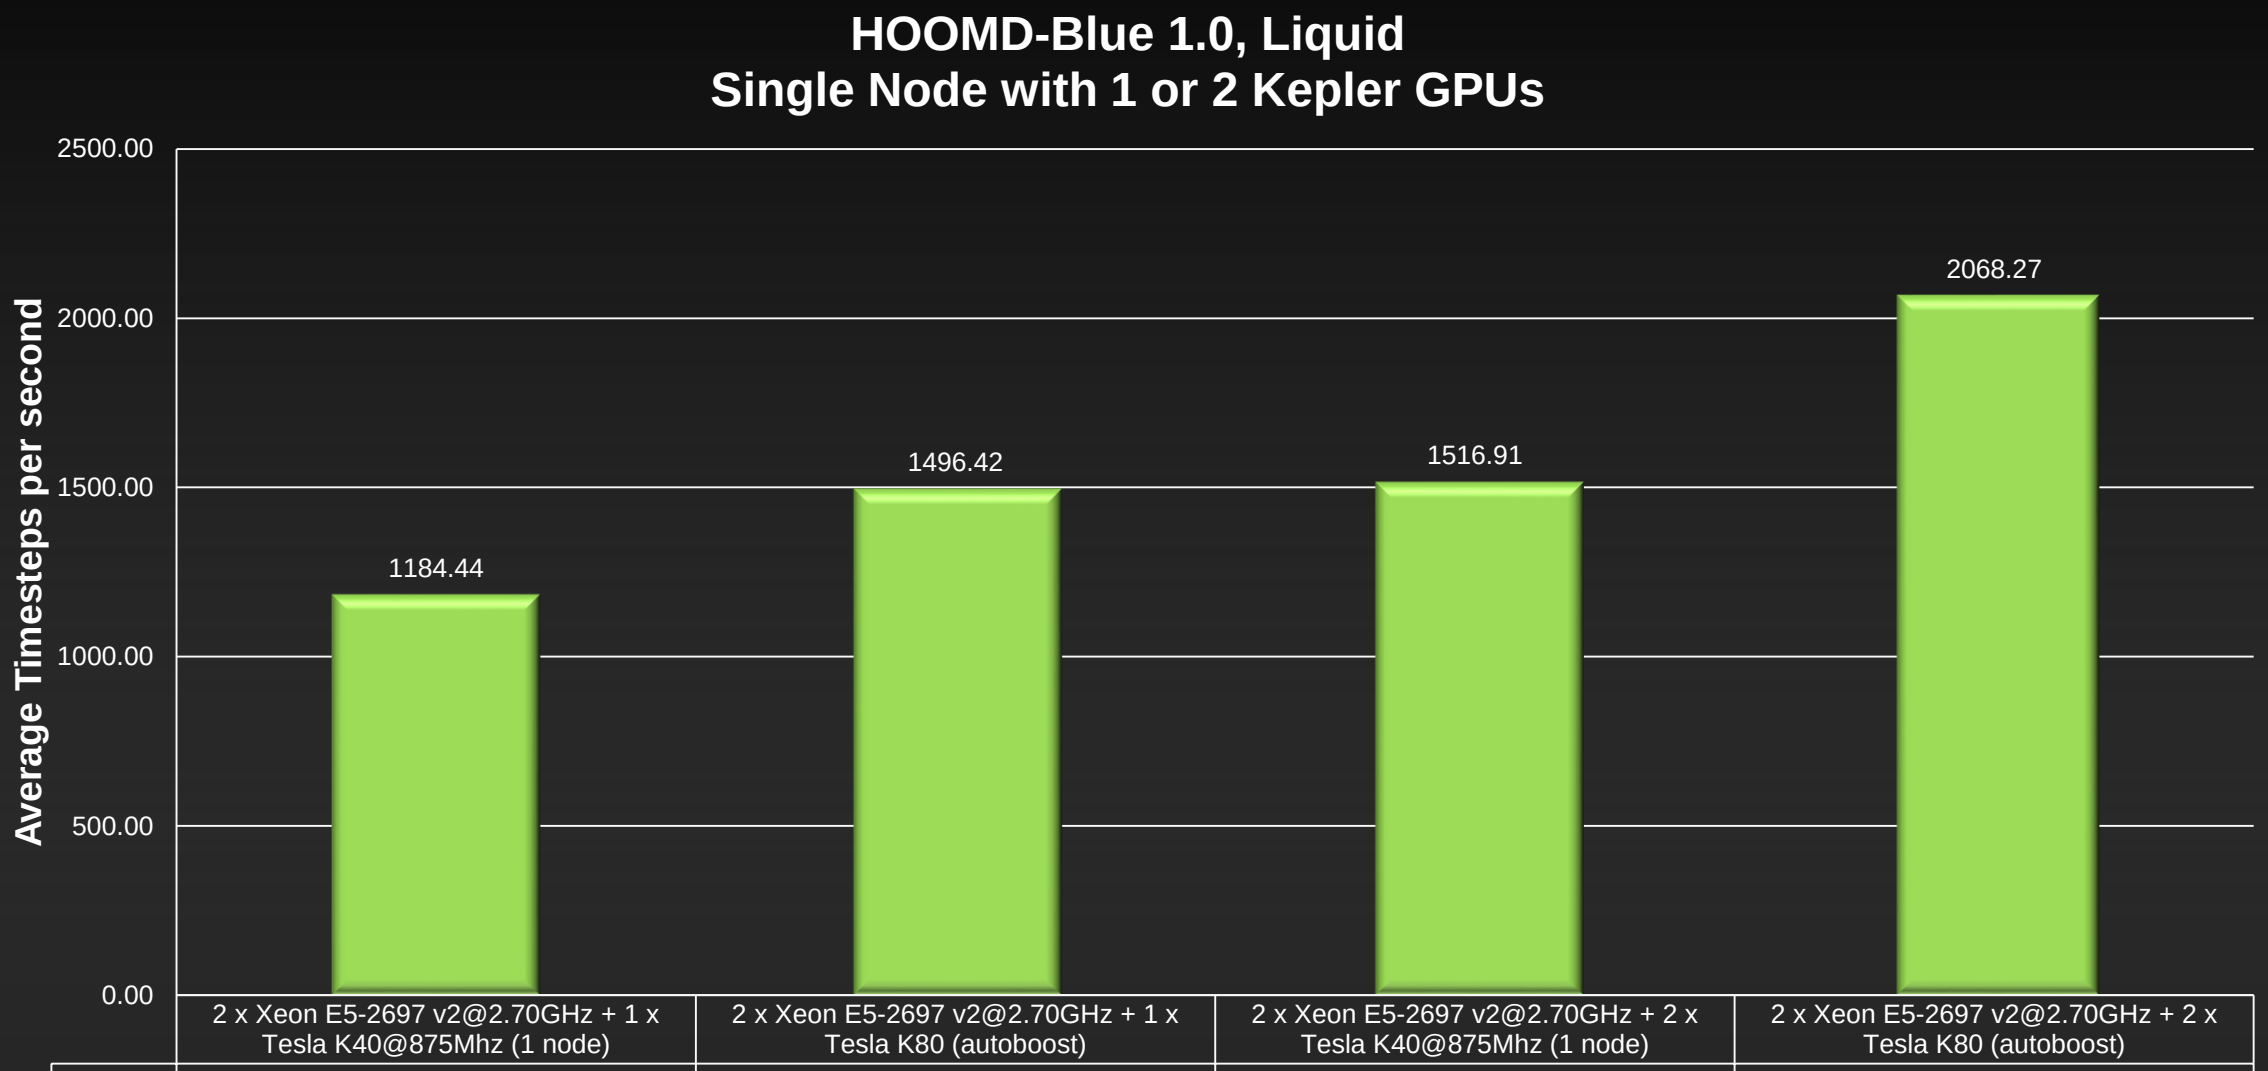

HOOMD-Blue 1.0, K40 & K80, Boost impact!

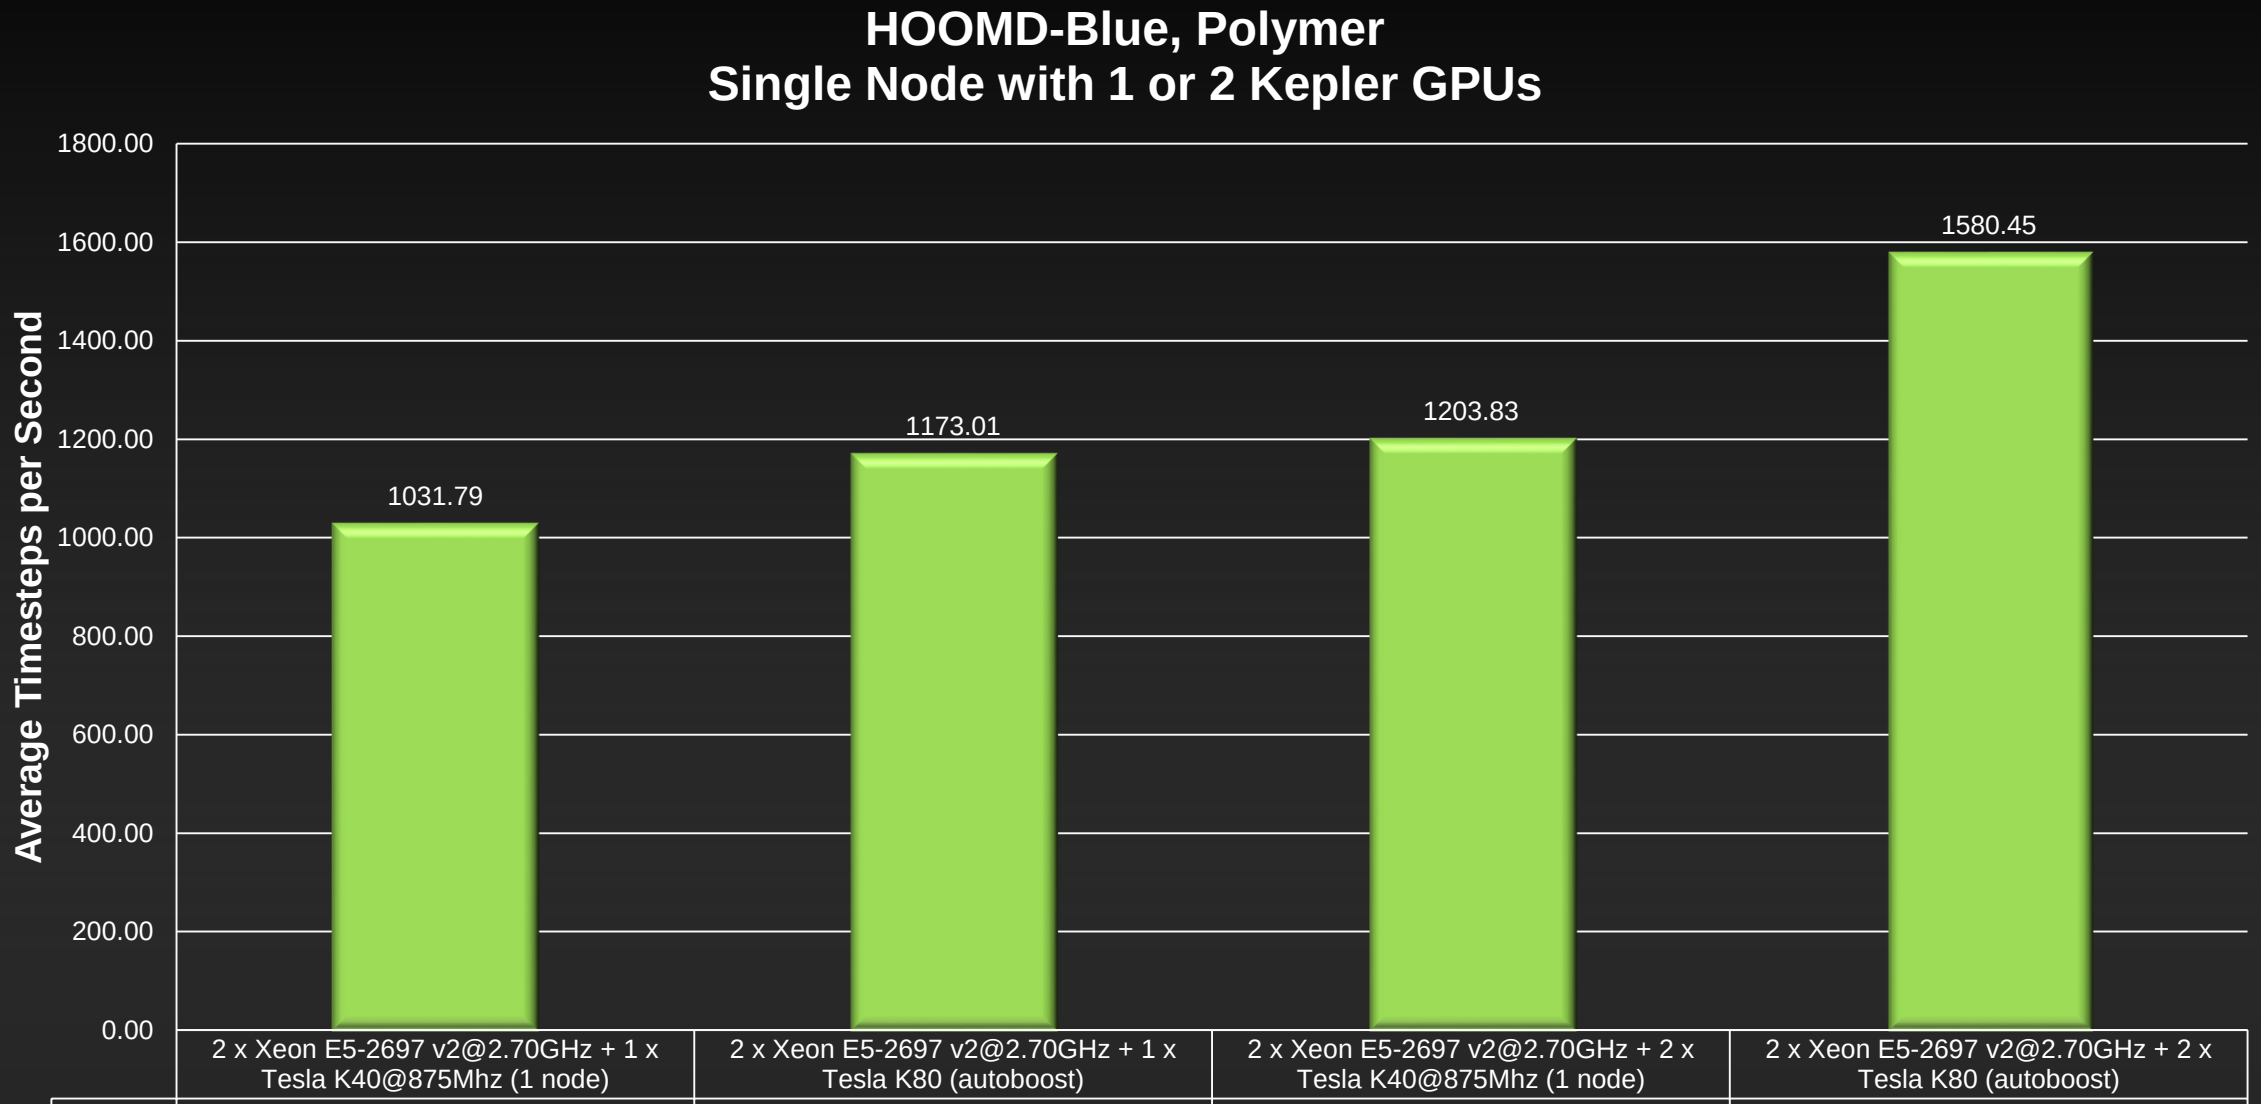

HOOMD-Blue 1.0.0 and K40, Boost impact!

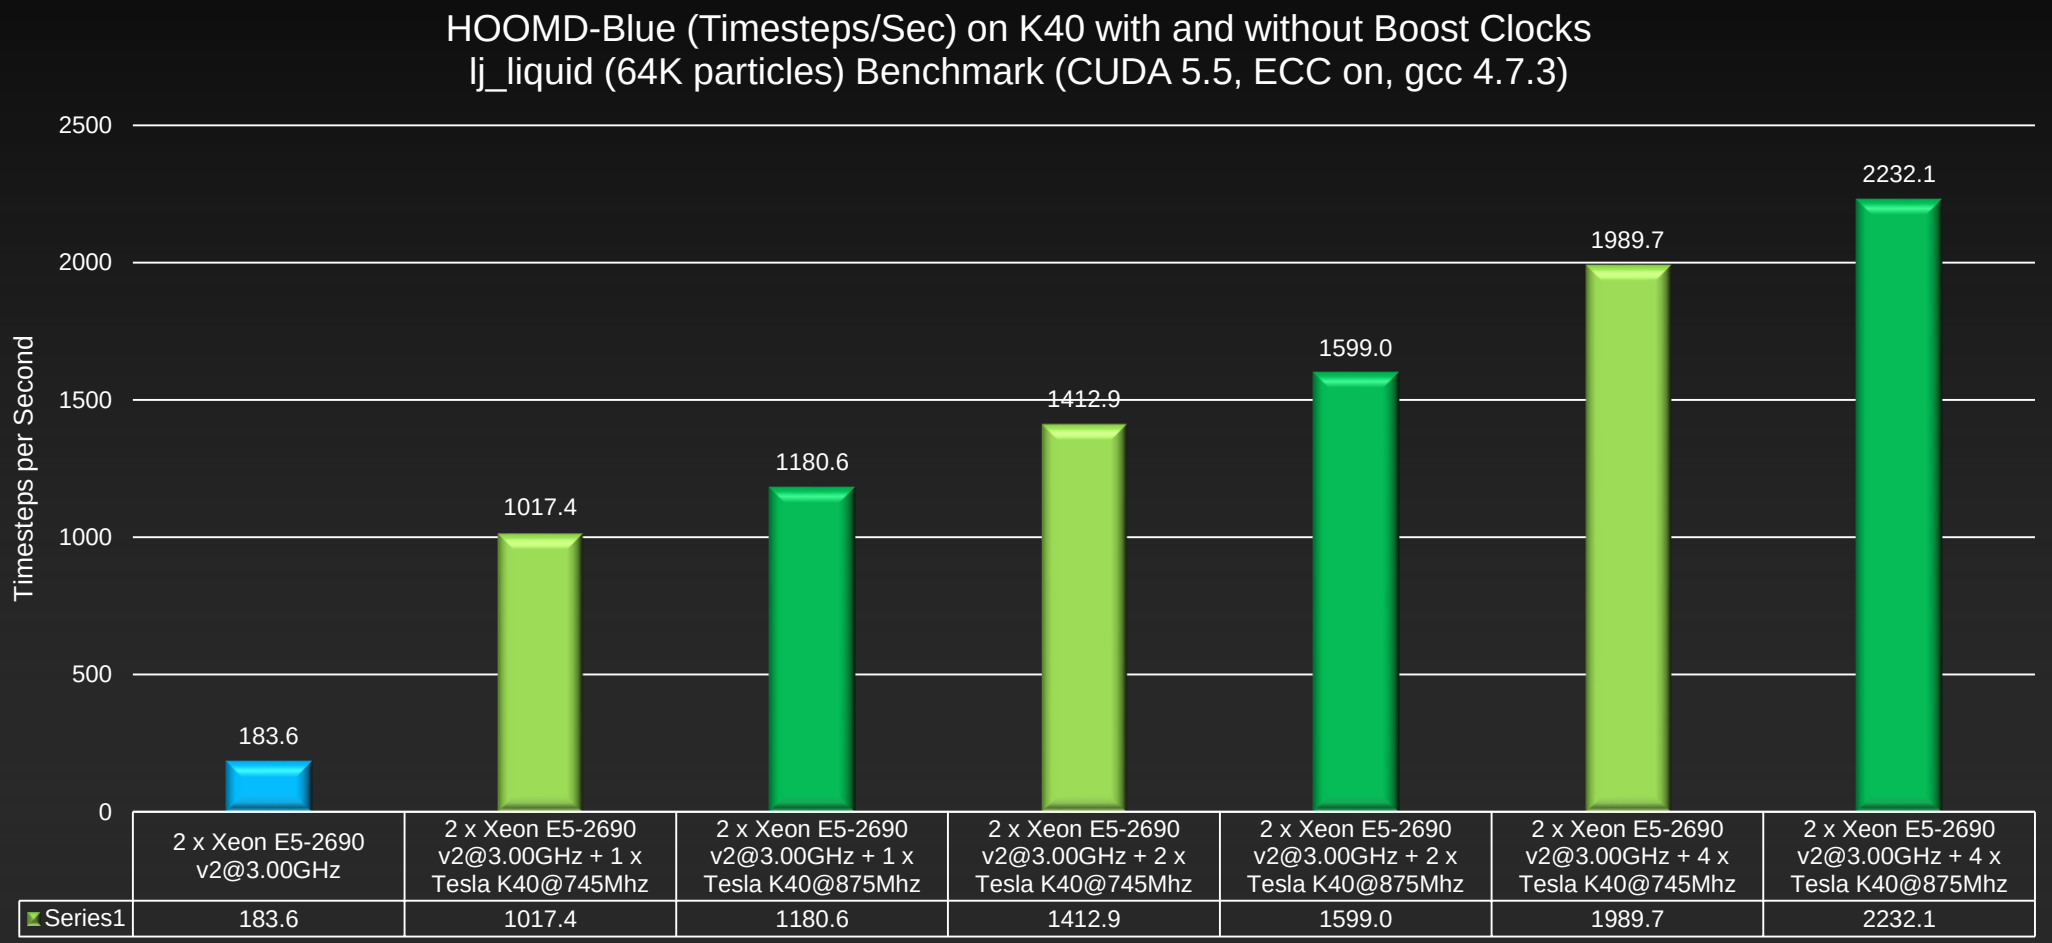

HOOMD-Blue 1.0.0 and K40, fastest GPU yet!

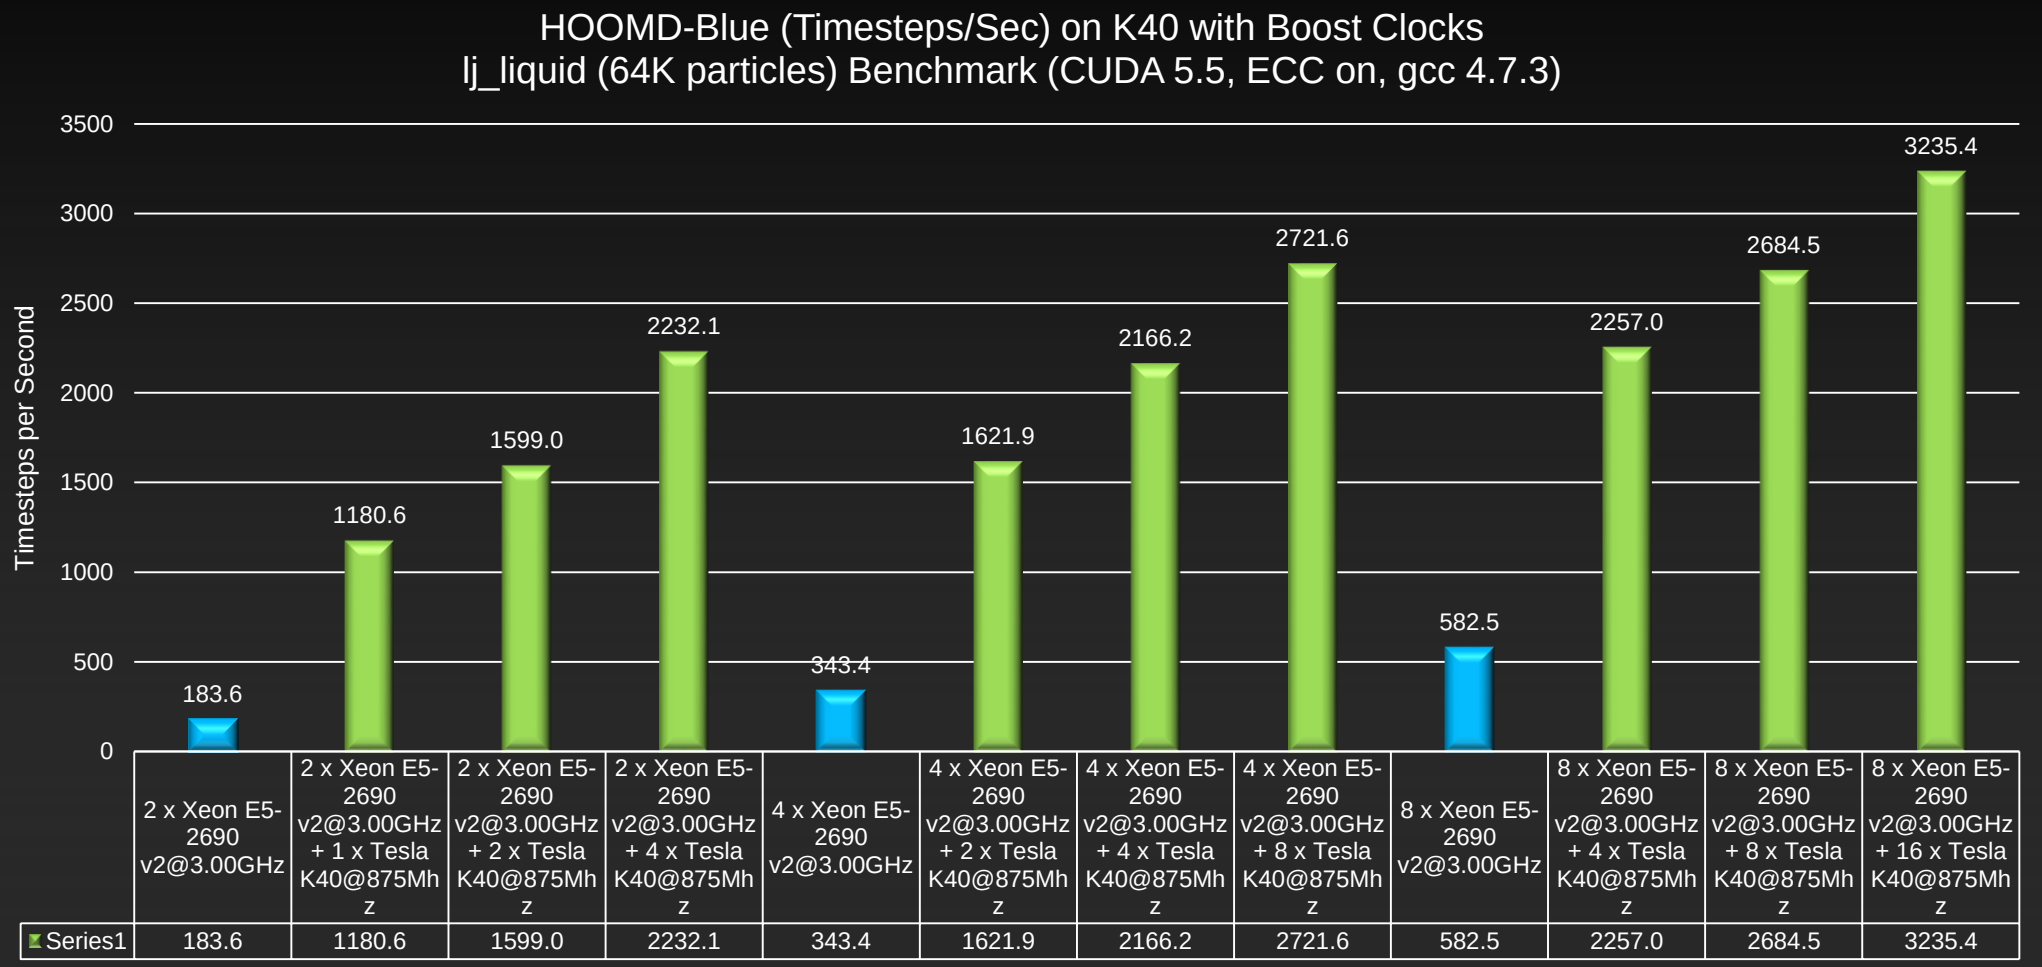

HOOMD-Blue 1.0.0 and K40, fastest GPU yet!

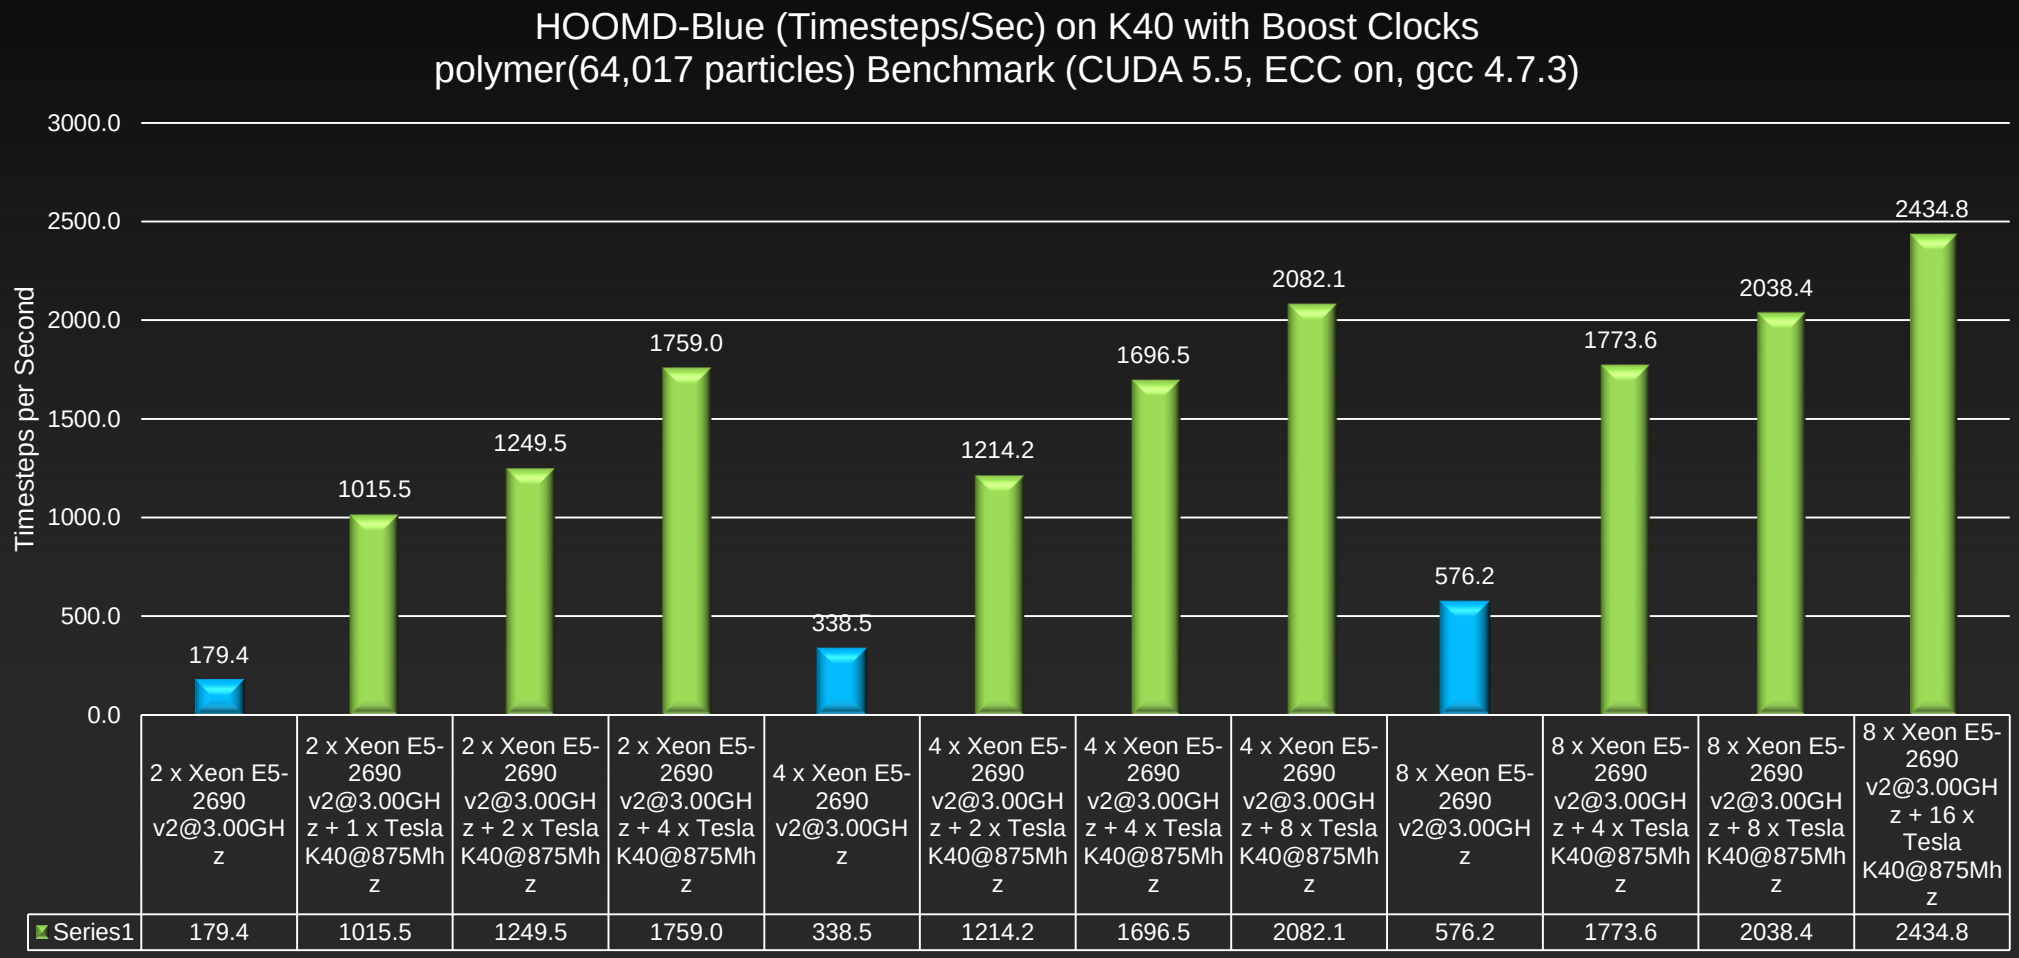

HOOMD-Blue 1.0.0 and K40, fastest GPU yet!

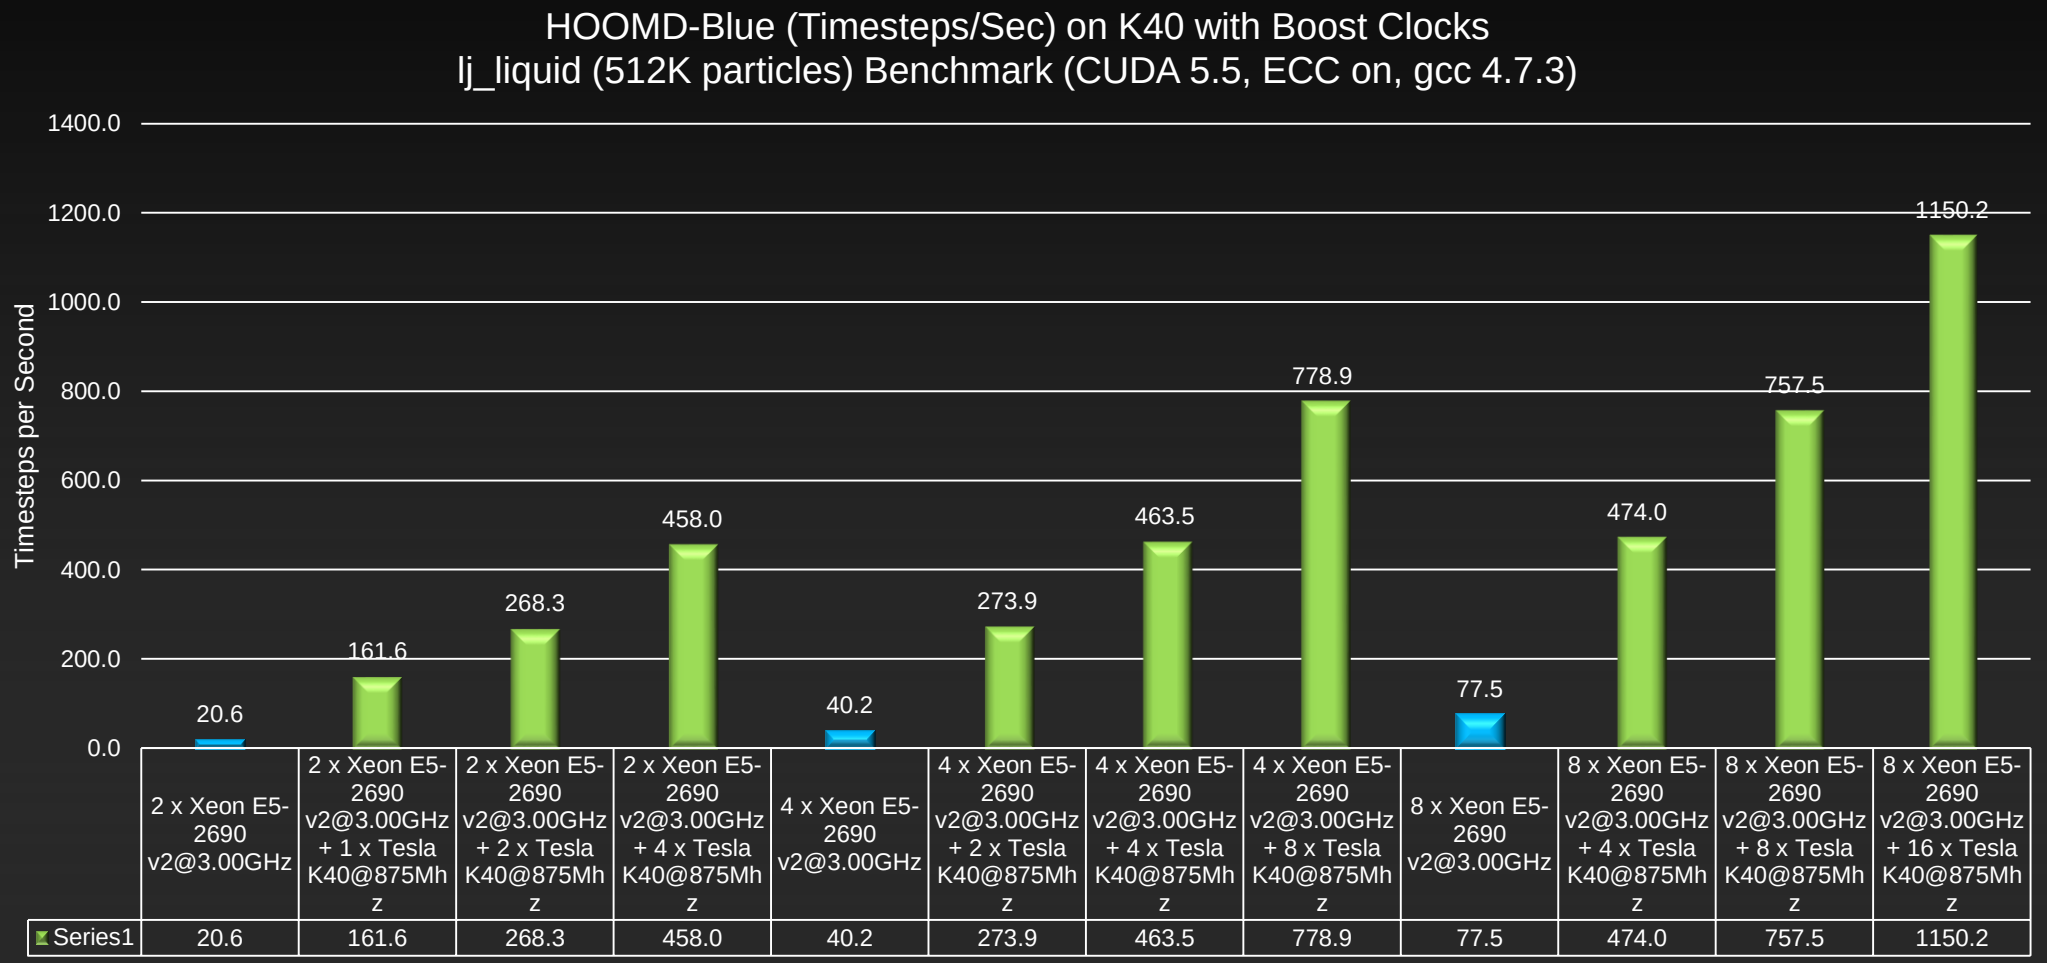

HOOMD-Blue on ARM vs. Ivy Bridge w/ & w/o K20
Equivalent Performance on ARM + K20

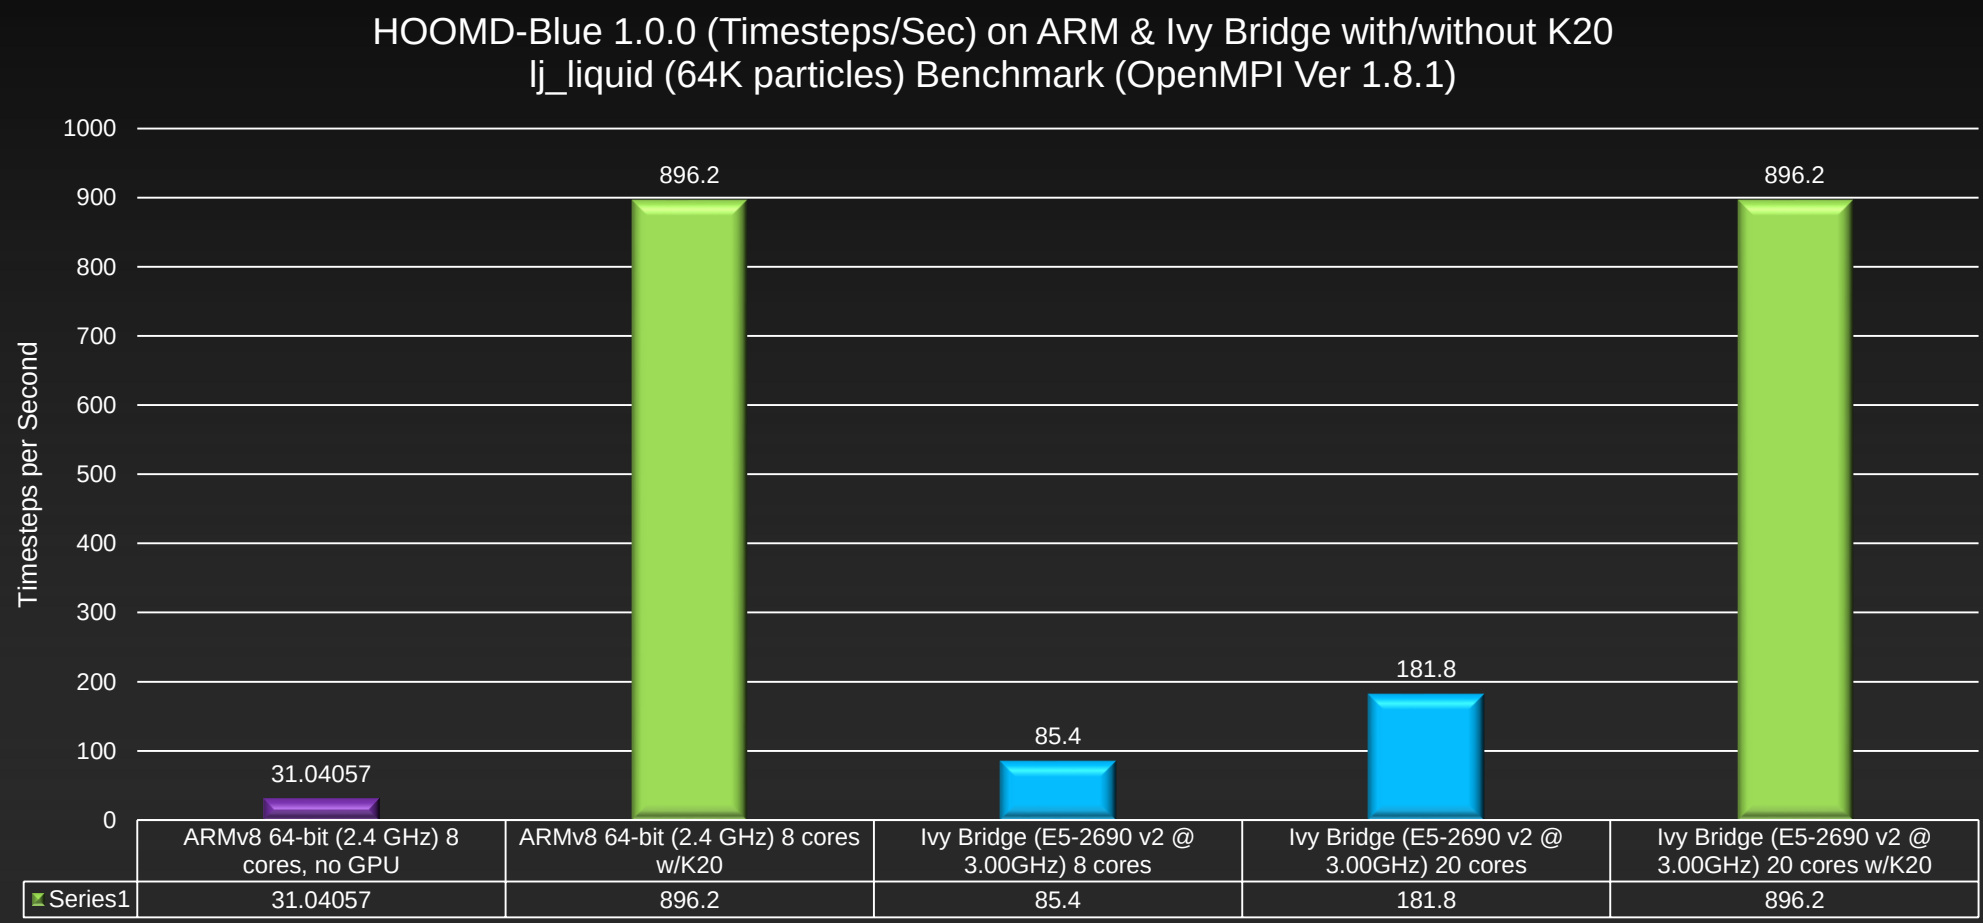

Application-Level Evaluation (HOOMD-blue)

HOOMD-blue Strong Scaling

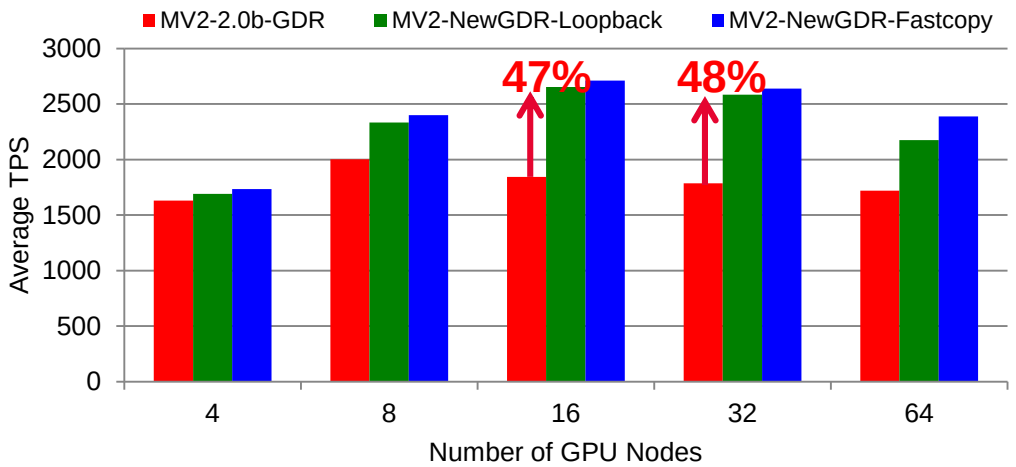

HOOMD-blue Weak Scaling

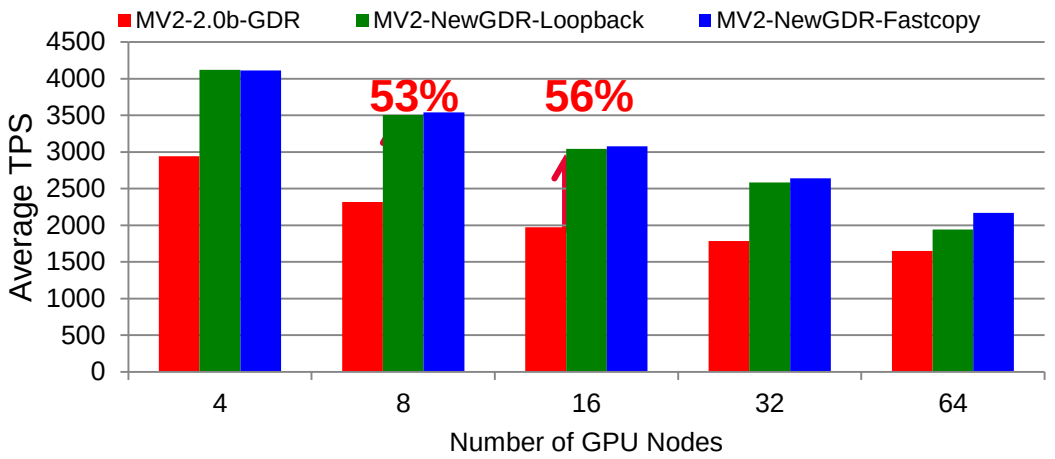

- Platform: Wilkes (Intel Ivy Bridge + NVIDIA Tesla K20c + Mellanox Connect-IB)
  - Strong Scaling: fixed 64K particles
  - Loopback and Fastcopy get up to 45% and 48% improvement for 32 GPUs
  - Weak Sailing: fixed 2K particles / GPU
  - Loopback and Fastcopy get up to 54% and 56% improvement for 16 GPUs

LAMMPS

October 2015

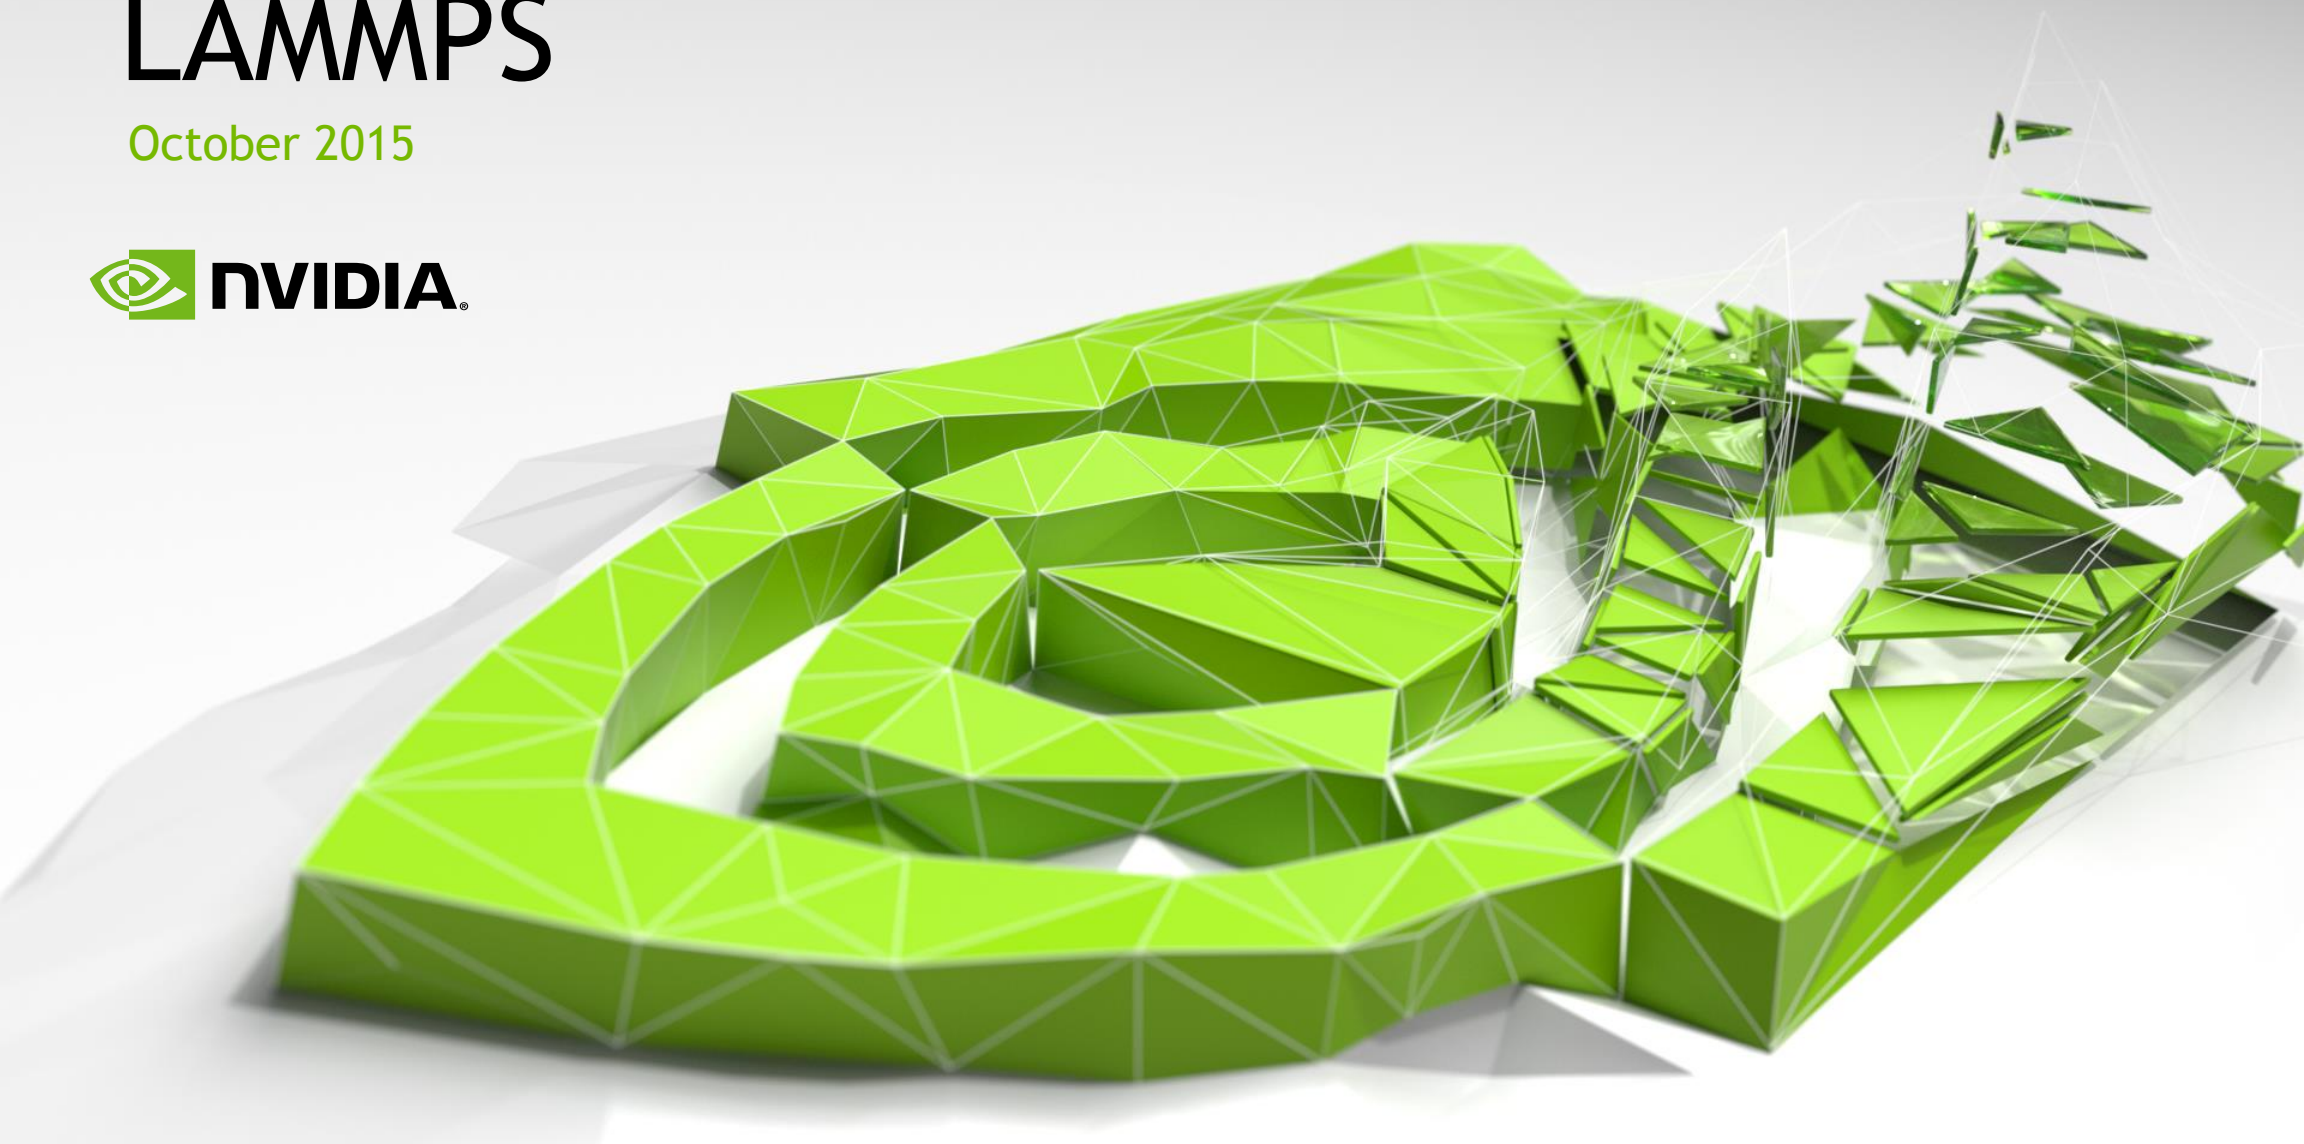

Lennard-Jones on K20X, K40s & K80s

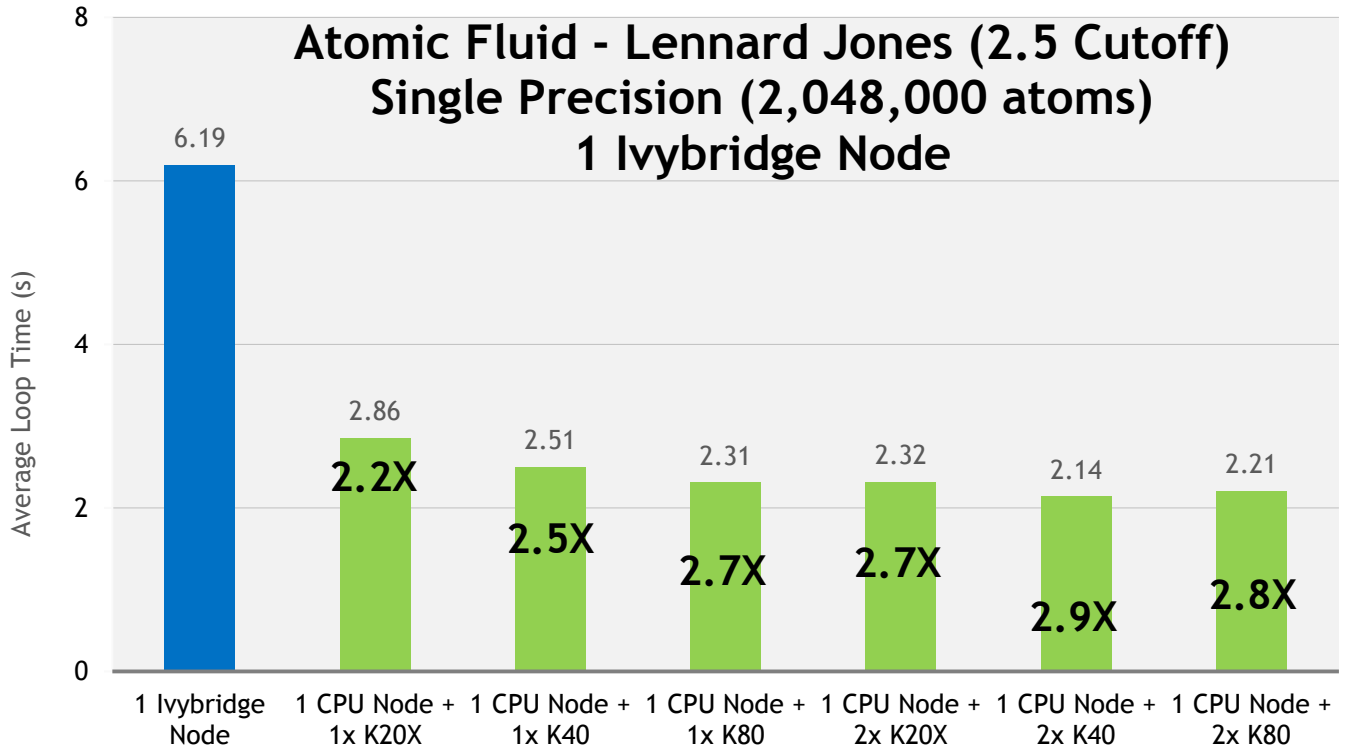

Running LAMMPS

The blue node contains Dual Intel Xeon E5-2697 v2@2.7GHz CPUs

The green nodes contain Dual Intel Xeon E5-2697 v2@2.7GHz CPUs + either NVIDIA Tesla K20X@732Mhz, Tesla K40@875Mhz or Tesla K80 (autoboost) GPUs

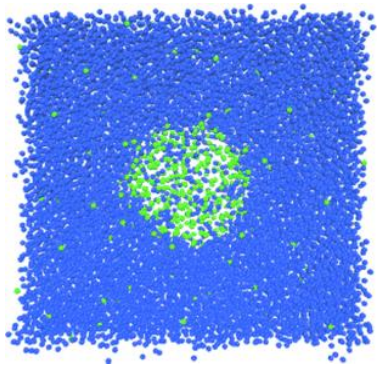

Lennard-Jones on K20X, K40s & K80s

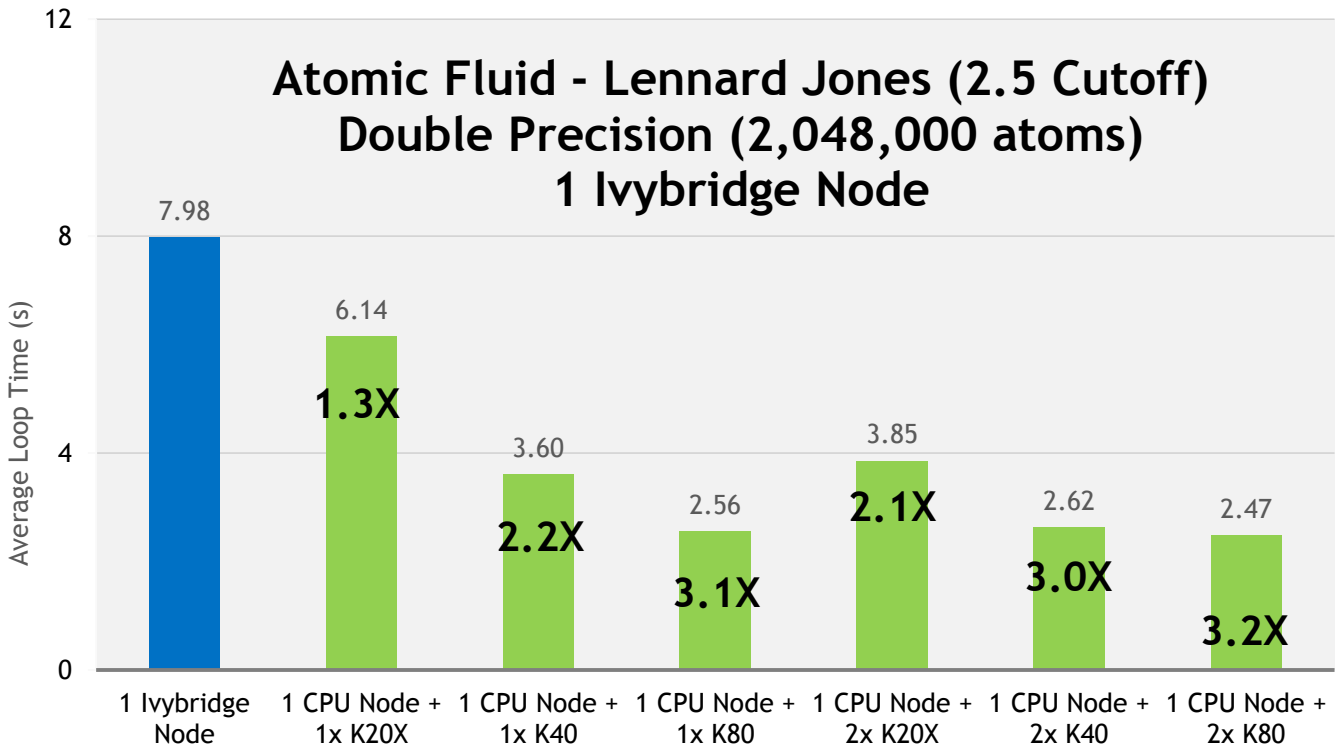

Running LAMMPS

The blue node contains Dual Intel Xeon E5-2697 v2@2.7GHz CPUs

The green nodes contain Dual Intel Xeon E5-2697 v2@2.7GHz CPUs + either NVIDIA Tesla K20X@732Mhz, Tesla K40@875Mhz or Tesla K80 (autoboost) GPUs

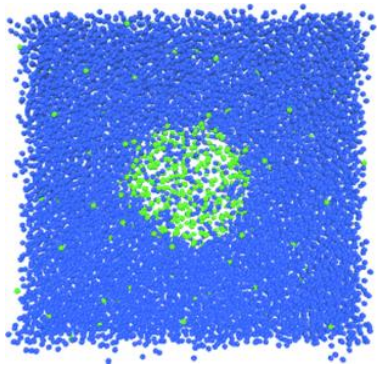

Lennard-Jones on K20X, K40s & K80s

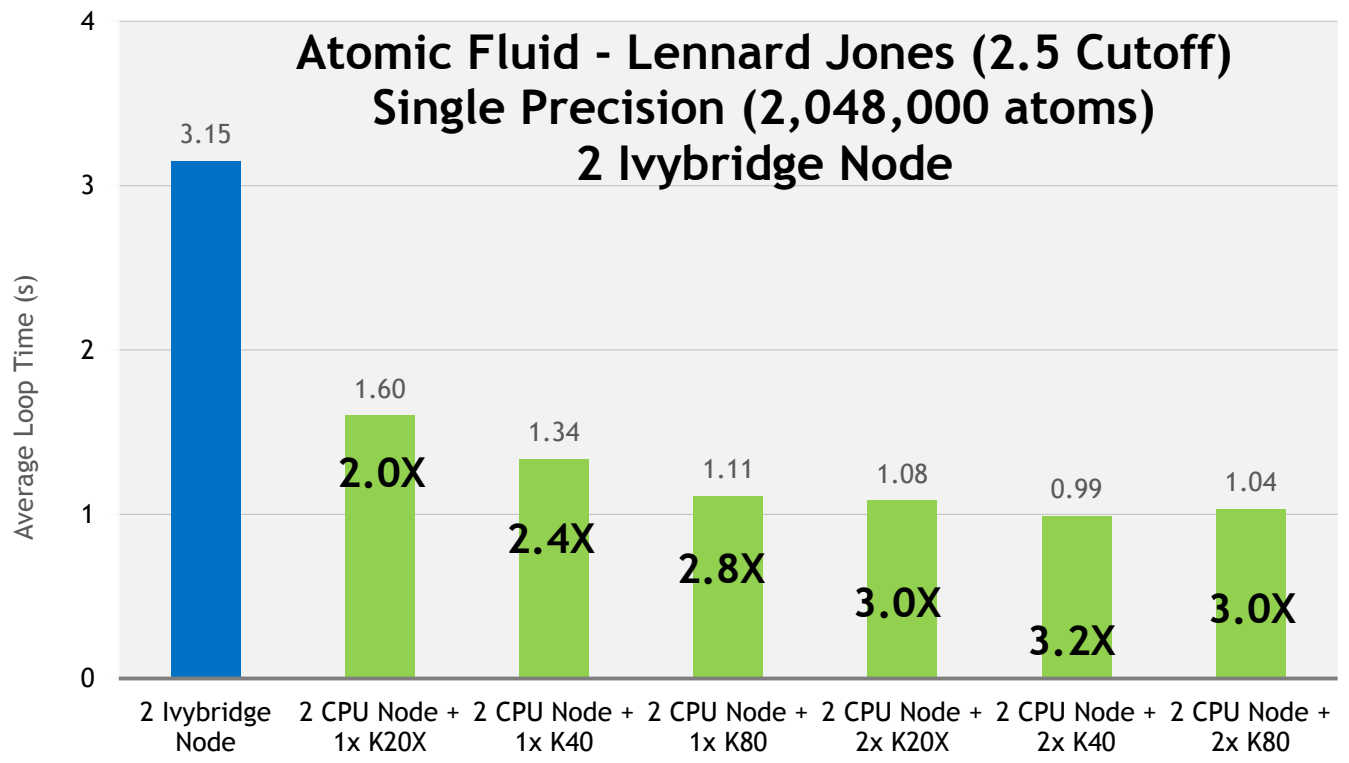

Running LAMMPS

The blue node contains Dual Intel Xeon E5-2697 v2@2.7GHz CPUs

The green nodes contain Dual Intel Xeon E5-2697 v2@2.7GHz CPUs + either NVIDIA Tesla K20X@732Mhz, Tesla K40@875Mhz or Tesla K80 (autoboost) GPUs

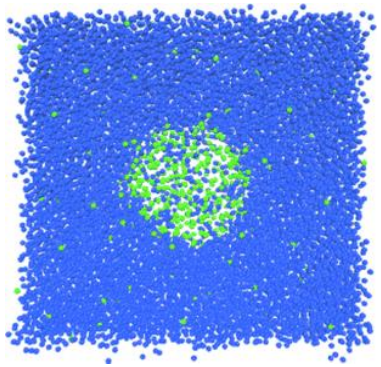

Lennard-Jones on K20X, K40s & K80s

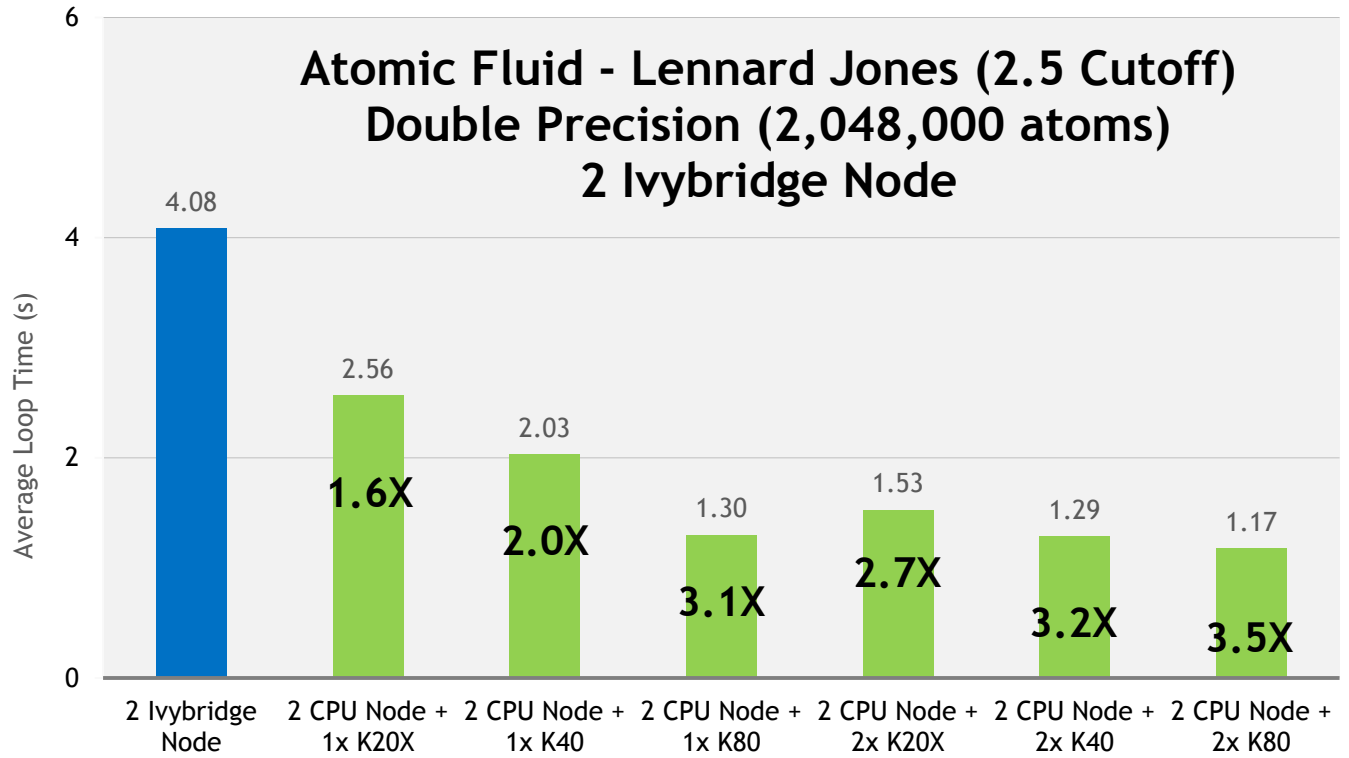

Running LAMMPS

The blue node contains Dual Intel Xeon E5-2697 v2@2.7GHz CPUs

The green nodes contain Dual Intel Xeon E5-2697 v2@2.7GHz CPUs + either NVIDIA Tesla K20X@732Mhz, Tesla K40@875Mhz or Tesla K80 (autoboost) GPUs

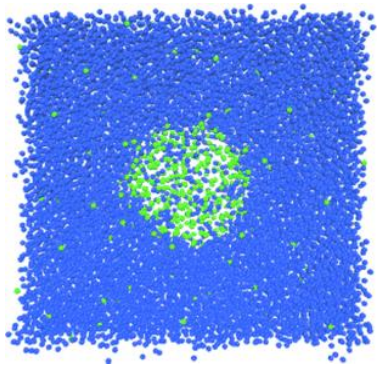

Lennard-Jones on K20X, K40s & K80s

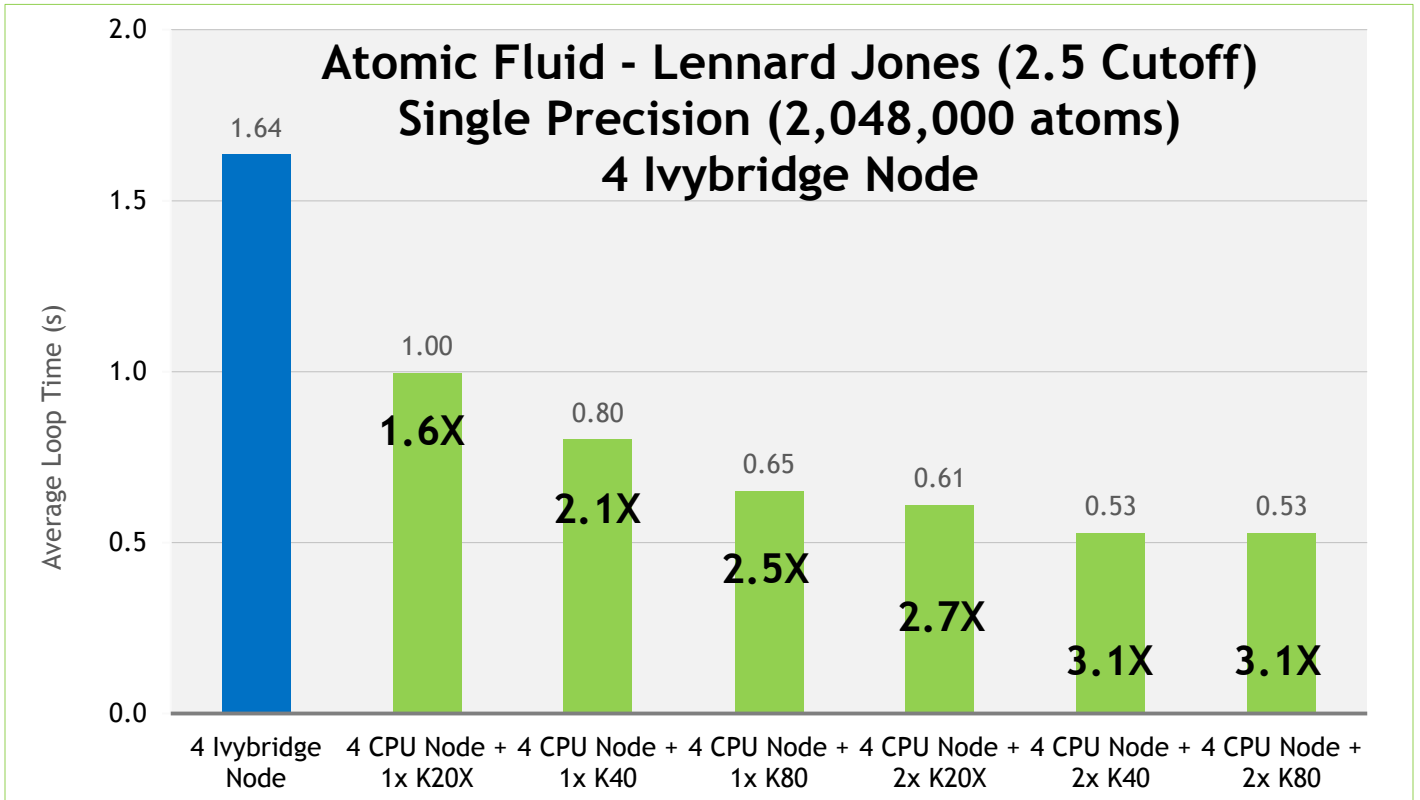

Running LAMMPS

The blue node contains Dual Intel Xeon E5-2697 v2@2.7GHz CPUs

The green nodes contain Dual Intel Xeon E5-2697 v2@2.7GHz CPUs + either NVIDIA Tesla K20X@732Mhz, Tesla K40@875Mhz or Tesla K80 (autoboost) GPUs

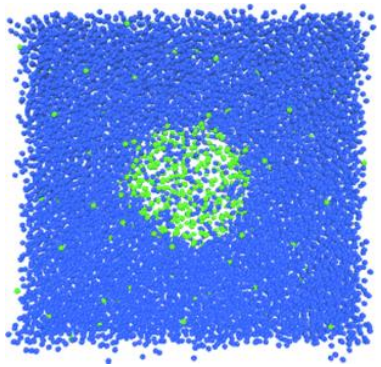

Lennard-Jones on K20X, K40s & K80s

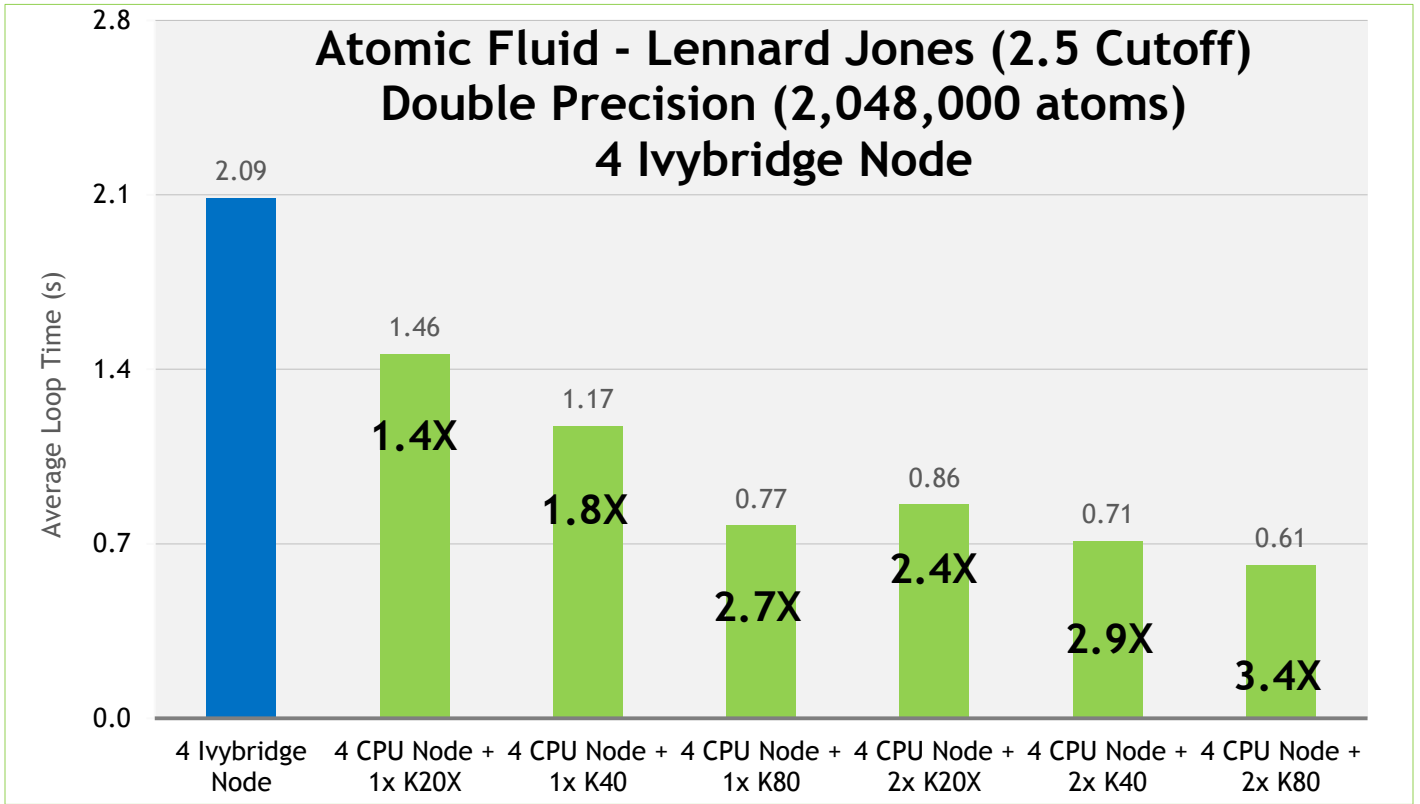

Running LAMMPS

The blue node contains Dual Intel Xeon E5-2697 v2@2.7GHz CPUs

The green nodes contain Dual Intel Xeon E5-2697 v2@2.7GHz CPUs + either NVIDIA Tesla K20X@732Mhz, Tesla K40@875Mhz or Tesla K80 (autoboost) GPUs

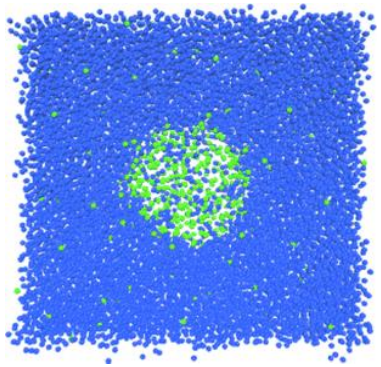

Lennard-Jones on K20X, K40s & K80s

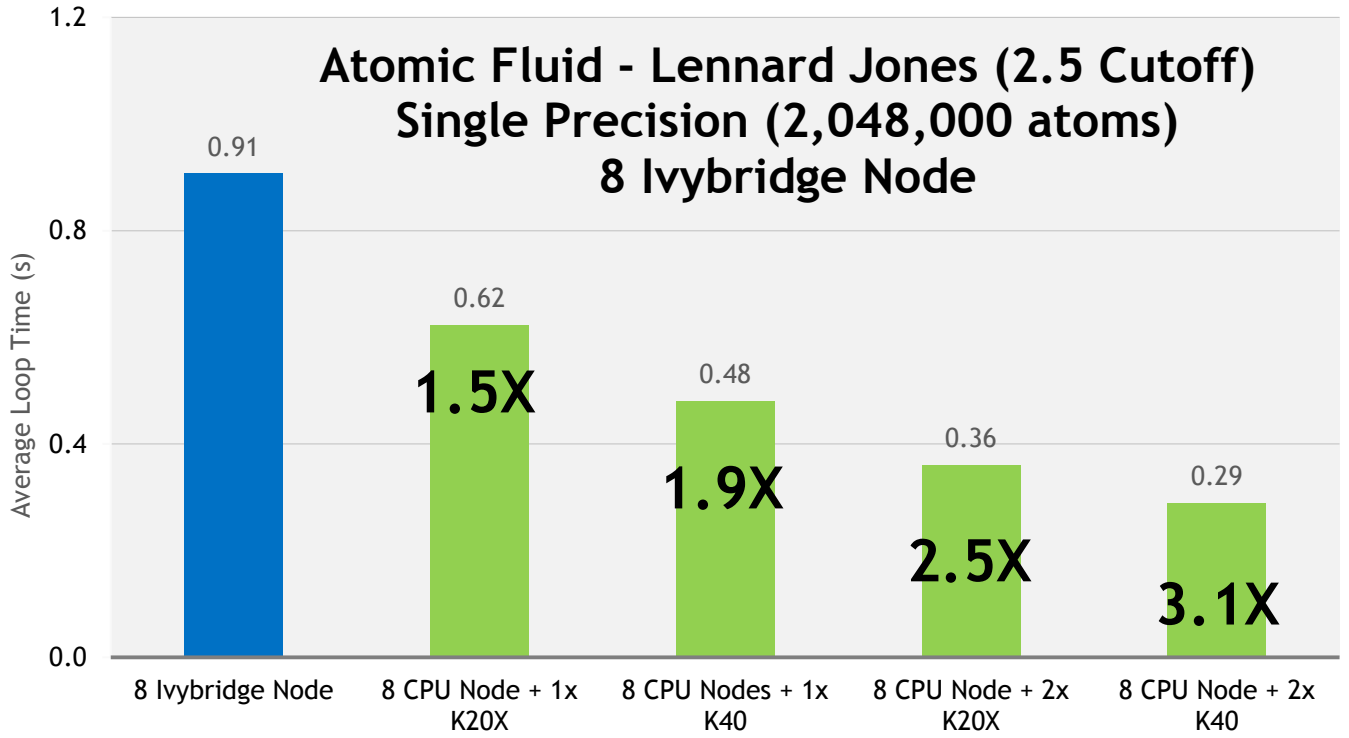

Running LAMMPS

The blue node contains Dual Intel Xeon E5-2697 v2@2.7GHz CPUs

The green nodes contain Dual Intel Xeon E5-2697 v2@2.7GHz CPUs + either NVIDIA Tesla K20X@732Mhz, Tesla K40@875Mhz or Tesla K80 (autoboost) GPUs

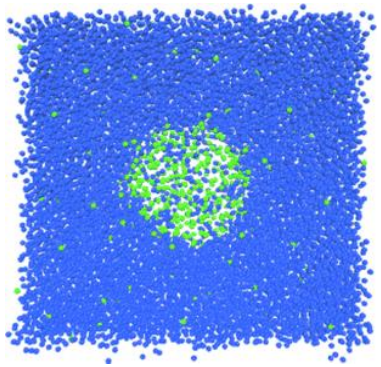

Lennard-Jones on K20X, K40s & K80s

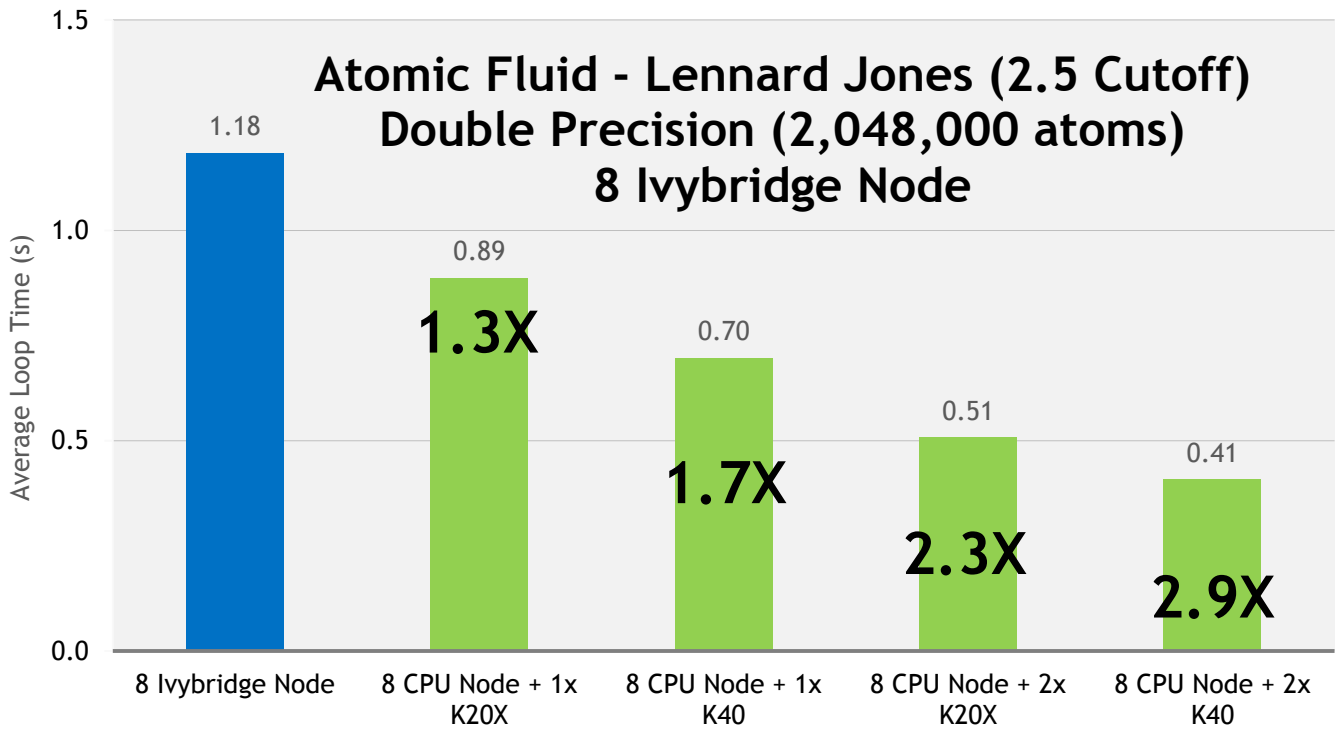

Running LAMMPS

The blue node contains Dual Intel Xeon E5-2697 v2@2.7GHz CPUs

The green nodes contain Dual Intel Xeon E5-2697 v2@2.7GHz CPUs + either NVIDIA Tesla K20X@732Mhz, Tesla K40@875Mhz or Tesla K80 (autoboost) GPUs

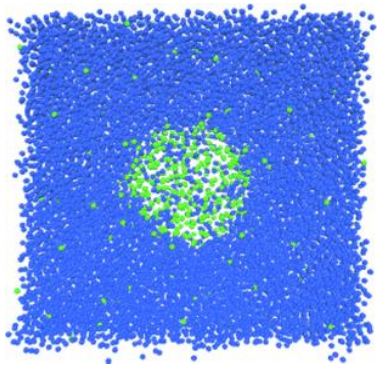

Lennard-Jones on K20X, K40s & K80s

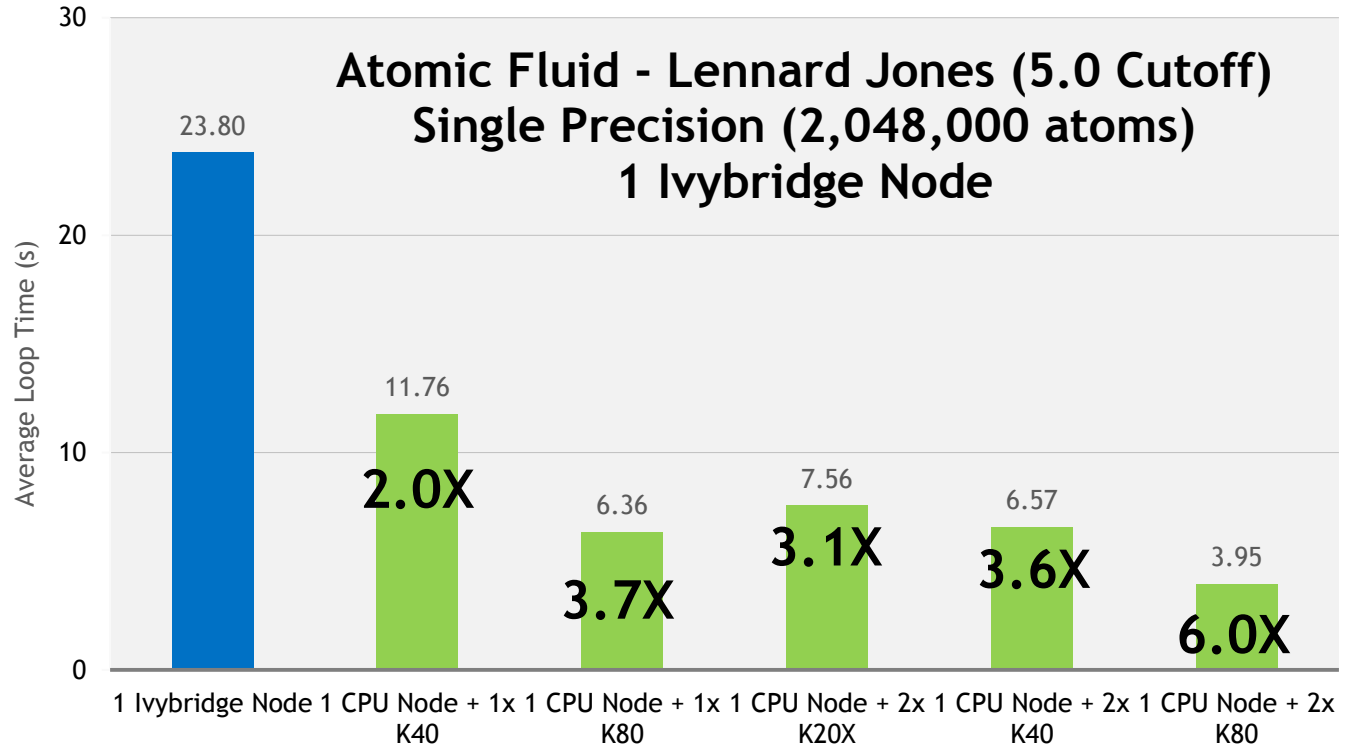

Running LAMMPS

The blue node contains Dual Intel Xeon E5-2697 v2@2.7GHz CPUs

The green nodes contain Dual Intel Xeon E5-2697 v2@2.7GHz CPUs + either NVIDIA Tesla K20X@732Mhz, Tesla K40@875Mhz or Tesla K80 (autoboost) GPUs

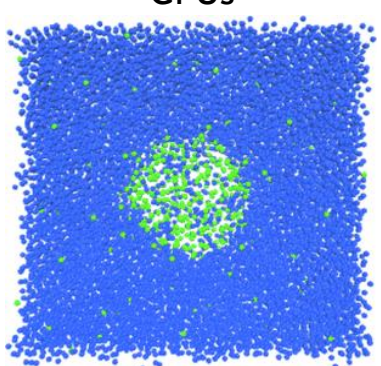

Lennard-Jones on K20X, K40s & K80s

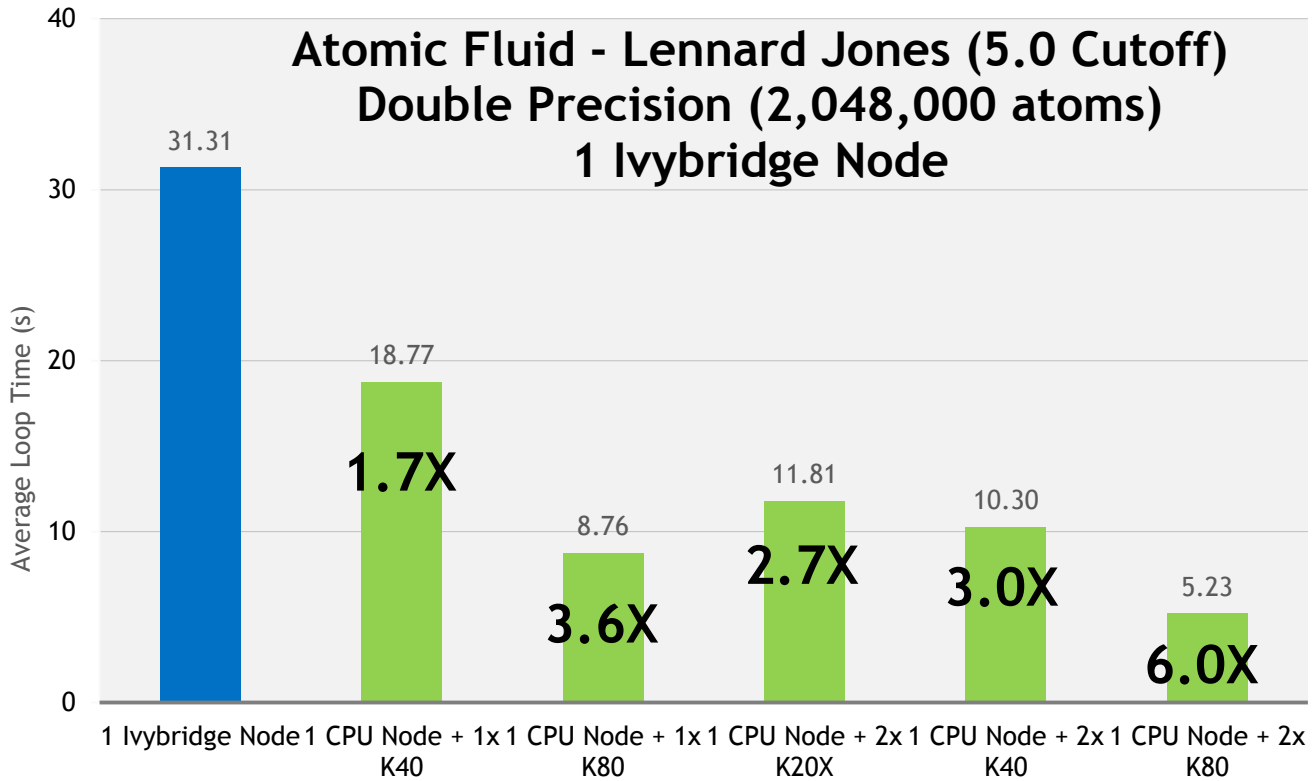

Running LAMMPS

The blue node contains Dual Intel Xeon E5-2697 v2@2.7GHz CPUs

The green nodes contain Dual Intel Xeon E5-2697 v2@2.7GHz CPUs + either NVIDIA Tesla K20X@732Mhz, Tesla K40@875Mhz or Tesla K80 (autoboost) GPUs

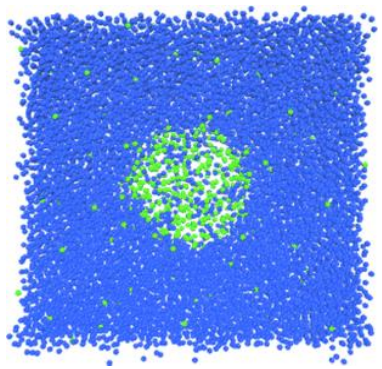

Lennard-Jones on K20X, K40s & K80s

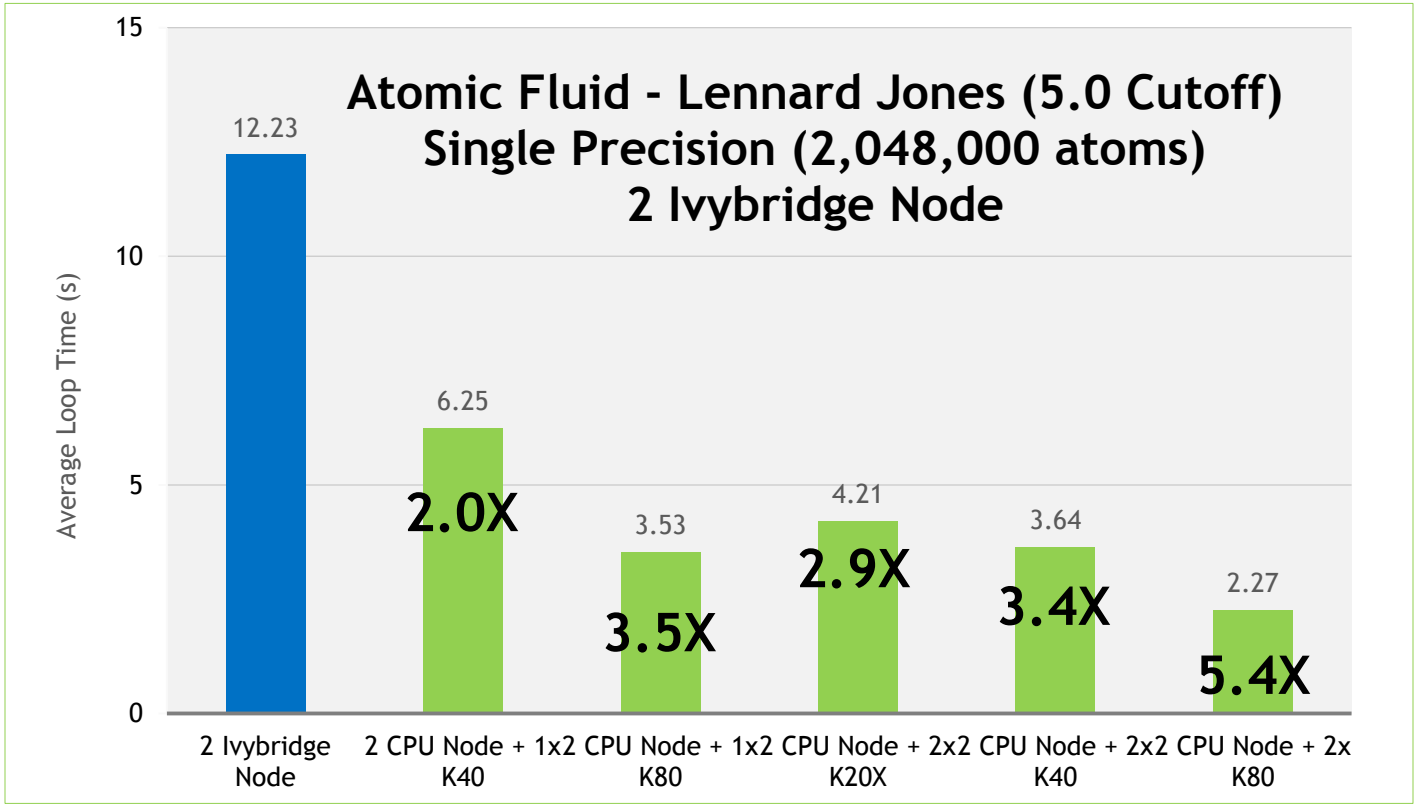

Running LAMMPS

The blue node contains Dual Intel Xeon E5-2697 v2@2.7GHz CPUs

The green nodes contain Dual Intel Xeon E5-2697 v2@2.7GHz CPUs + either NVIDIA Tesla K20X@732Mhz, Tesla K40@875Mhz or Tesla K80 (autoboost) GPUs

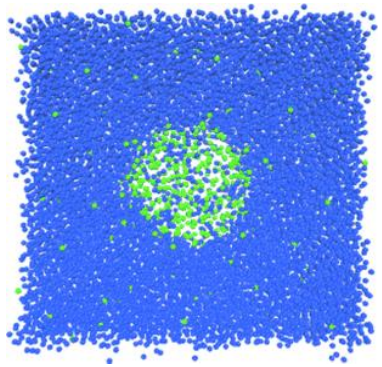

Lennard-Jones on K20X, K40s & K80s

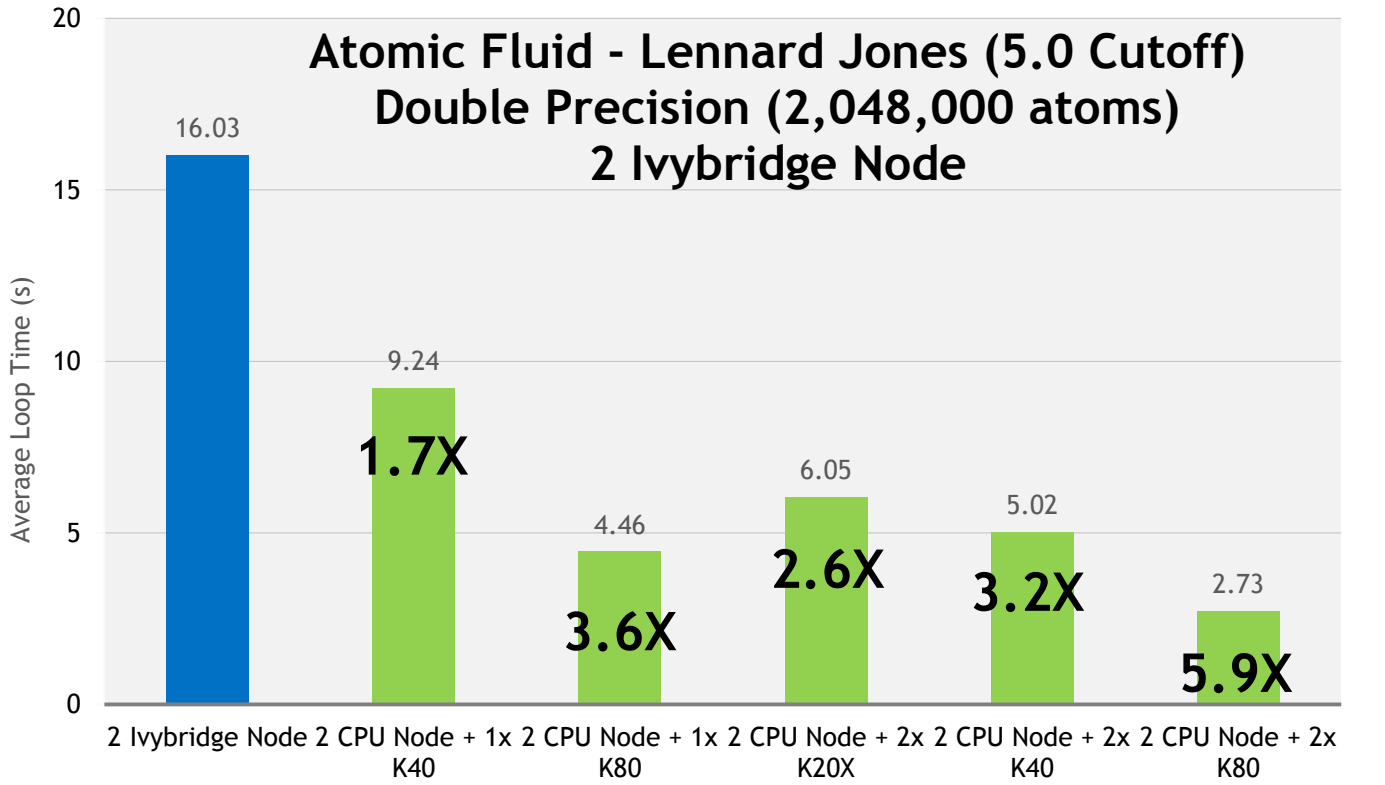

Running LAMMPS

The blue node contains Dual Intel Xeon E5-2697 v2@2.7GHz CPUs

The green nodes contain Dual Intel Xeon E5-2697 v2@2.7GHz CPUs + either NVIDIA Tesla K20X@732Mhz, Tesla K40@875Mhz or Tesla K80 (autoboost) GPUs

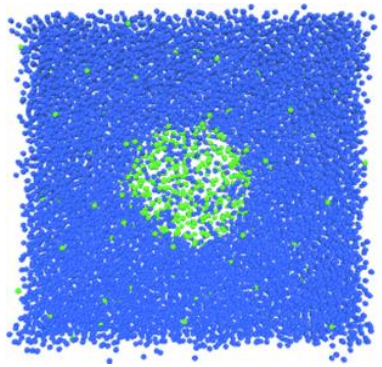

Lennard-Jones on K20X, K40s & K80s

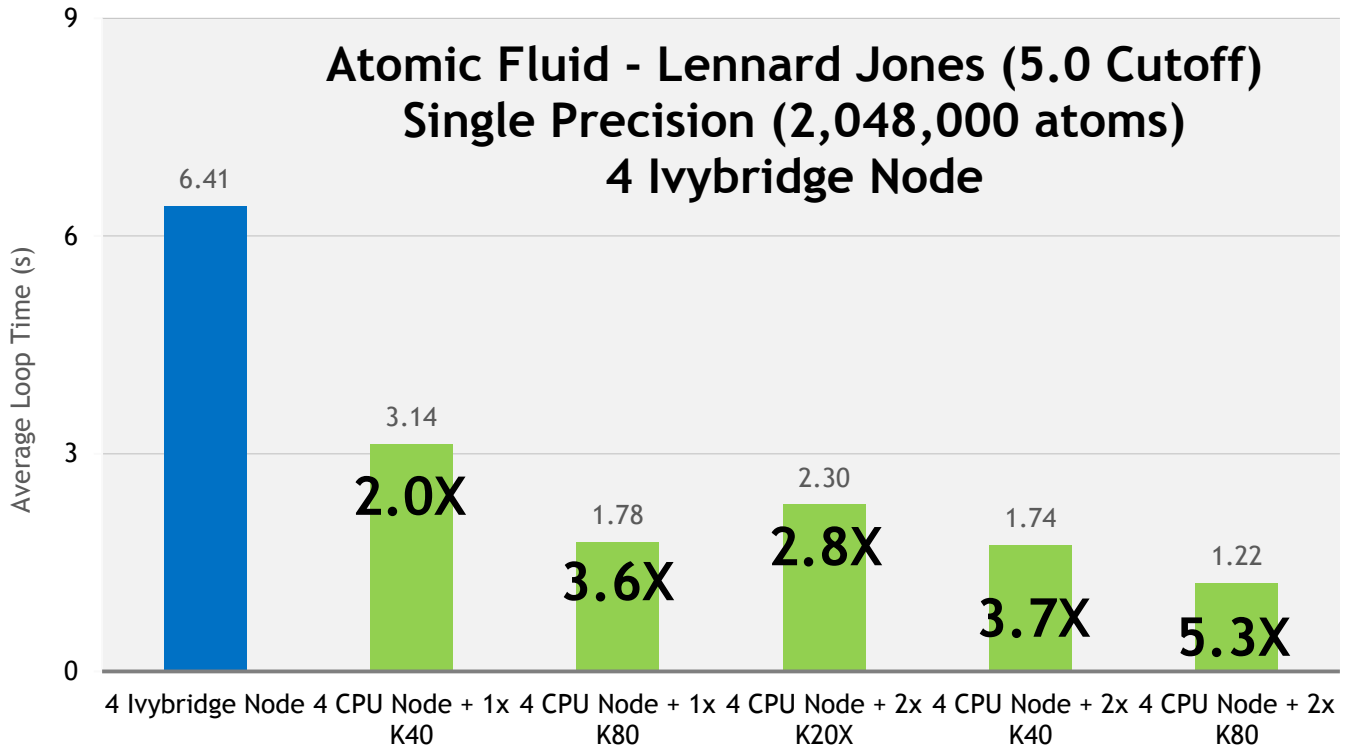

Running LAMMPS

The blue node contains Dual Intel Xeon E5-2697 v2@2.7GHz CPUs

The green nodes contain Dual Intel Xeon E5-2697 v2@2.7GHz CPUs + either NVIDIA Tesla K20X@732Mhz, Tesla K40@875Mhz or Tesla K80 (autoboost) GPUs

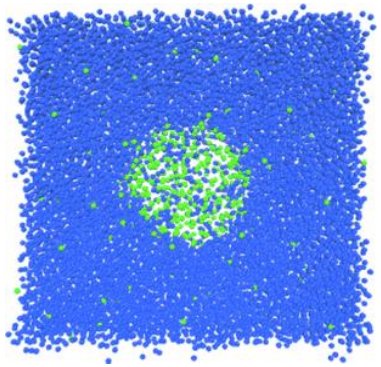

Lennard-Jones on K20X, K40s & K80s

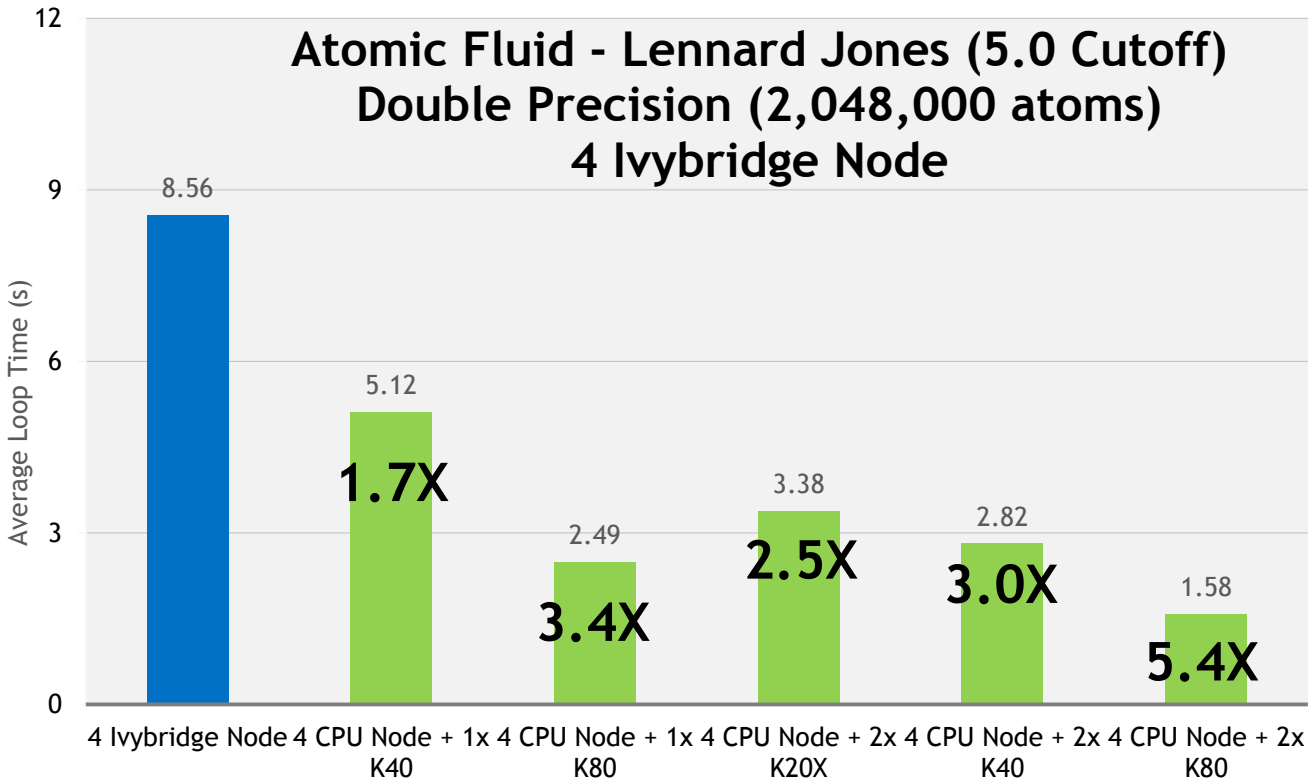

Running LAMMPS

The blue node contains Dual Intel Xeon E5-2697 v2@2.7GHz CPUs

The green nodes contain Dual Intel Xeon E5-2697 v2@2.7GHz CPUs + either NVIDIA Tesla K20X@732Mhz, Tesla K40@875Mhz or Tesla K80 (autoboost) GPUs

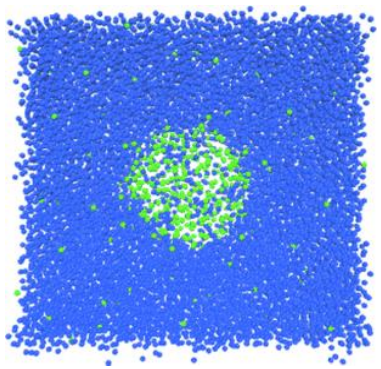

Lennard-Jones on K20X and K40s

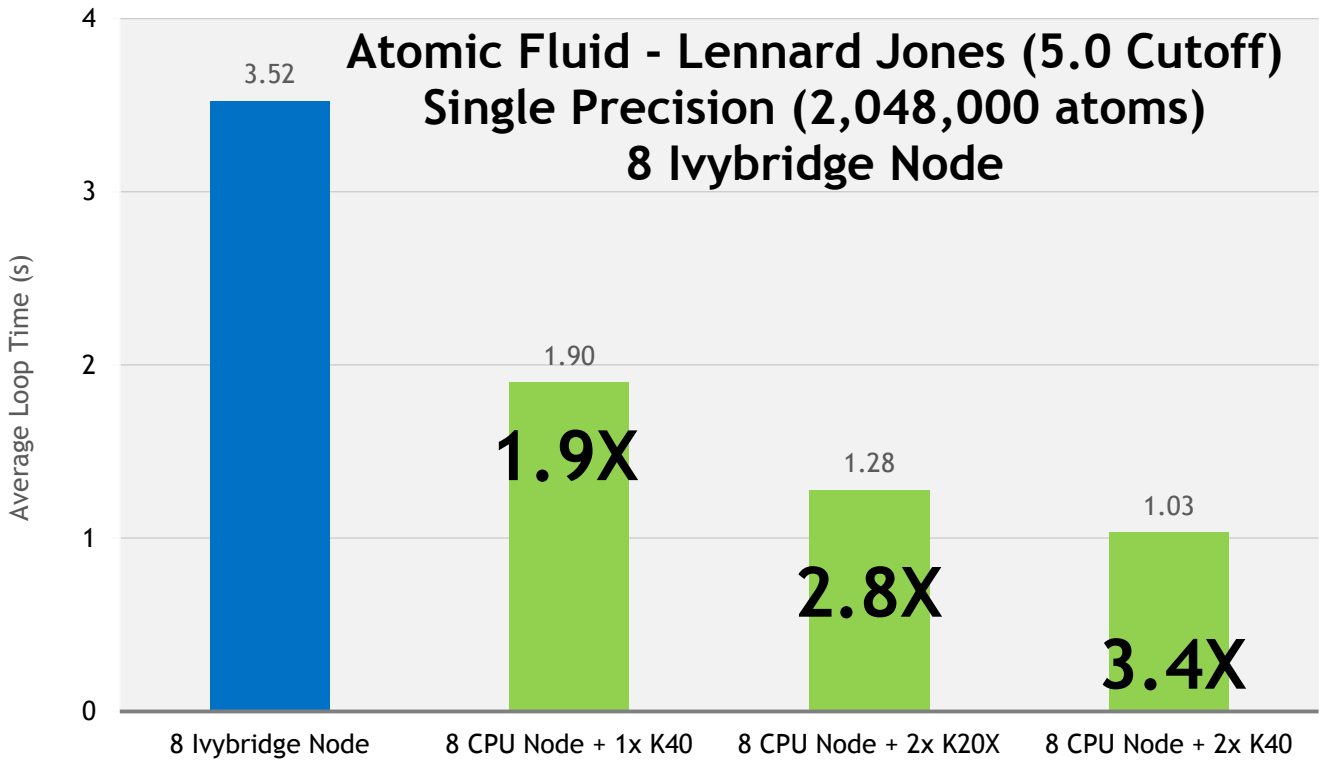

Running LAMMPS

The blue node contains Dual Intel Xeon E5-2697 v2@2.7GHz CPUs

The green nodes contain Dual Intel Xeon E5-2697 v2@2.7GHz CPUs + either NVIDIA Tesla K20X@732Mhz, Tesla K40@875Mhz or Tesla K80 (autoboost) GPUs

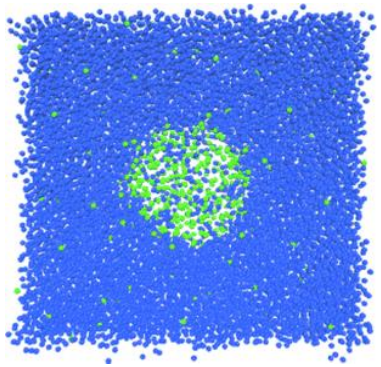

Lennard-Jones on K20X and K40s

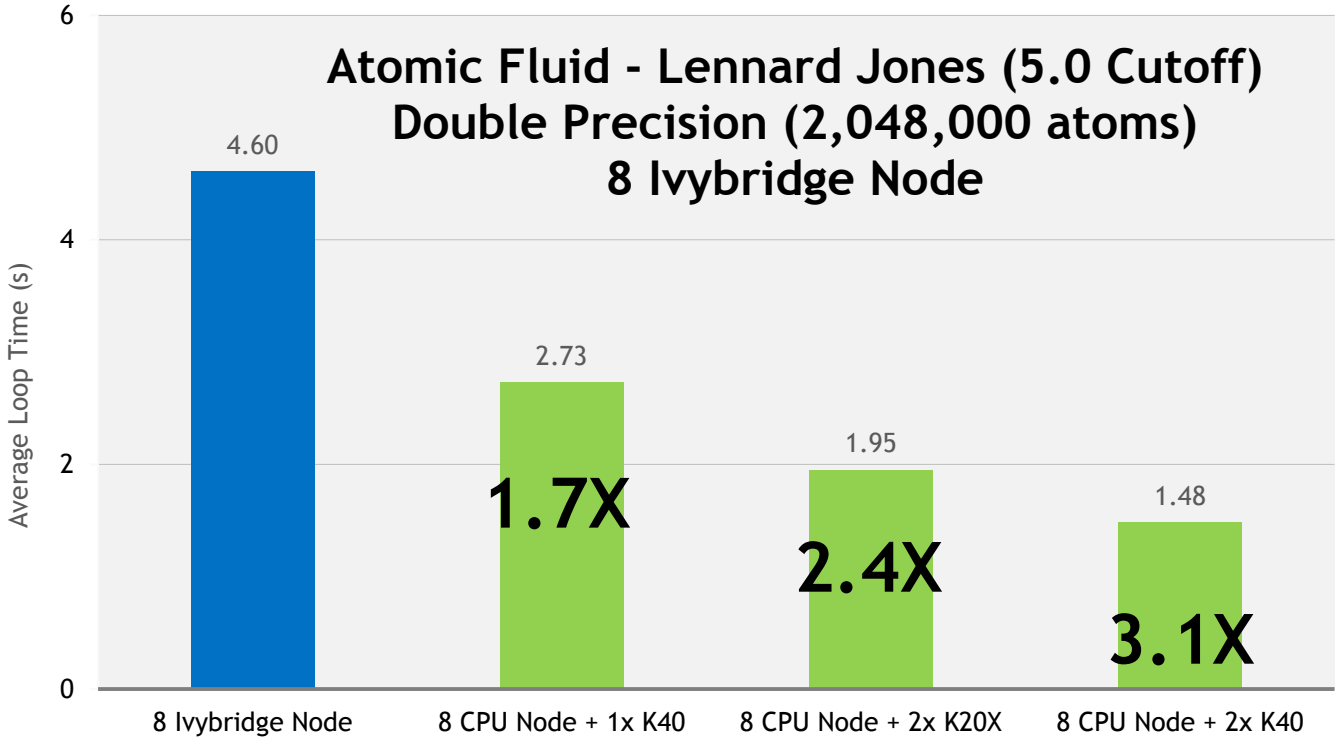

Running LAMMPS

The blue node contains Dual Intel Xeon E5-2697 v2@2.7GHz CPUs

The green nodes contain Dual Intel Xeon E5-2697 v2@2.7GHz CPUs + either NVIDIA Tesla K20X@732Mhz, Tesla K40@875Mhz or Tesla K80 (autoboost) GPUs

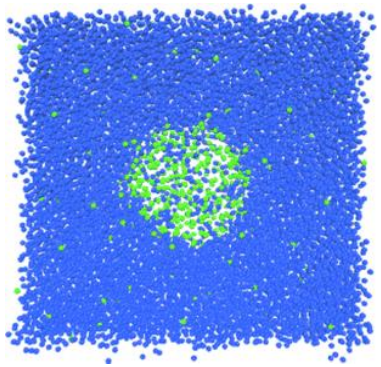

Lennard-Jones single/multi-node throughout

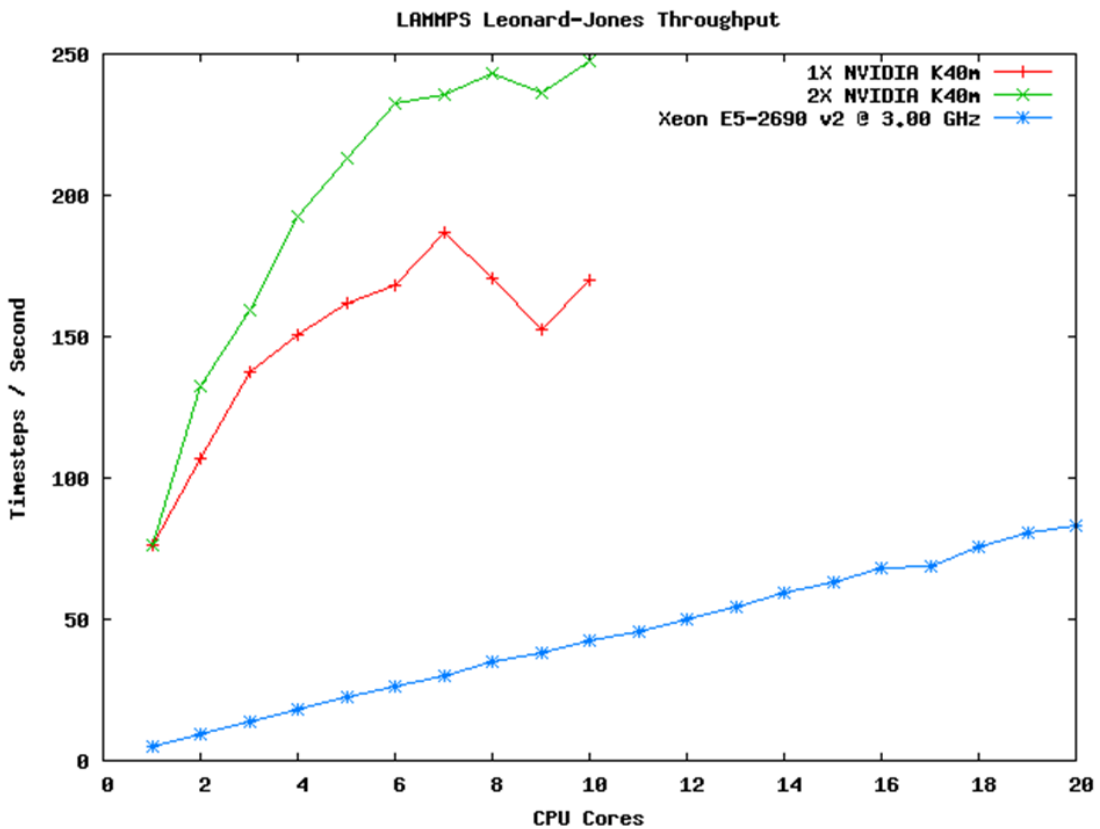

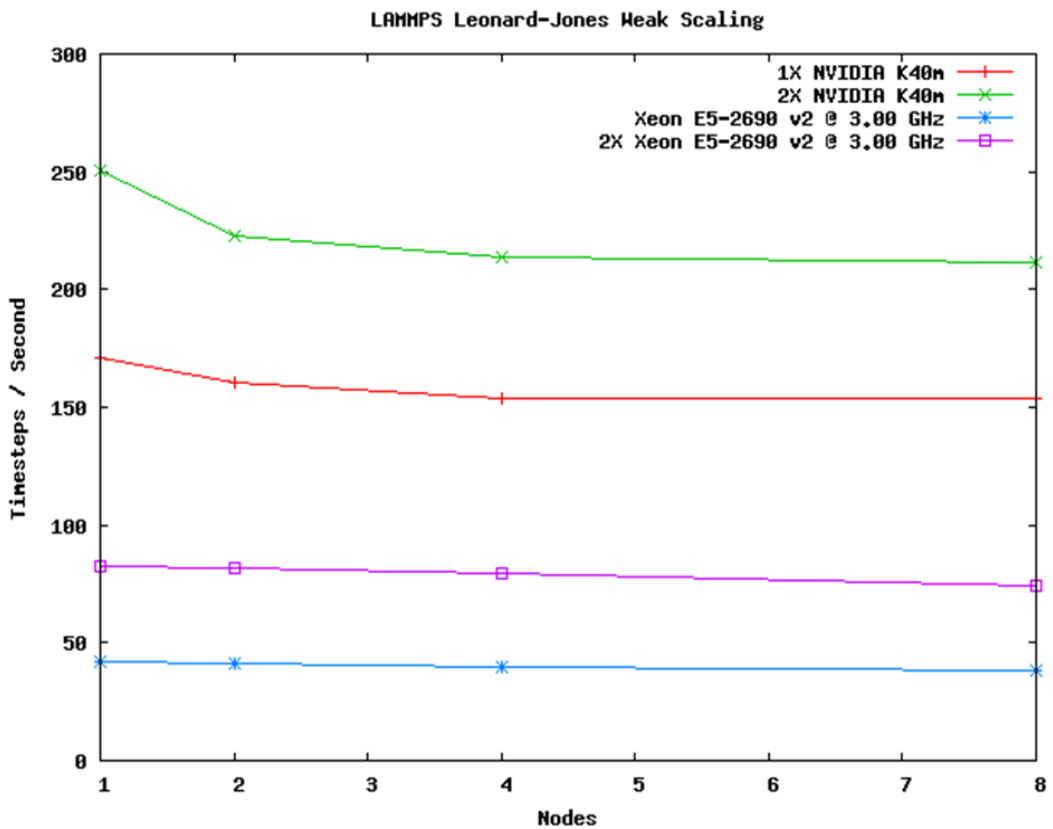

Figure 1: Leonard-Jones single-node throughout (strong scaling)
Figure 2: Leonard-Jones multi-node throughout

EAM on K20X, K40s & K80s

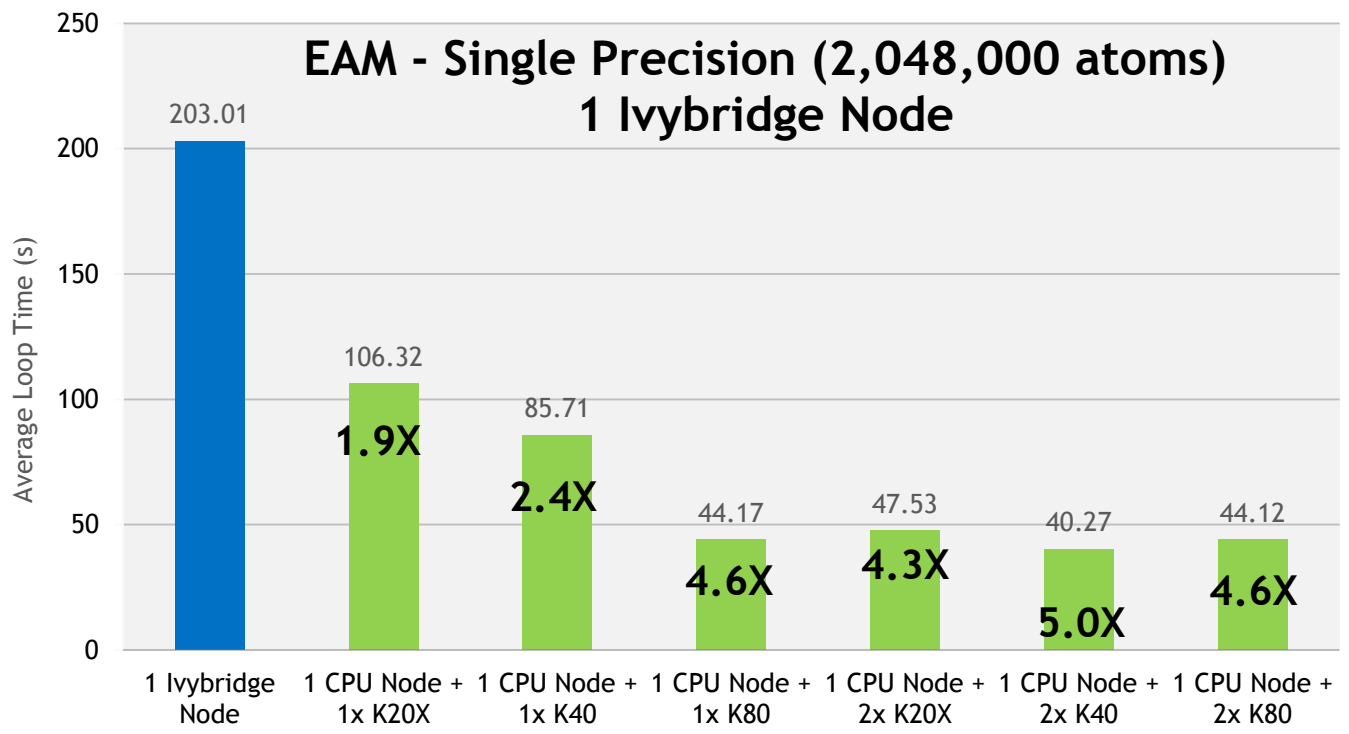

Running LAMMPS

The blue node contains Dual Intel Xeon E5-2697 v2@2.7GHz CPUs

The green nodes contain Dual Intel Xeon E5-2697 v2@2.7GHz CPUs + either NVIDIA Tesla K20X@732Mhz, Tesla K40@875Mhz or Tesla K80 (autoboost) GPUs

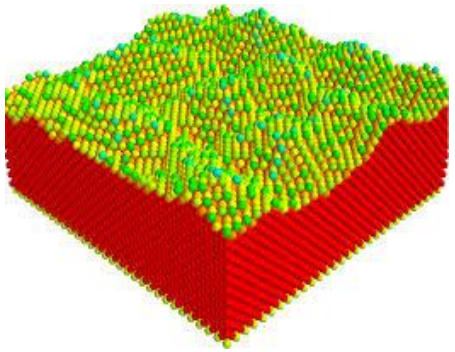

EAM on K20X, K40s & K80s

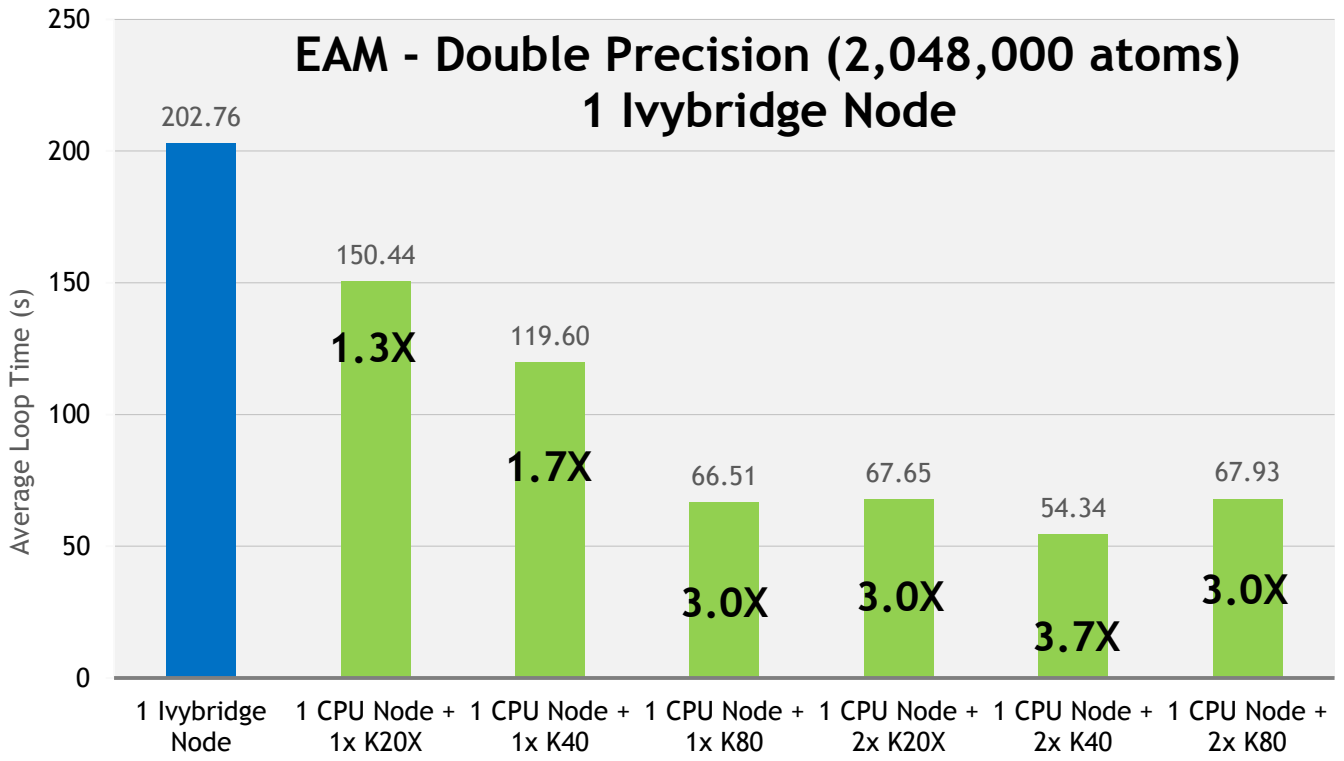

Running LAMMPS

The blue node contains Dual Intel Xeon E5-2697 v2@2.7GHz CPUs

The green nodes contain Dual Intel Xeon E5-2697 v2@2.7GHz CPUs + either NVIDIA Tesla K20X@732Mhz, Tesla K40@875Mhz or Tesla K80 (autoboost) GPUs

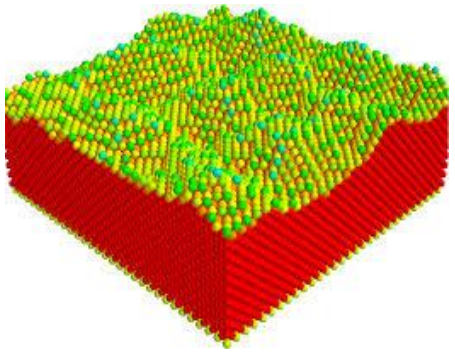

EAM on K20X, K40s & K80s

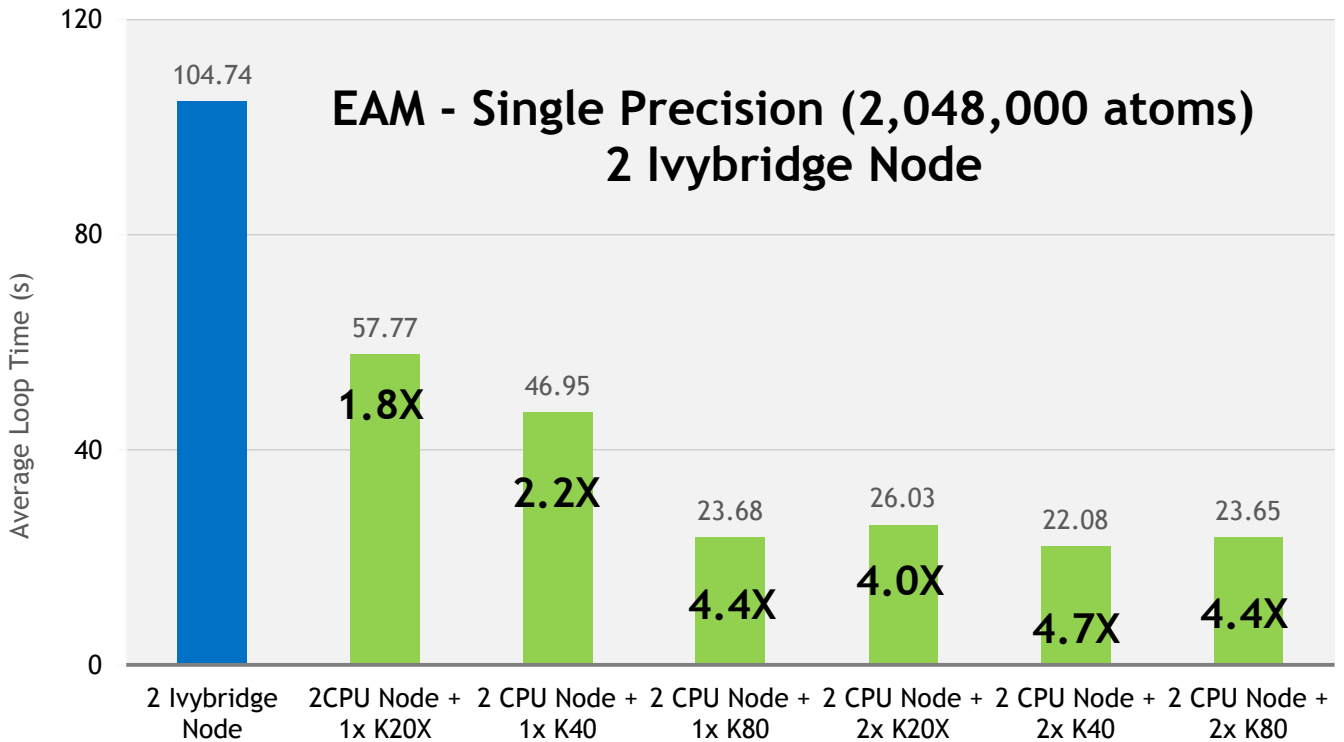

Running LAMMPS

The blue node contains Dual Intel Xeon E5-2697 v2@2.7GHz CPUs

The green nodes contain Dual Intel Xeon E5-2697 v2@2.7GHz CPUs + either NVIDIA Tesla K20X@732Mhz, Tesla K40@875Mhz or Tesla K80 (autoboost) GPUs

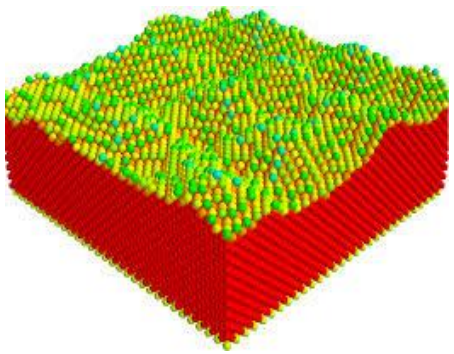

EAM on K20X, K40s & K80s

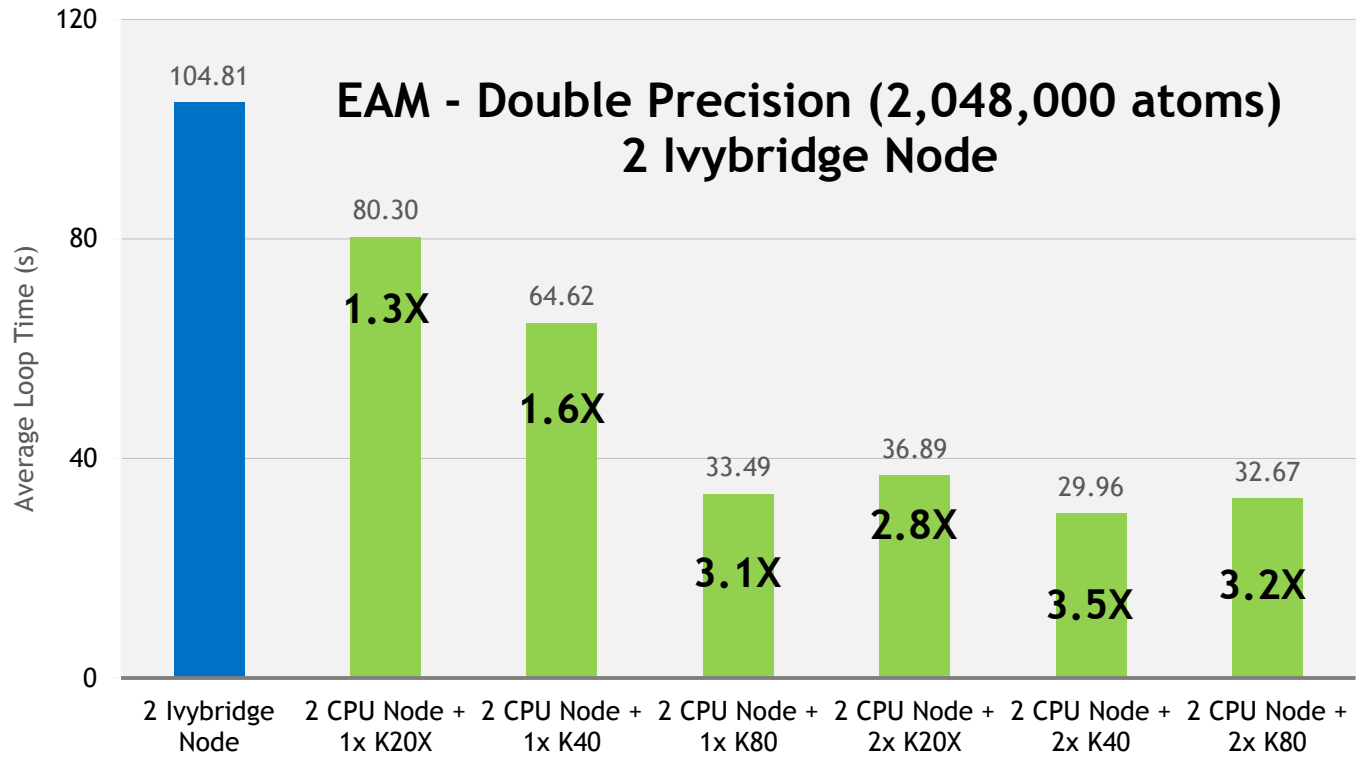

Running LAMMPS

The blue node contains Dual Intel Xeon E5-2697 v2@2.7GHz CPUs

The green nodes contain Dual Intel Xeon E5-2697 v2@2.7GHz CPUs + either NVIDIA Tesla K20X@732Mhz, Tesla K40@875Mhz or Tesla K80 (autoboost) GPUs

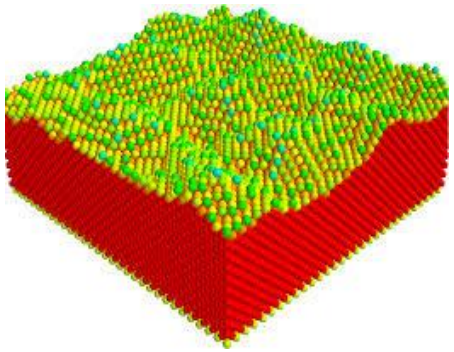

EAM on K20X, K40s & K80s

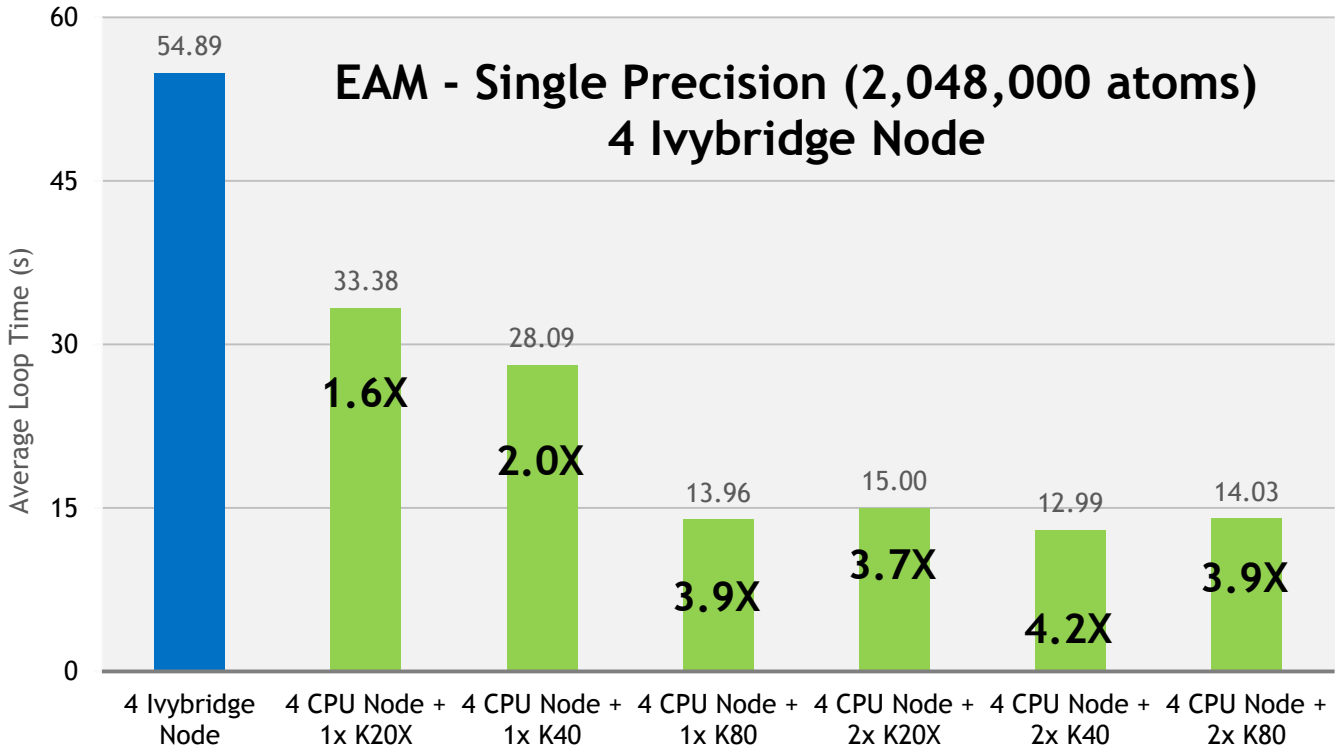

Running LAMMPS

The blue node contains Dual Intel Xeon E5-2697 v2@2.7GHz CPUs

The green nodes contain Dual Intel Xeon E5-2697 v2@2.7GHz CPUs + either NVIDIA Tesla K20X@732Mhz, Tesla K40@875Mhz or Tesla K80 (autoboost) GPUs

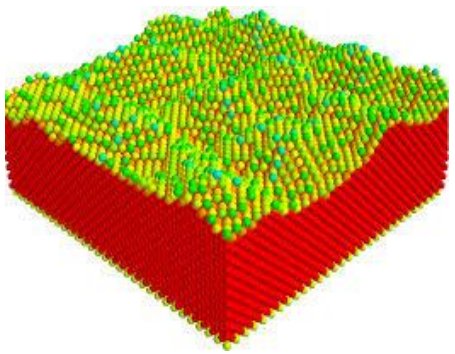

EAM on K20X, K40s & K80s

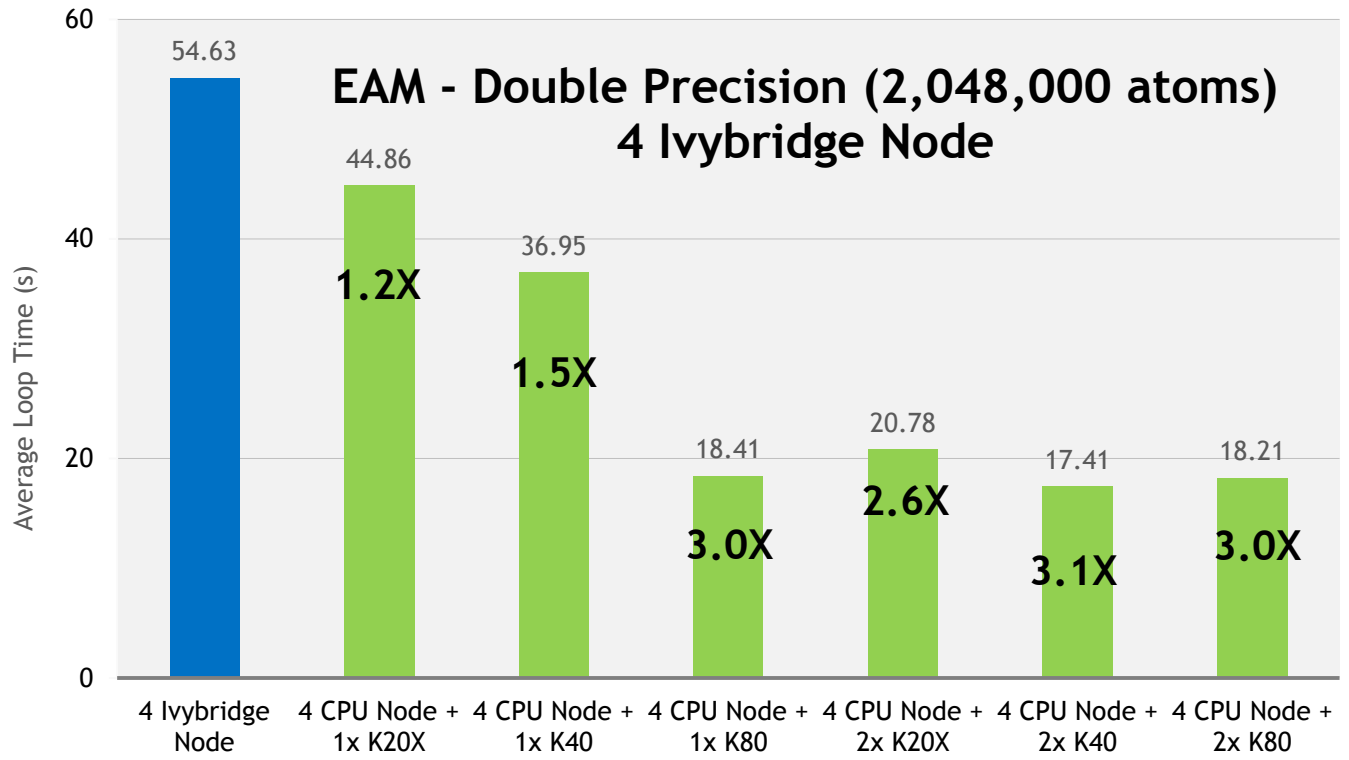

Running LAMMPS

The blue node contains Dual Intel Xeon E5-2697 v2@2.7GHz CPUs

The green nodes contain Dual Intel Xeon E5-2697 v2@2.7GHz CPUs + either NVIDIA Tesla K20X@732Mhz, Tesla K40@875Mhz or Tesla K80 (autoboost) GPUs

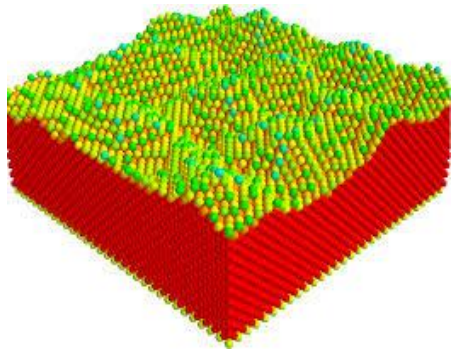

EAM on K20X and K40s

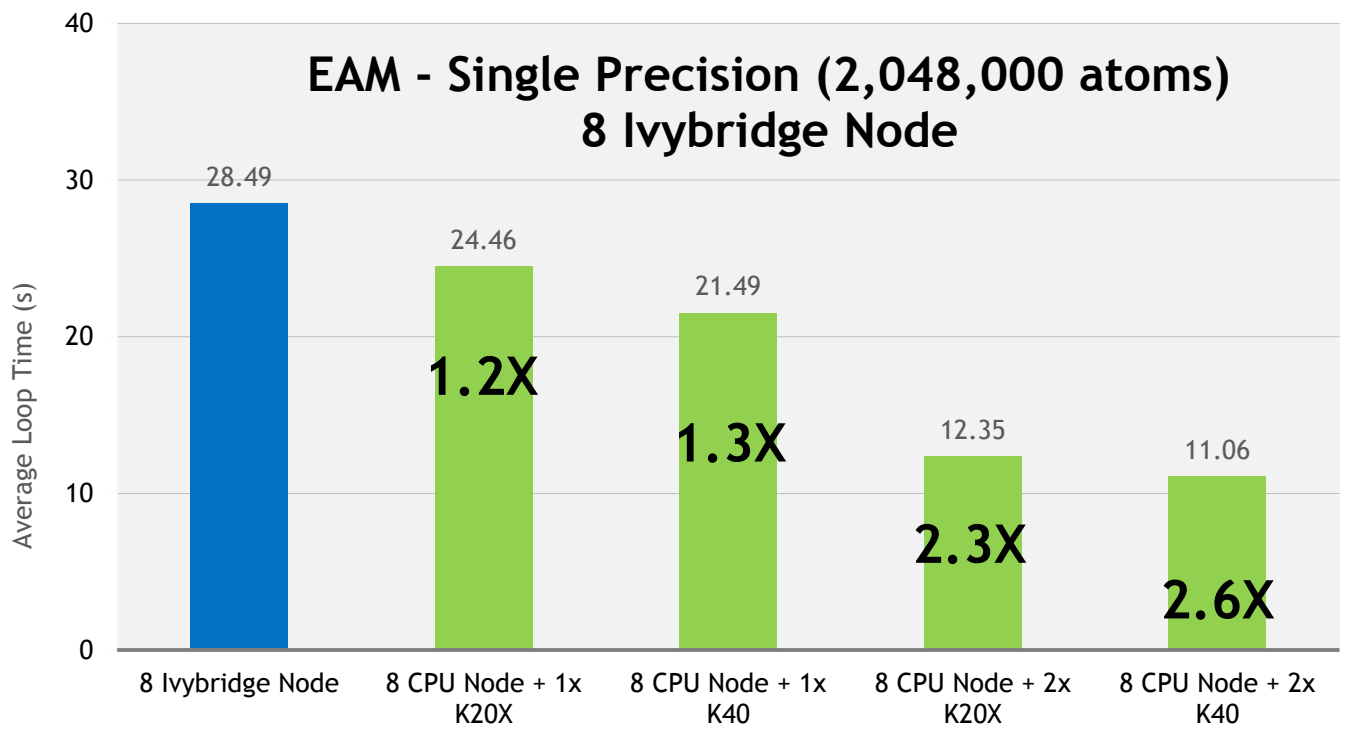

Running LAMMPS

The blue node contains Dual Intel Xeon E5-2697 v2@2.7GHz CPUs

The green nodes contain Dual Intel Xeon E5-2697 v2@2.7GHz CPUs + either NVIDIA Tesla K20X@732Mhz, Tesla K40@875Mhz or Tesla K80 (autoboost) GPUs

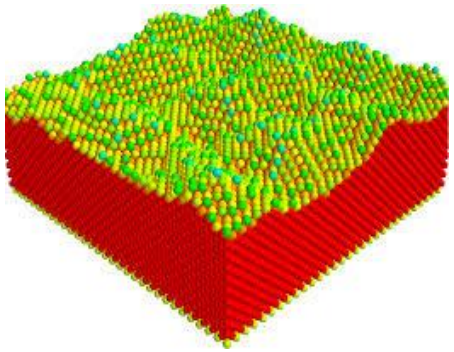

EAM on K20X and K40s

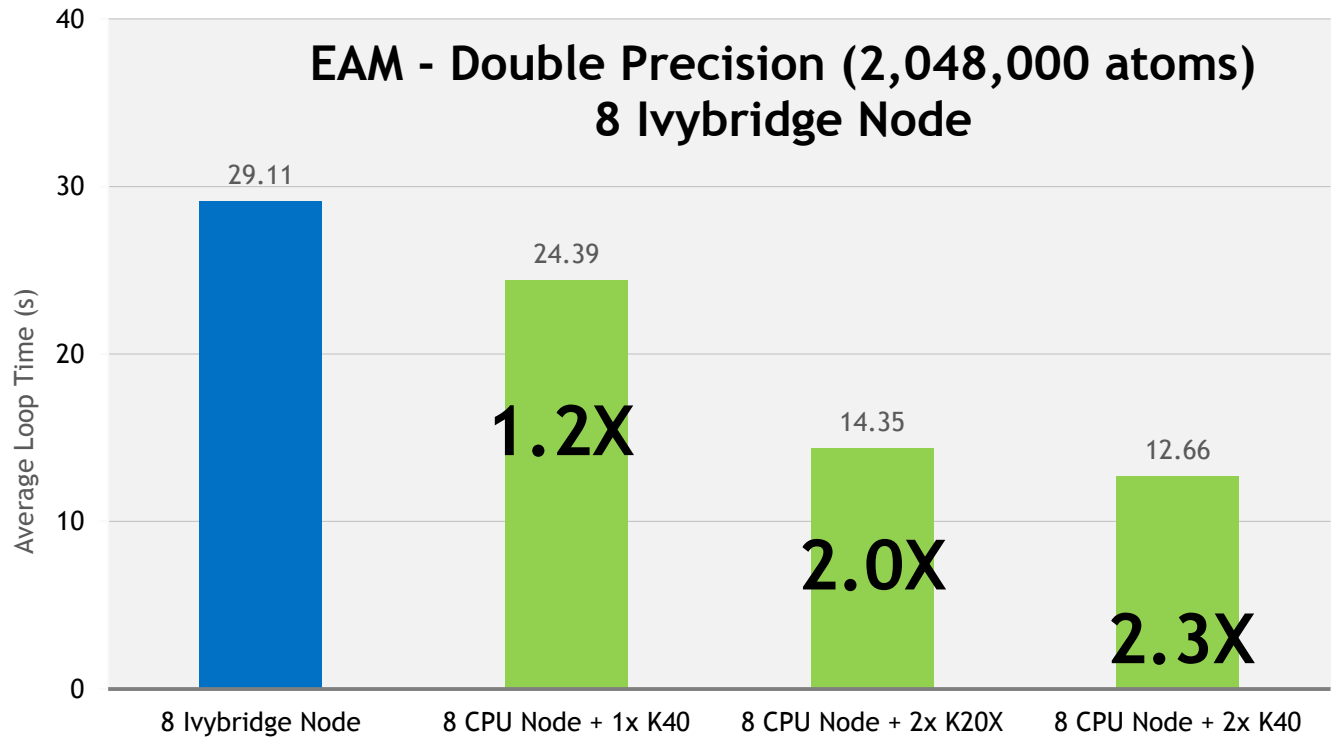

Running LAMMPS

The blue node contains Dual Intel Xeon E5-2697 v2@2.7GHz CPUs

The green nodes contain Dual Intel Xeon E5-2697 v2@2.7GHz CPUs + either NVIDIA Tesla K20X@732Mhz, Tesla K40@875Mhz or Tesla K80 (autoboost) GPUs

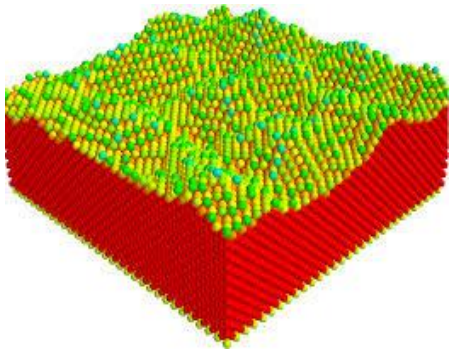

EAM single/multi-node throughput

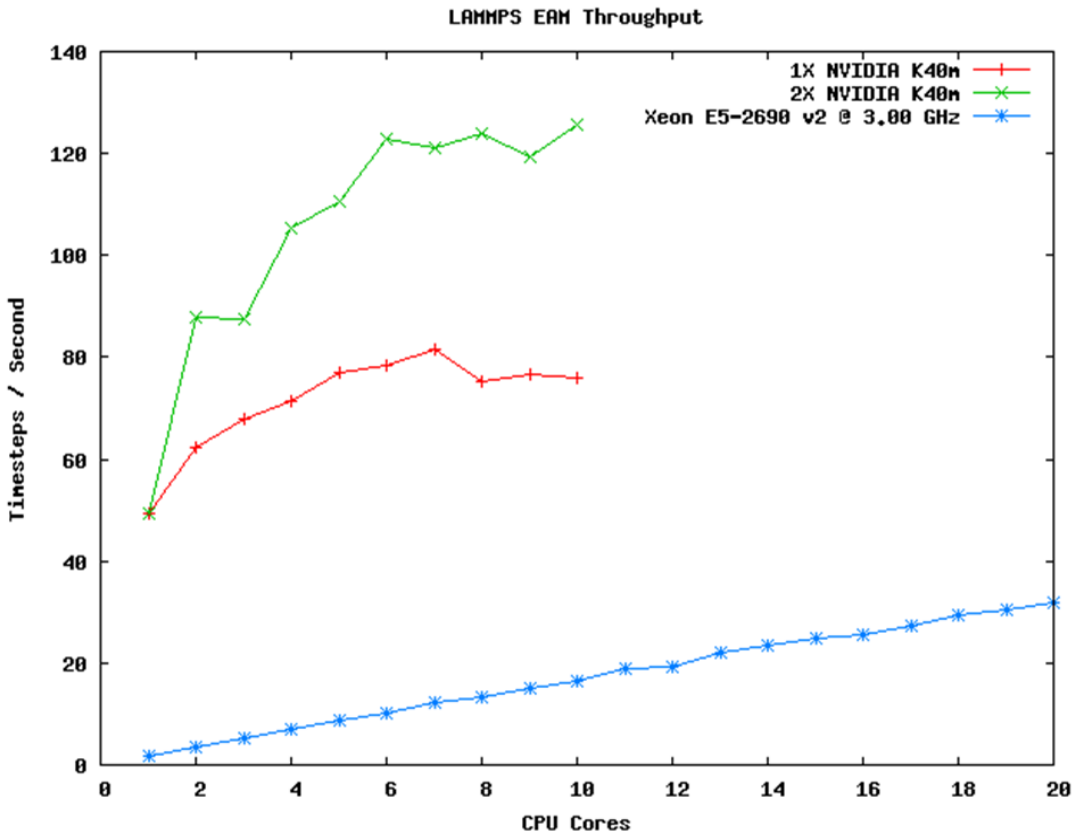

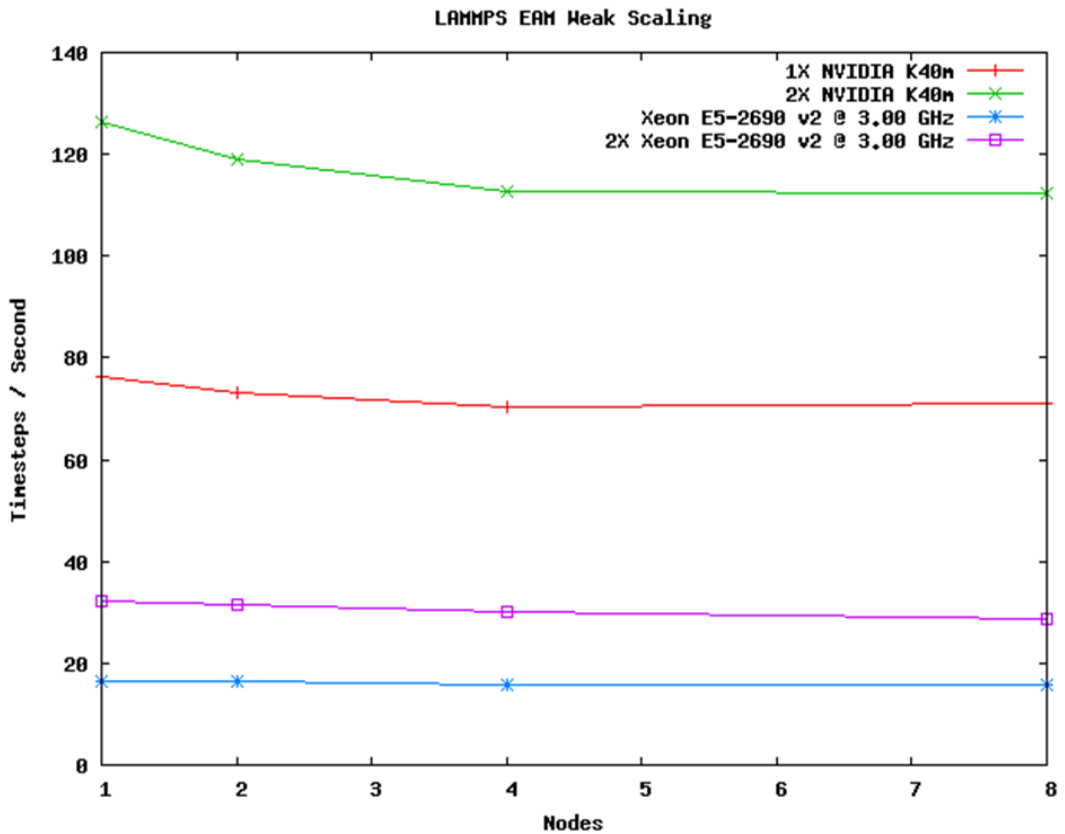

Figure 1: EAM single-node throughout (strong scaling)
Figure 2: EAM multi-scaling throughout

Gay-Berne on K20X, K40s & K80s

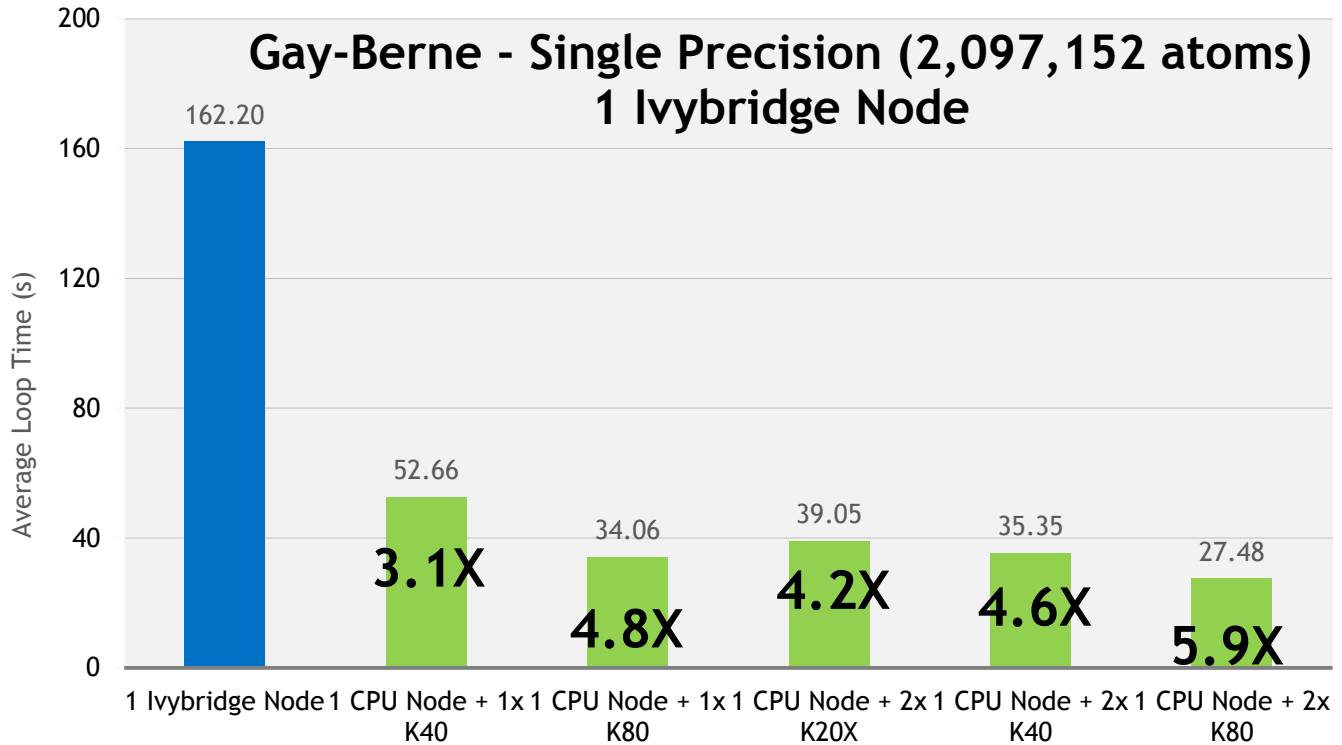

Running LAMMPS

The blue node contains Dual Intel Xeon E5-2697 v2@2.7GHz CPUs

The green nodes contain Dual Intel Xeon E5-2697 v2@2.7GHz CPUs + either NVIDIA Tesla K20X@732Mhz, Tesla K40@875Mhz or Tesla K80 (autoboost) GPUs

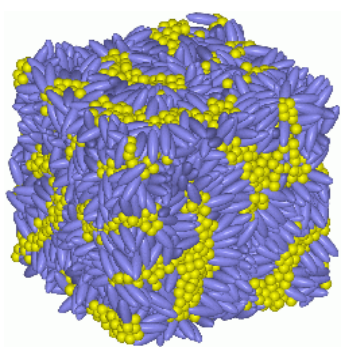

Gay-Berne on K20X, K40s & K80s

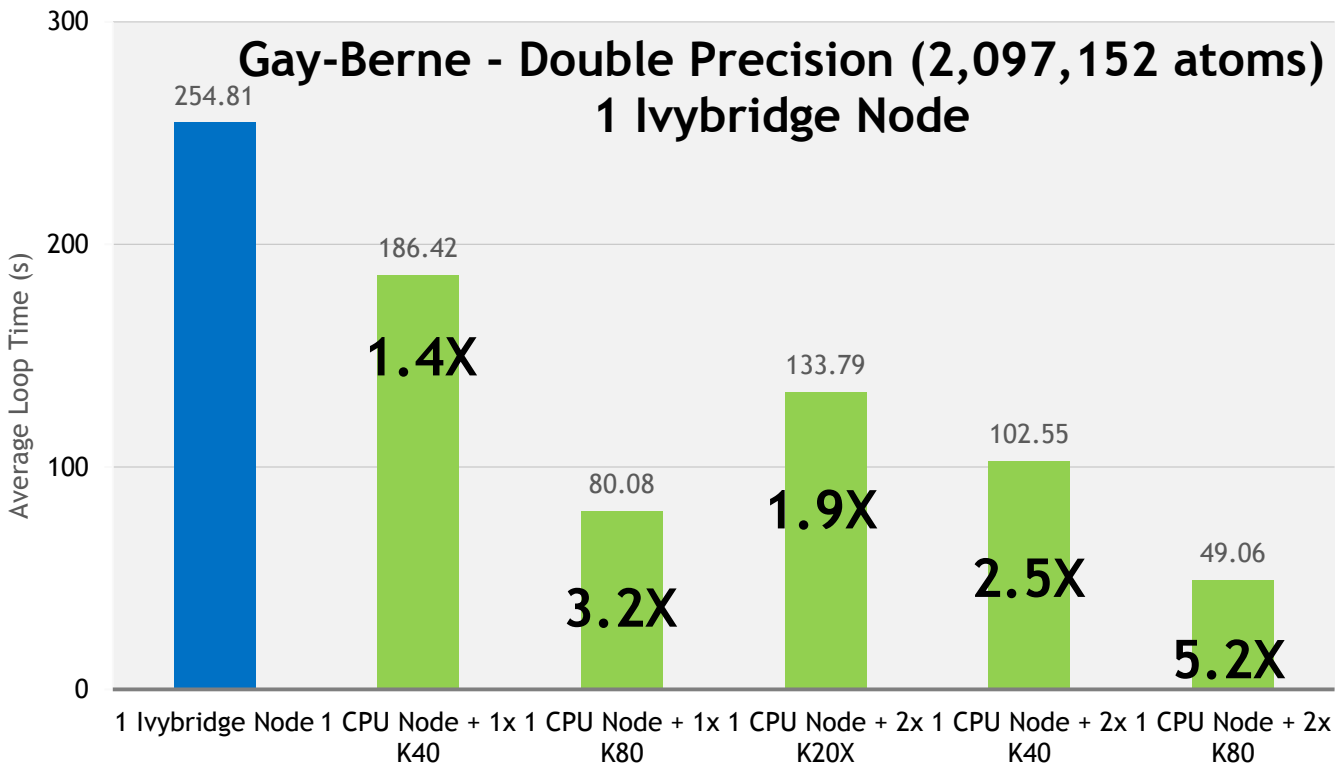

Running LAMMPS

The blue node contains Dual Intel Xeon E5-2697 v2@2.7GHz CPUs

The green nodes contain Dual Intel Xeon E5-2697 v2@2.7GHz CPUs + either NVIDIA Tesla K20X@732Mhz, Tesla K40@875Mhz or Tesla K80 (autoboost) GPUs

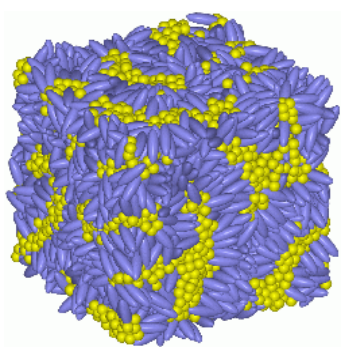

Gay-Berne on K20X, K40s & K80s

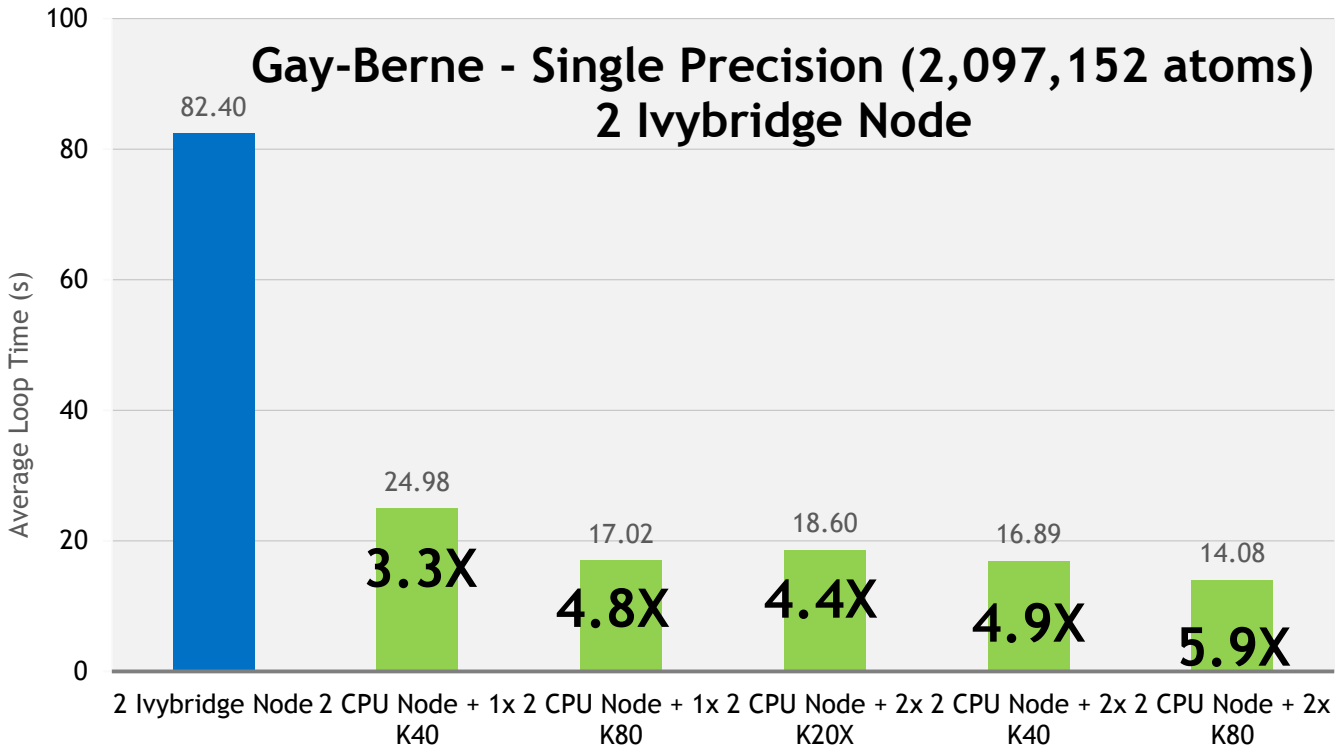

Running LAMMPS

The blue node contains Dual Intel Xeon E5-2697 v2@2.7GHz CPUs

The green nodes contain Dual Intel Xeon E5-2697 v2@2.7GHz CPUs + either NVIDIA Tesla K20X@732Mhz, Tesla K40@875Mhz or Tesla K80 (autoboost) GPUs

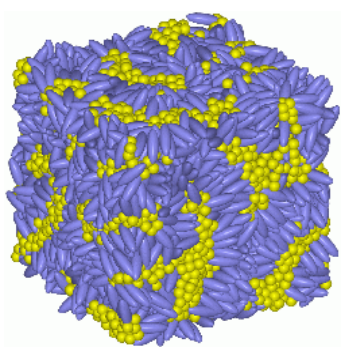

Gay-Berne on K20X, K40s & K80s

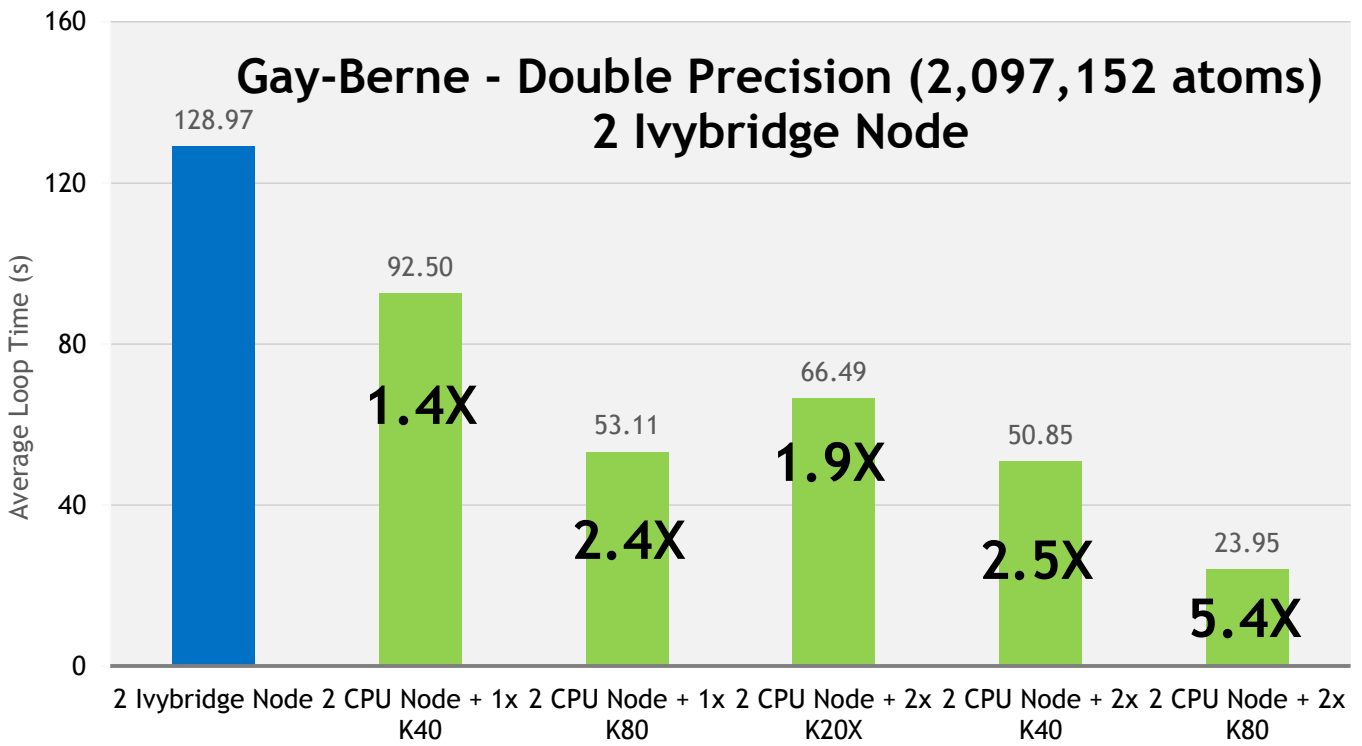

Running LAMMPS

The blue node contains Dual Intel Xeon E5-2697 v2@2.7GHz CPUs

The green nodes contain Dual Intel Xeon E5-2697 v2@2.7GHz CPUs + either NVIDIA Tesla K20X@732Mhz, Tesla K40@875Mhz or Tesla K80 (autoboost) GPUs

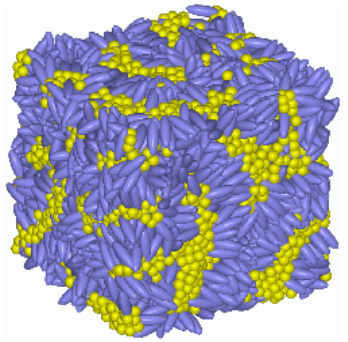

Gay-Berne on K20X, K40s & K80s

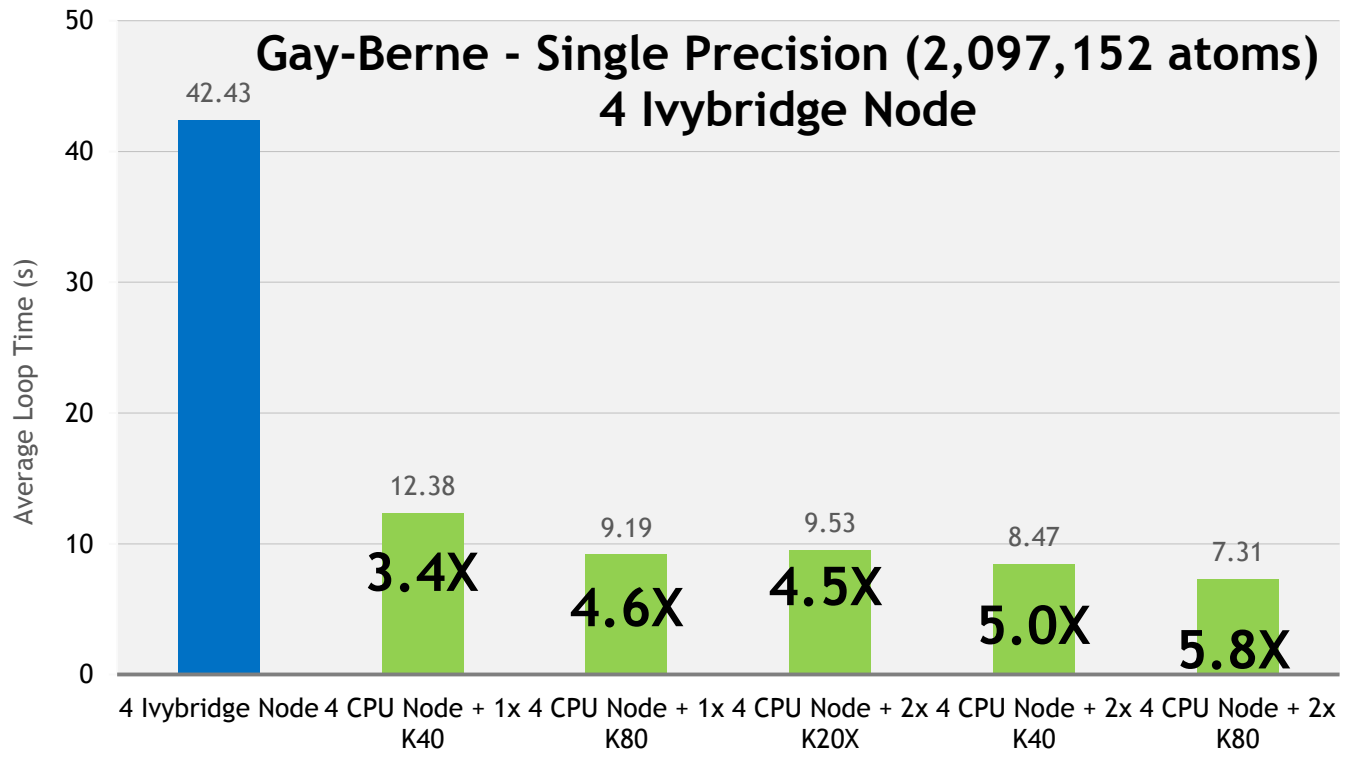

Running LAMMPS

The blue node contains Dual Intel Xeon E5-2697 v2@2.7GHz CPUs

The green nodes contain Dual Intel Xeon E5-2697 v2@2.7GHz CPUs + either NVIDIA Tesla K20X@732Mhz, Tesla K40@875Mhz or Tesla K80 (autoboost) GPUs

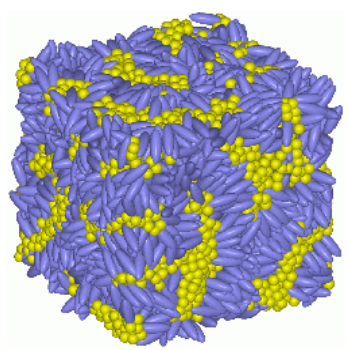

Gay-Berne on K20X, K40s & K80s

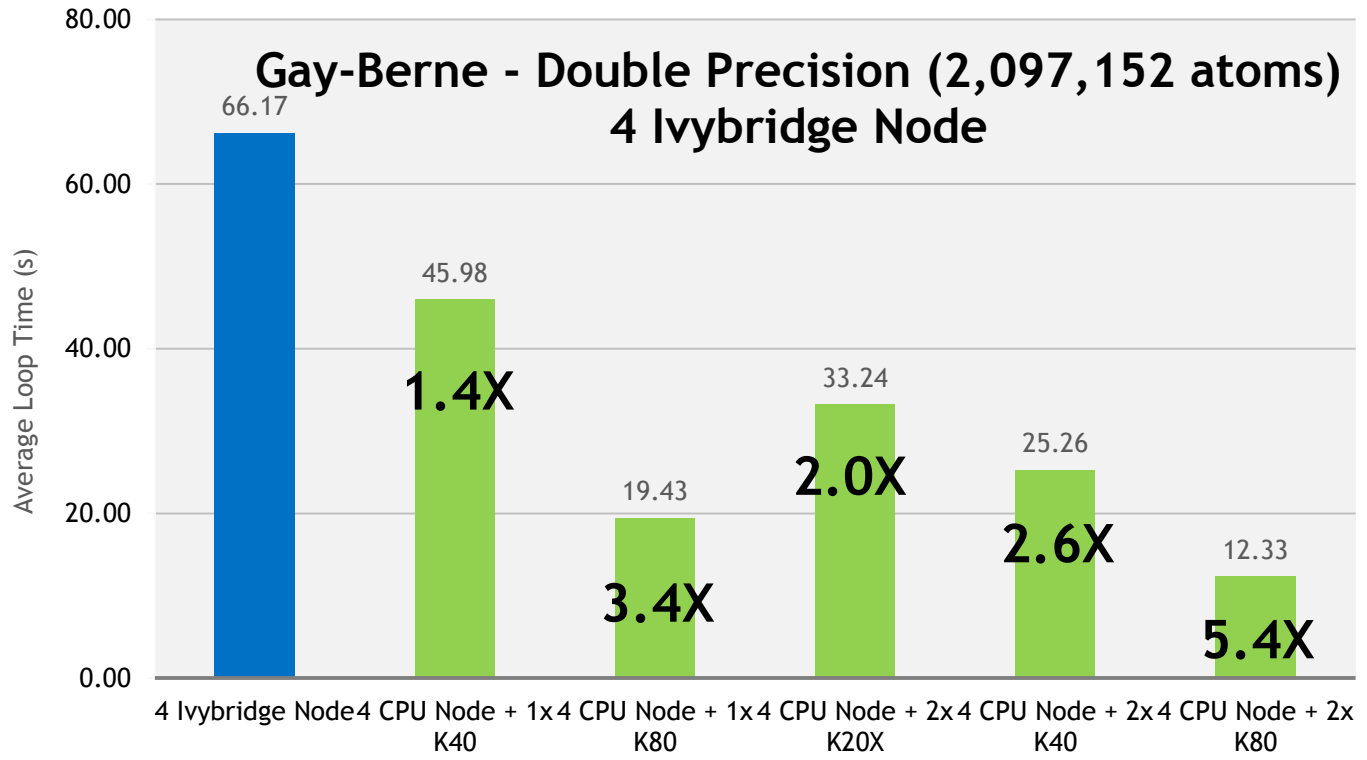

Running LAMMPS

The blue node contains Dual Intel Xeon E5-2697 v2@2.7GHz CPUs

The green nodes contain Dual Intel Xeon E5-2697 v2@2.7GHz CPUs + either NVIDIA Tesla K20X@732Mhz, Tesla K40@875Mhz or Tesla K80 (autoboost) GPUs

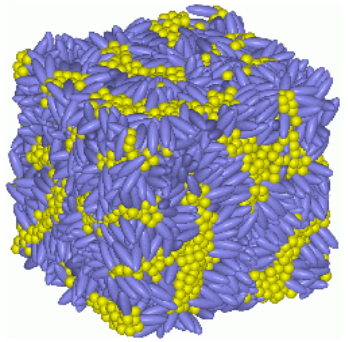

Gay-Berne on K20X and K40s

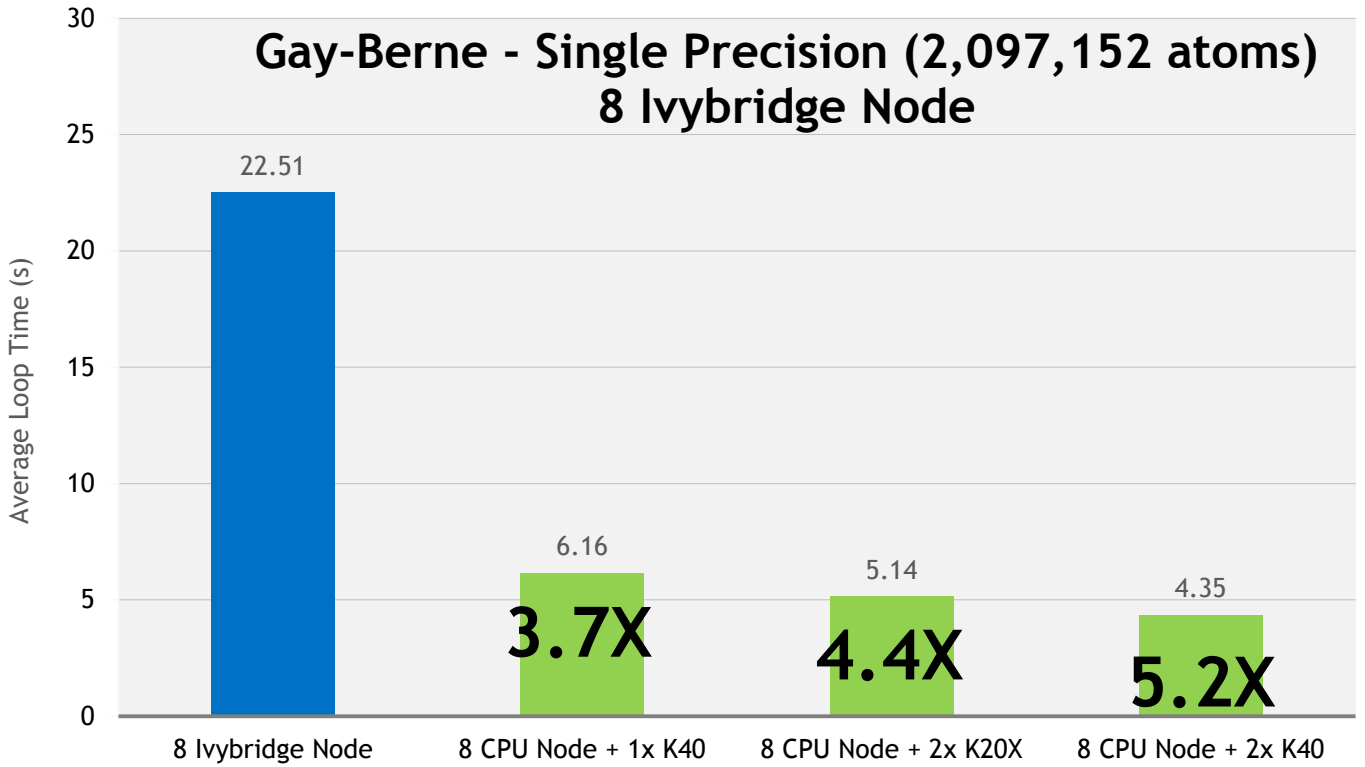

Running LAMMPS

The blue node contains Dual Intel Xeon E5-2697 v2@2.7GHz CPUs

The green nodes contain Dual Intel Xeon E5-2697 v2@2.7GHz CPUs + either NVIDIA Tesla K20X@732Mhz, Tesla K40@875Mhz or Tesla K80 (autoboost) GPUs

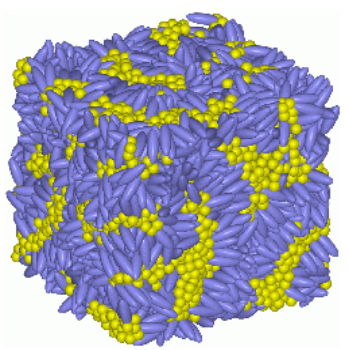

Gay-Berne on K20X and K40s

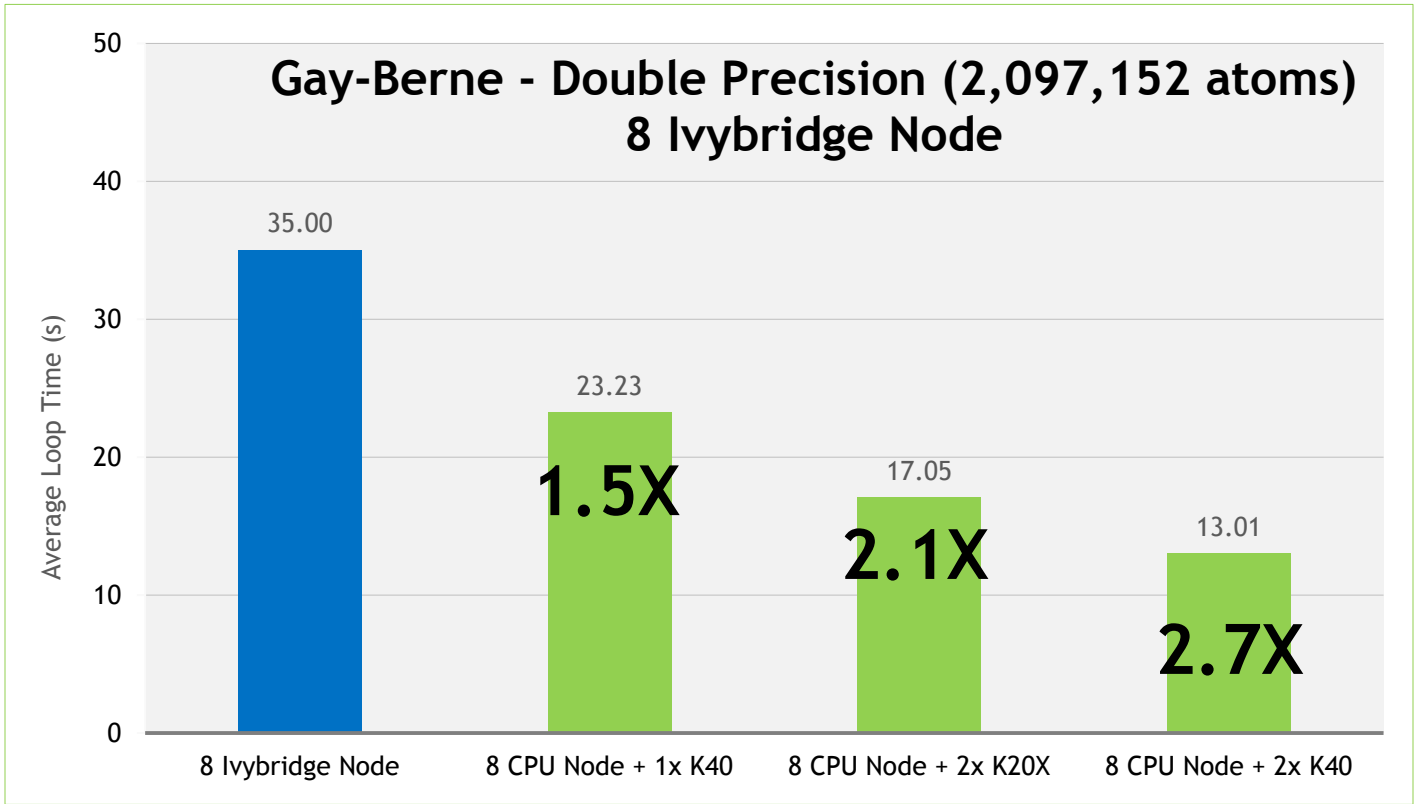

Running LAMMPS

The blue node contains Dual Intel Xeon E5-2697 v2@2.7GHz CPUs

The green nodes contain Dual Intel Xeon E5-2697 v2@2.7GHz CPUs + either NVIDIA Tesla K20X@732Mhz, Tesla K40@875Mhz or Tesla K80 (autoboost) GPUs

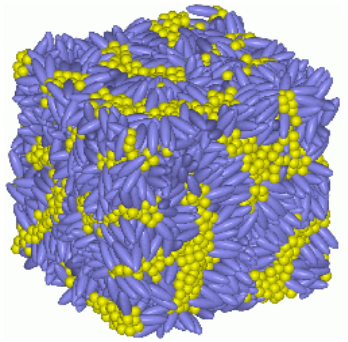

Rhodopsin on K20X, K40s & K80s

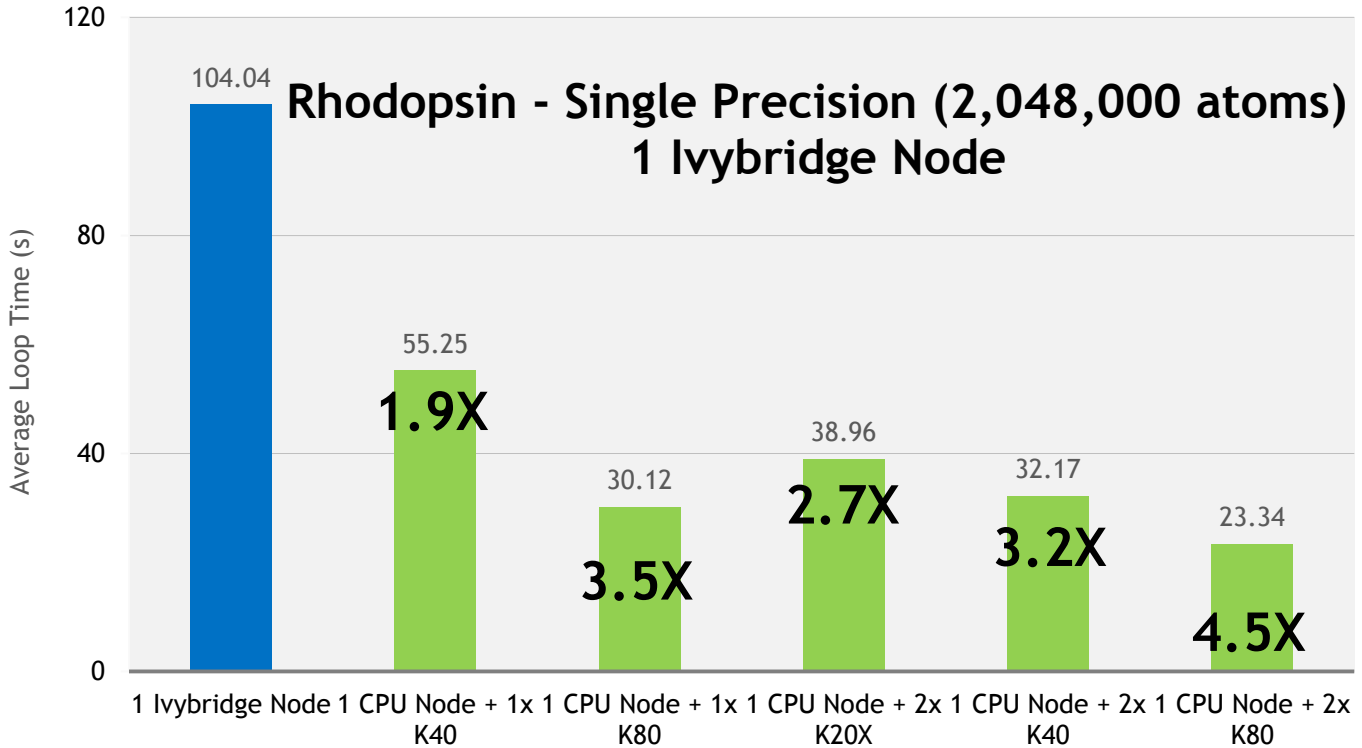

Running LAMMPS

The blue node contains Dual Intel Xeon E5-2697 v2@2.7GHz CPUs

The green nodes contain Dual Intel Xeon E5-2697 v2@2.7GHz CPUs + either NVIDIA Tesla K20X@732Mhz, Tesla K40@875Mhz or Tesla K80 (autoboost) GPUs

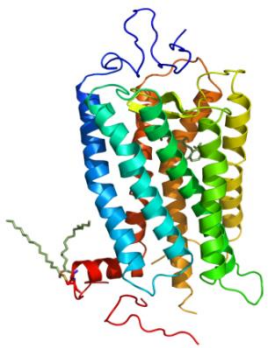

Rhodopsin on K20X, K40s & K80s

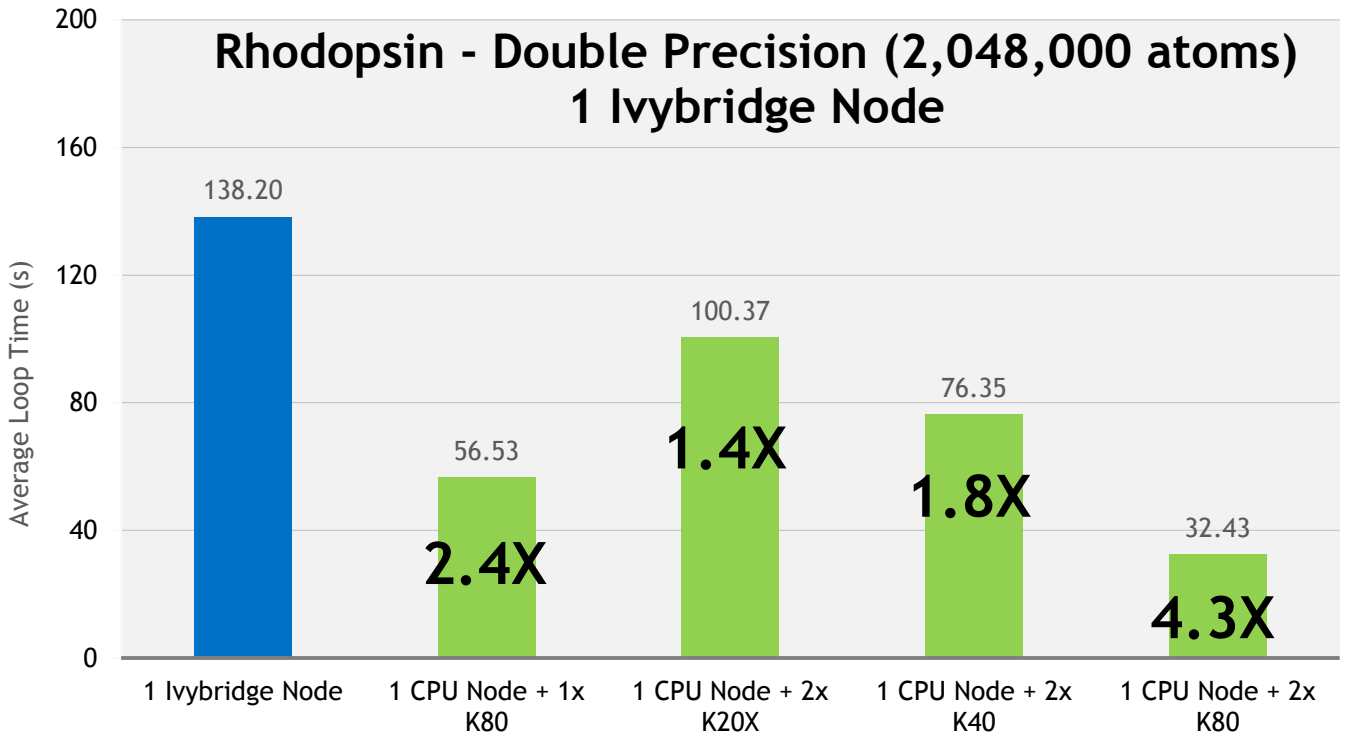

Running LAMMPS

The blue node contains Dual Intel Xeon E5-2697 v2@2.7GHz CPUs

The green nodes contain Dual Intel Xeon E5-2697 v2@2.7GHz CPUs + either NVIDIA Tesla K20X@732Mhz, Tesla K40@875Mhz or Tesla K80 (autoboost) GPUs

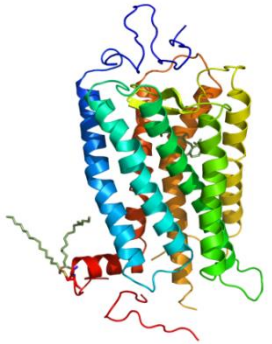

Rhodopsin on K20X, K40s & K80s

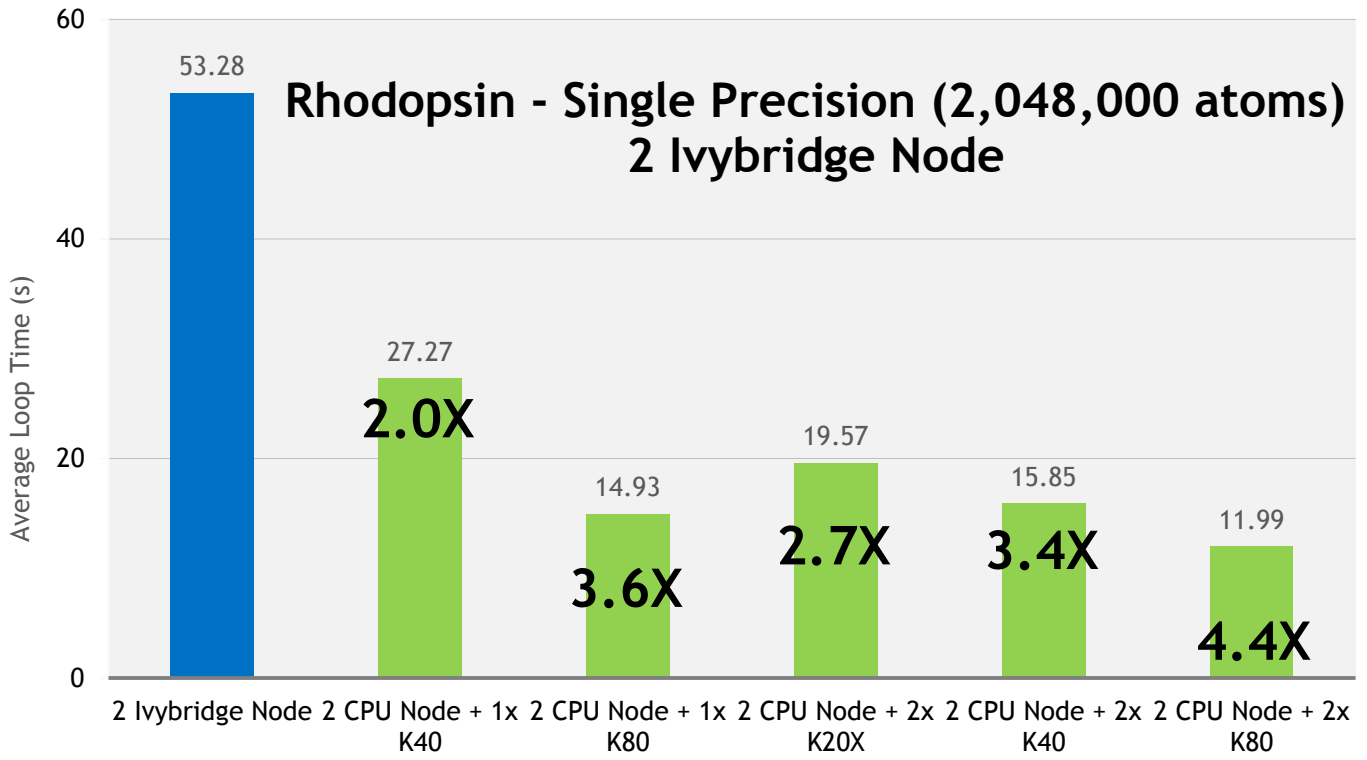

Running LAMMPS

The blue node contains Dual Intel Xeon E5-2697 v2@2.7GHz CPUs

The green nodes contain Dual Intel Xeon E5-2697 v2@2.7GHz CPUs + either NVIDIA Tesla K20X@732Mhz, Tesla K40@875Mhz or Tesla K80 (autoboost) GPUs

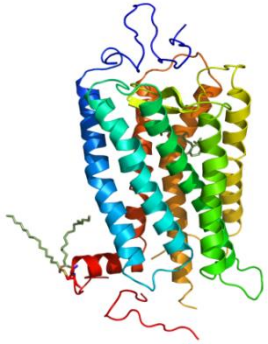

Rhodopsin on K20X, K40s & K80s

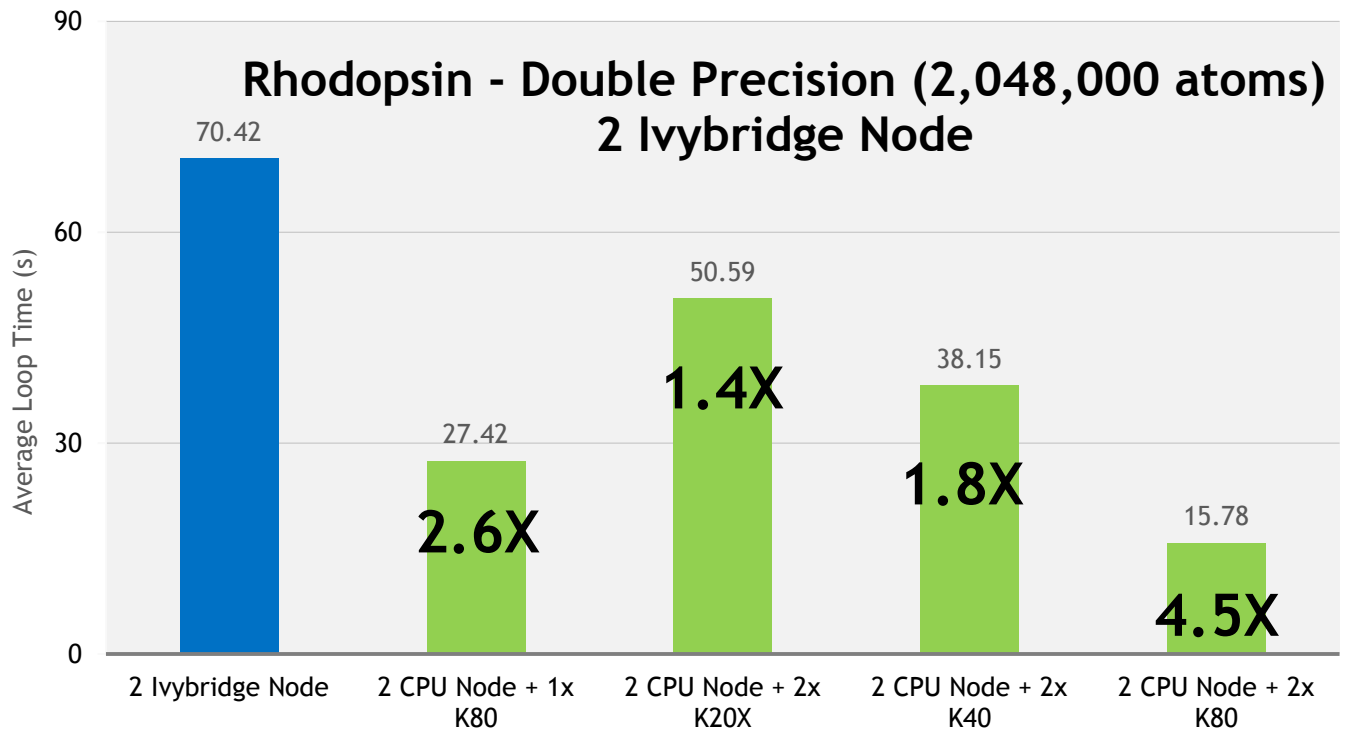

Running LAMMPS

The blue node contains Dual Intel Xeon E5-2697 v2@2.7GHz CPUs

The green nodes contain Dual Intel Xeon E5-2697 v2@2.7GHz CPUs + either NVIDIA Tesla K20X@732Mhz, Tesla K40@875Mhz or Tesla K80 (autoboost) GPUs

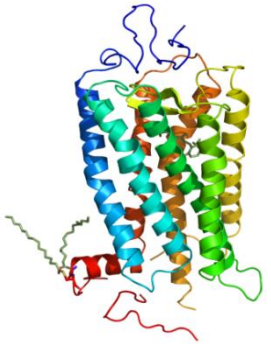

Rhodopsin on K20X, K40s & K80s

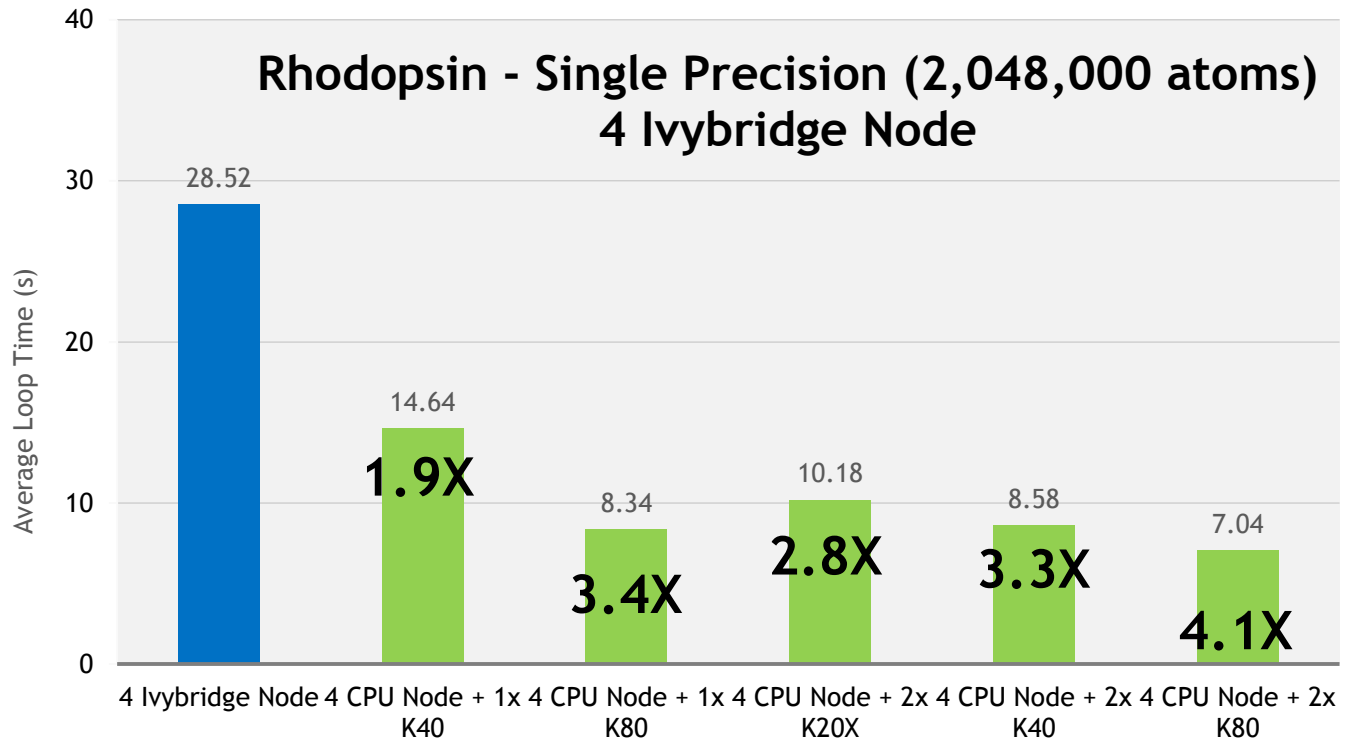

Running LAMMPS

The blue node contains Dual Intel Xeon E5-2697 v2@2.7GHz CPUs

The green nodes contain Dual Intel Xeon E5-2697 v2@2.7GHz CPUs + either NVIDIA Tesla K20X@732Mhz, Tesla K40@875Mhz or Tesla K80 (autoboost) GPUs

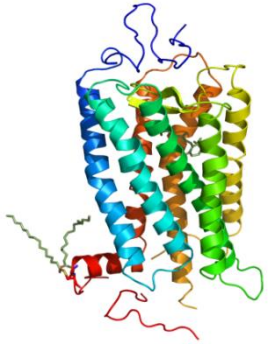

Rhodopsin on K20X, K40s & K80s

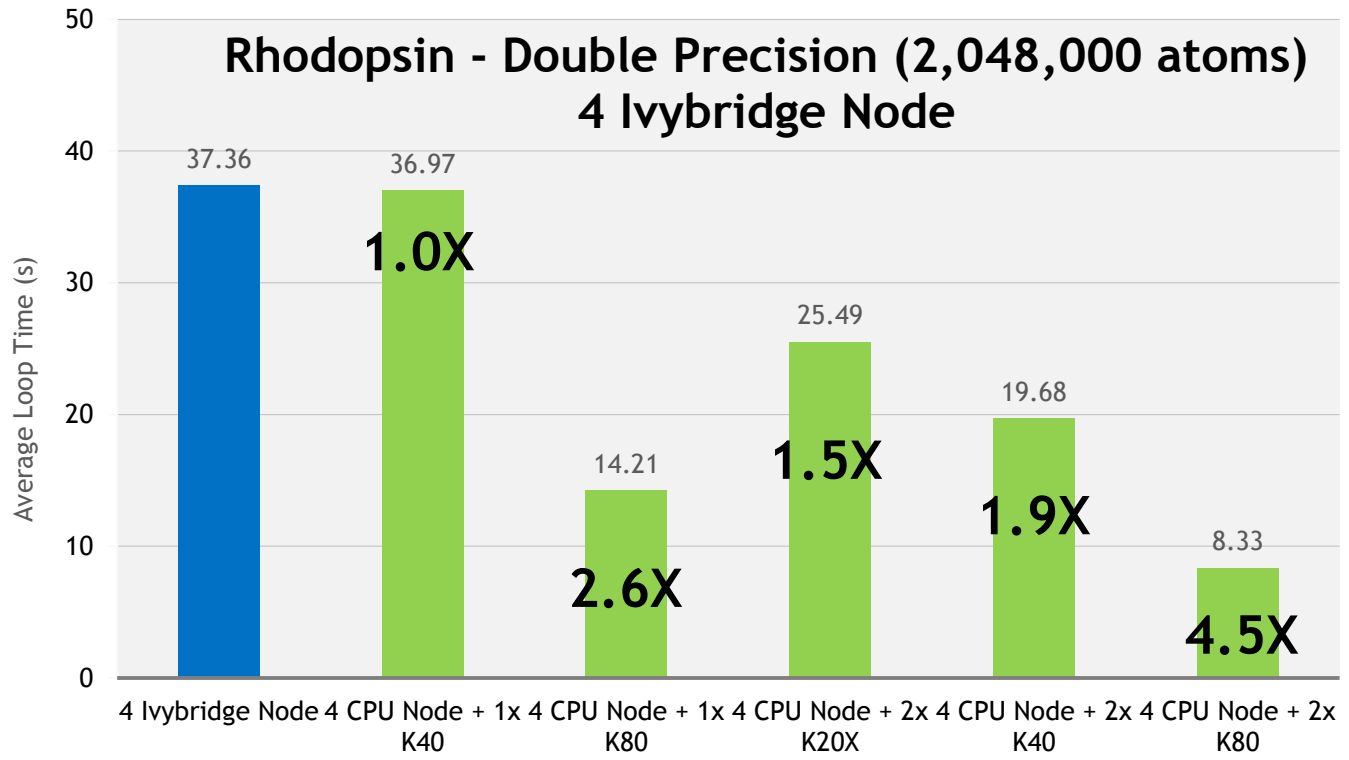

Running LAMMPS

The blue node contains Dual Intel E5-2698 v3@2.3GHz CPUs

The green nodes contain Dual Intel E5-2698 v3@2.3GHz CPUs + either NVIDIA Tesla K20X@732Mhz, Tesla K40@875Mhz or Tesla K80@562Mhz GPUs

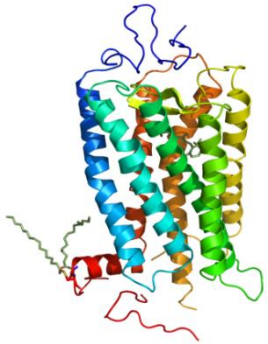

Rhodopsin on K20X and K40s

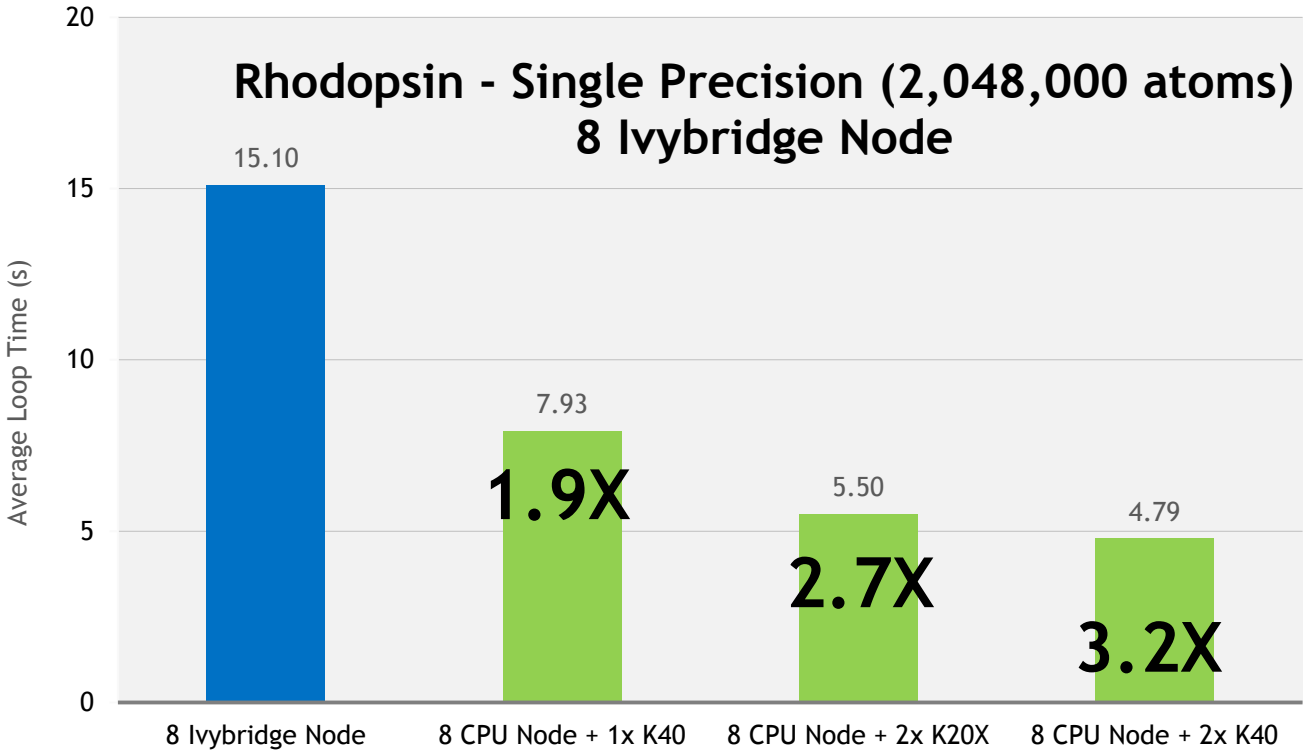

Running LAMMPS

The blue node contains Dual Intel Xeon E5-2697 v2@2.7GHz CPUs

The green nodes contain Dual Intel Xeon E5-2697 v2@2.7GHz CPUs + either NVIDIA Tesla K20X@732Mhz, Tesla K40@875Mhz or Tesla K80 (autoboost) GPUs

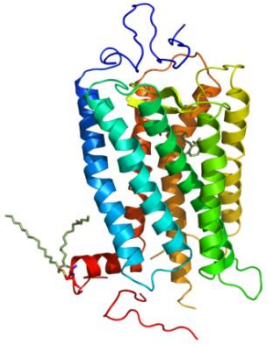

Rhodopsin on K20X and K40s

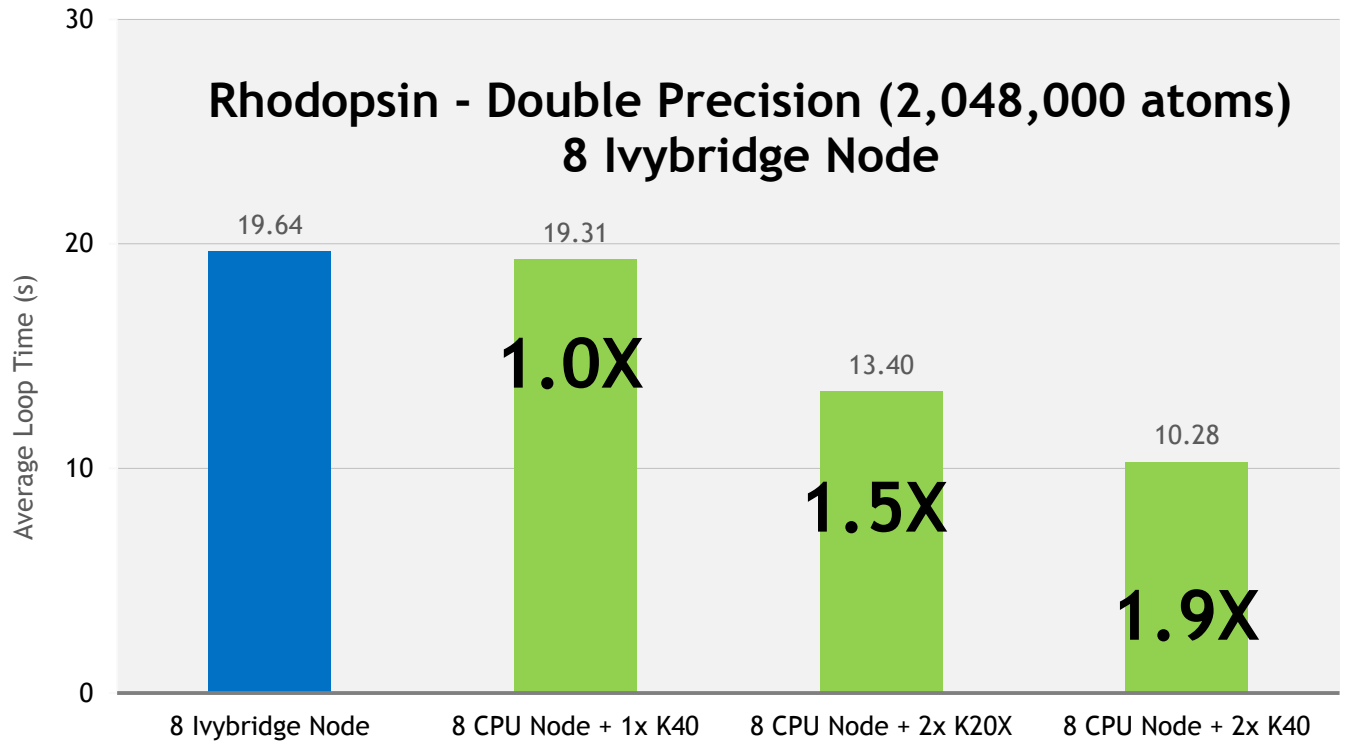

Running LAMMPS

The blue node contains Dual Intel Xeon E5-2697 v2@2.7GHz CPUs

The green nodes contain Dual Intel Xeon E5-2697 v2@2.7GHz CPUs + either NVIDIA Tesla K20X@732Mhz, Tesla K40@875Mhz or Tesla K80 (autoboost) GPUs

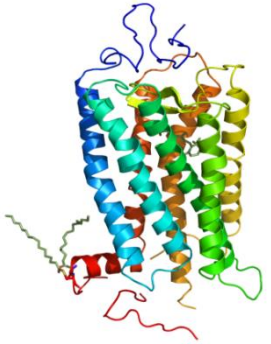

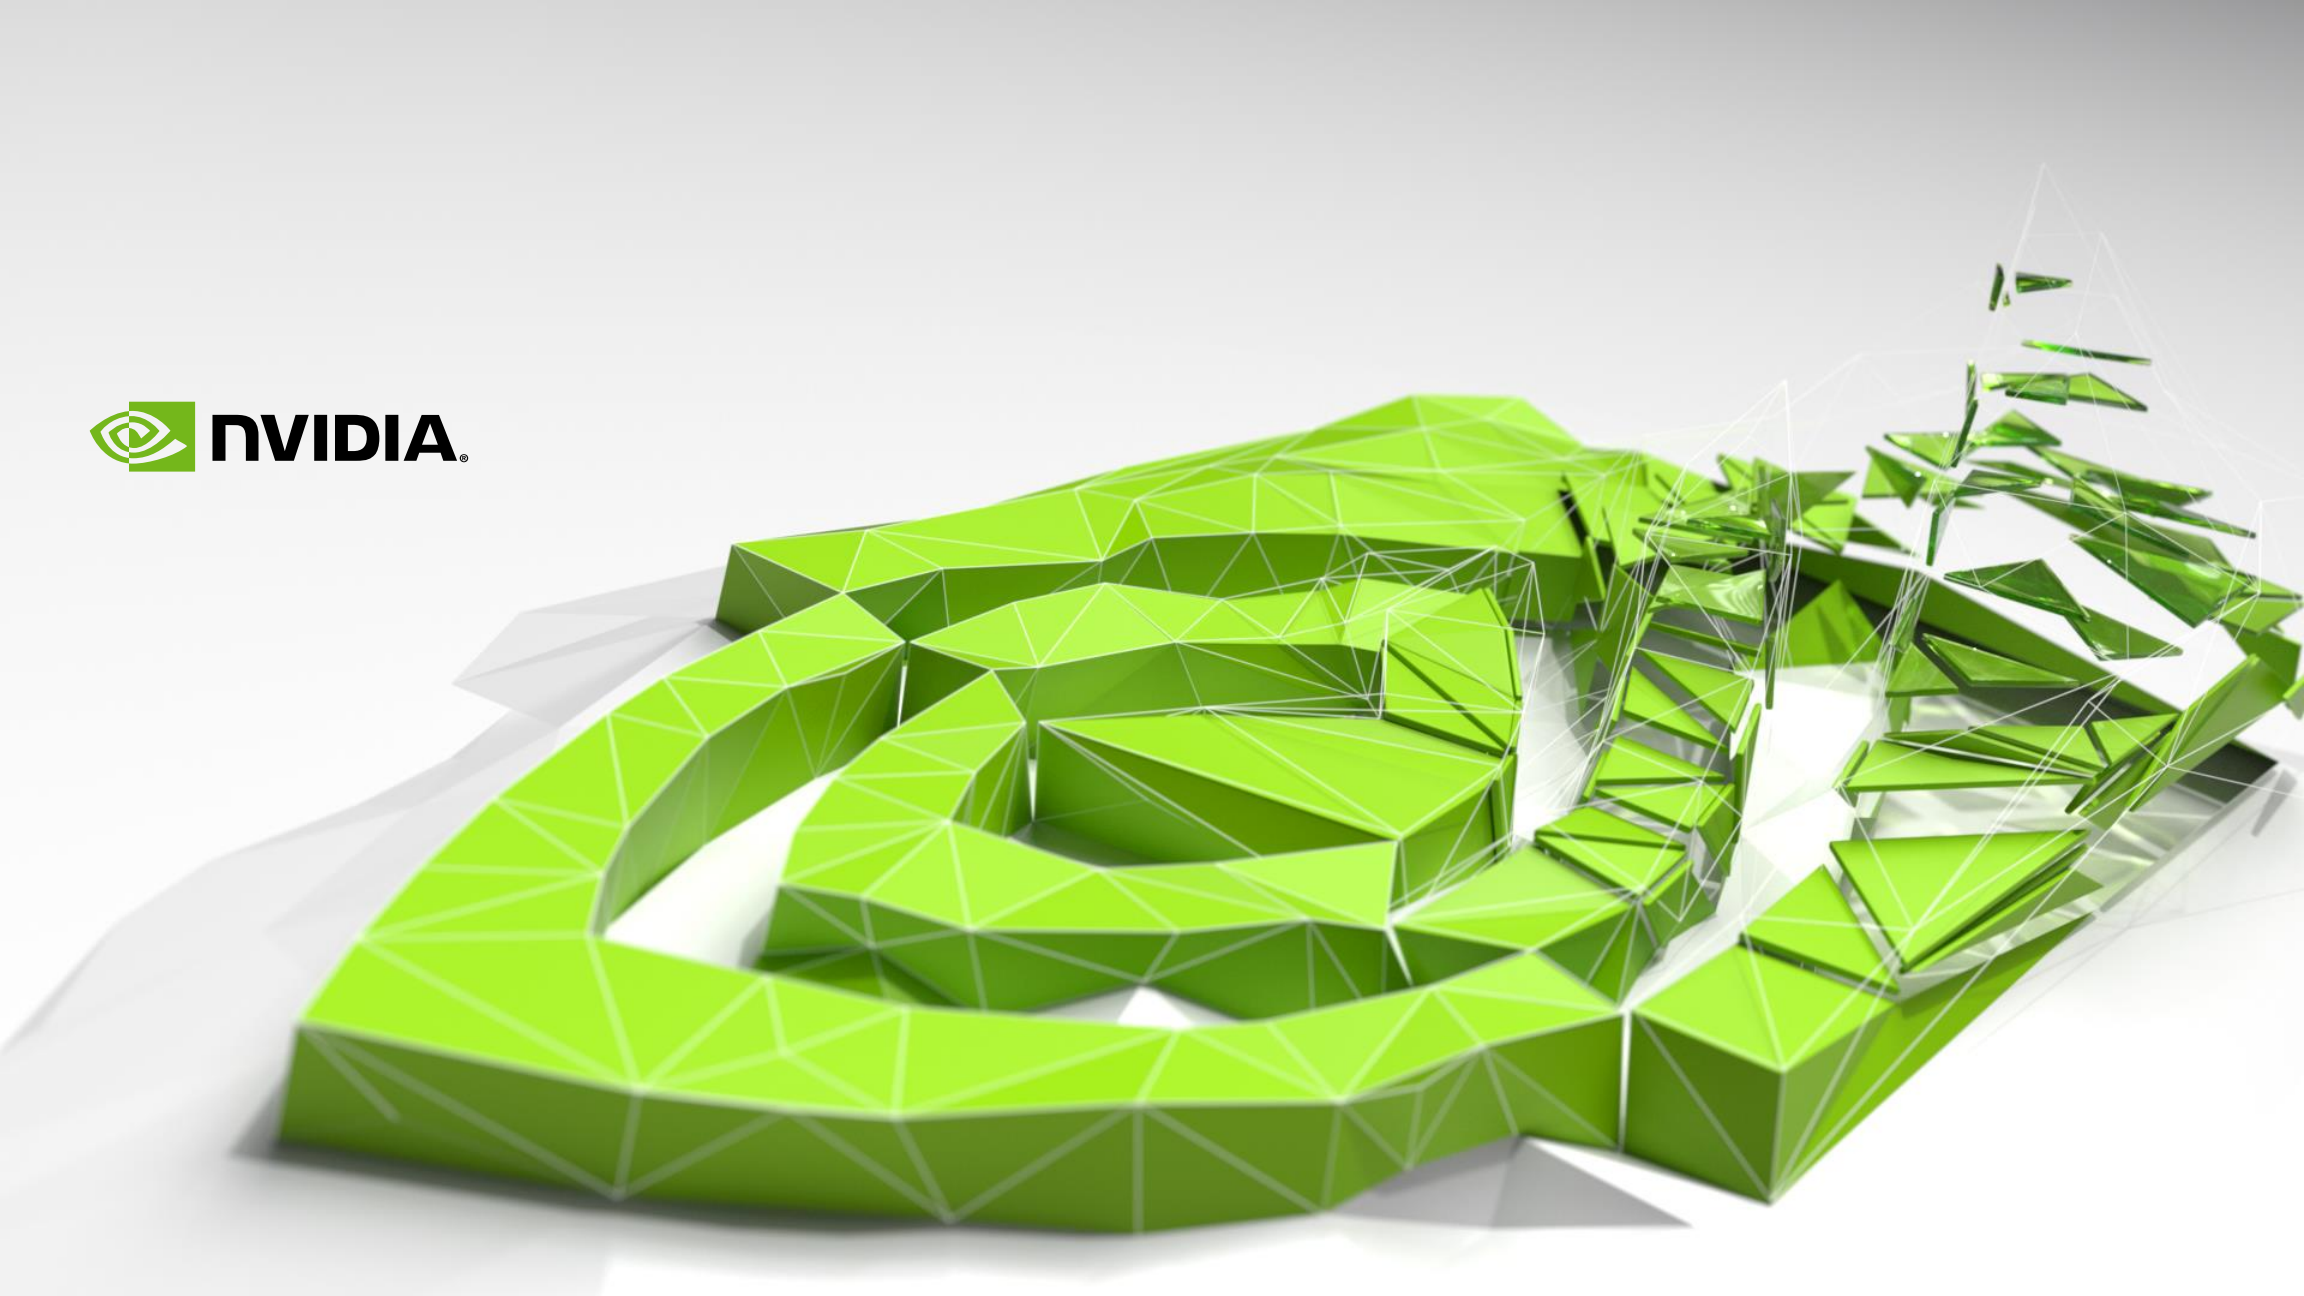
NVIDIA

More Science for Your Money (LAMMPS)

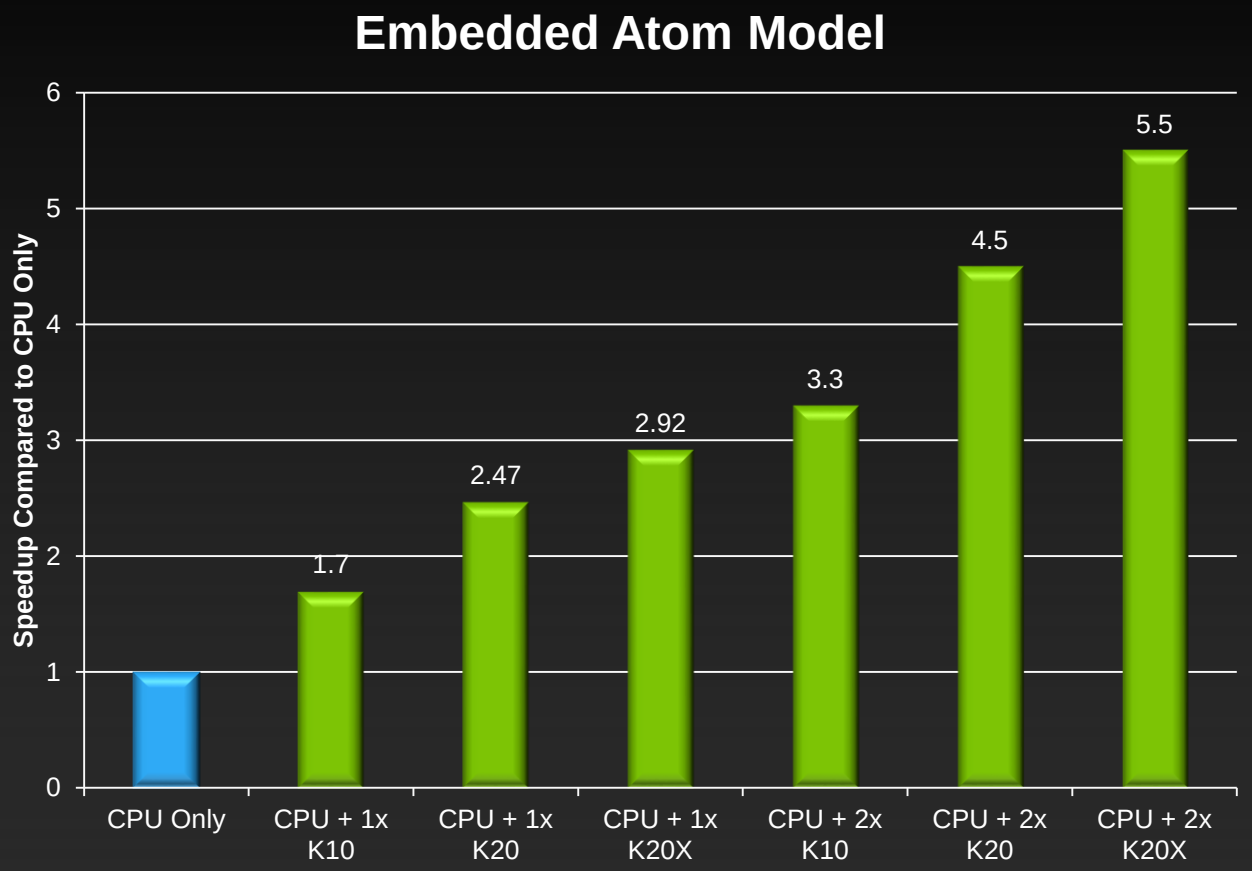

Blue node uses 2x E5-2687W (8 Cores and 150W per CPU).

Green nodes have 2x E5-2687W and 1 or 2 NVIDIA K10, K20, or K20X GPUs (235W).

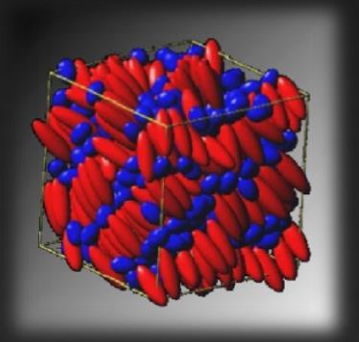

Experience performance increases of up to 5.5x with Kepler GPU nodes.

Excellent Strong Scaling on Large Clusters

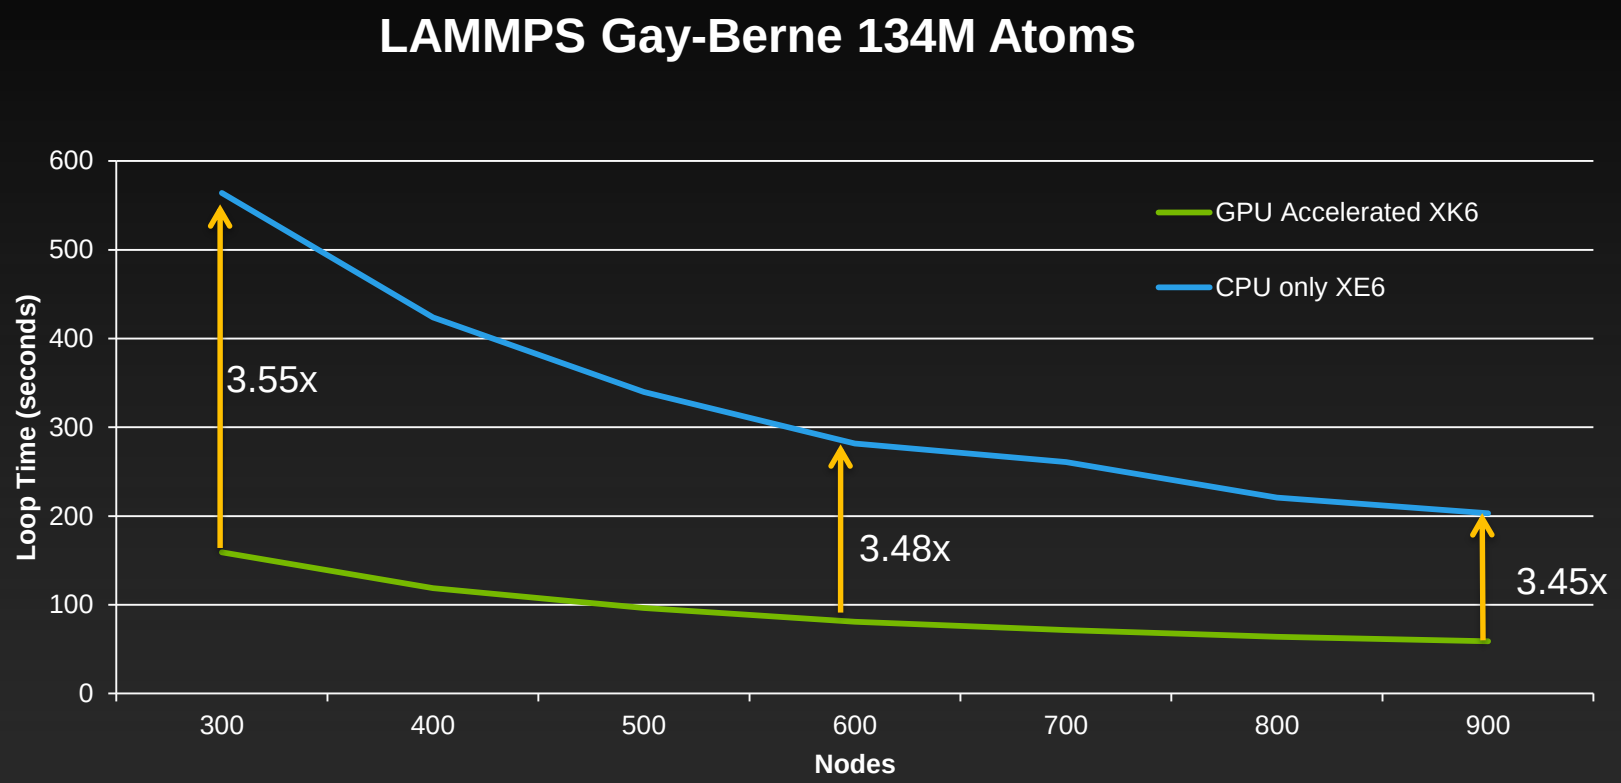

From 300-900 nodes, the NVIDIA GPU-powered XK6 maintained 3.5x performance compared to XE6 CPU nodes

Each blue Cray XE6 Nodes have 2x AMD Opteron CPUs (16 Cores per CPU)
Each green Cray XK6 Node has 1x AMD Opteron 1600 CPU (16 Cores per CPU) and 1x NVIDIA X2090

GPUs Sustain 5x Performance for Weak Scaling

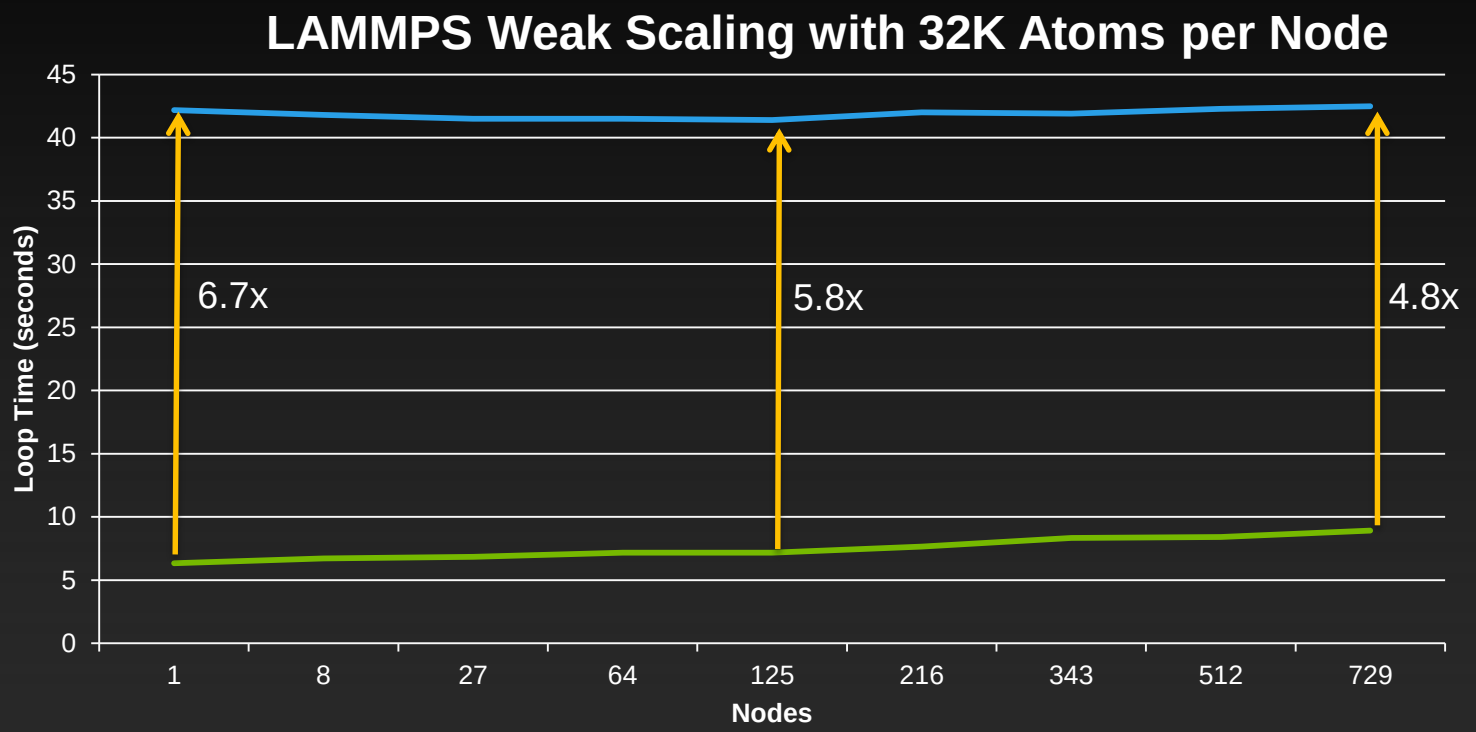

Performance of 4.8x-6.7x with GPU-accelerated nodes when compared to CPUs alone

Each blue Cray XE6 Node have 2x AMD Opteron CPUs (16 Cores per CPU)
Each green Cray XK6 Node has 1x AMD Opteron 1600 CPU (16 Core per CPU) and 1x NVIDIA X2090

Faster, Greener – Worth It! (LAMMPS)

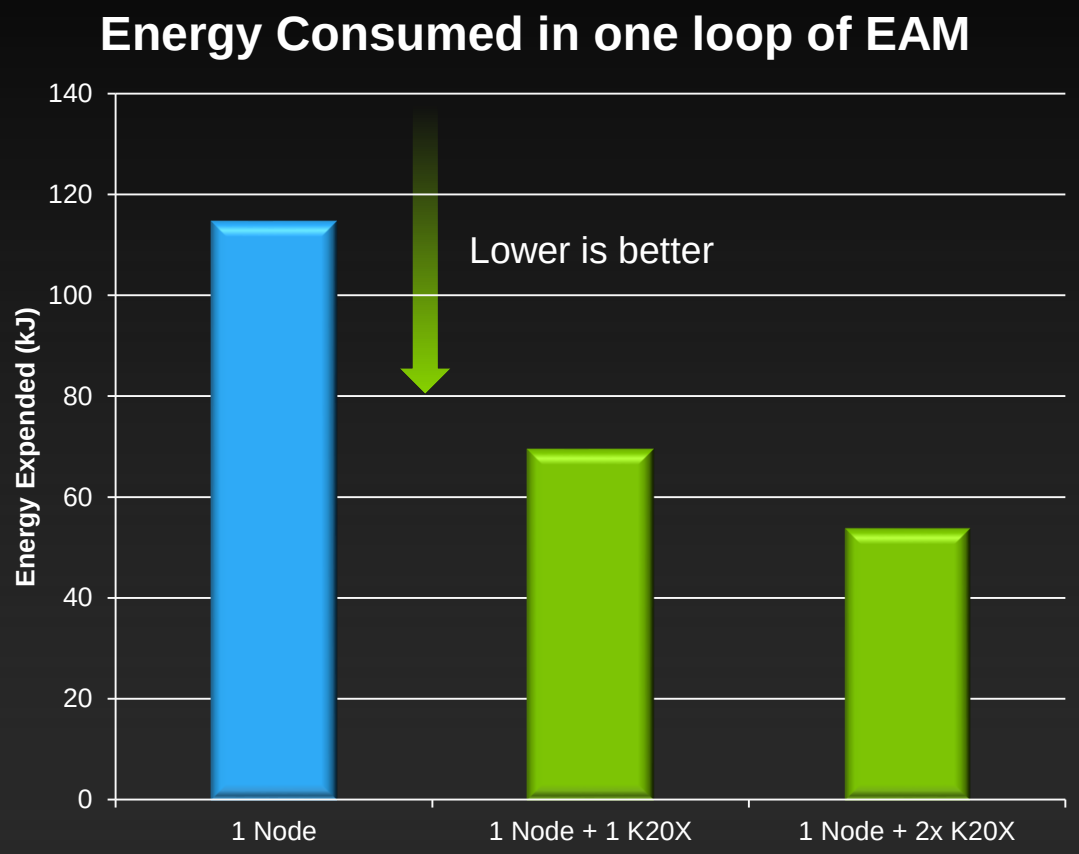

GPU-accelerated computing uses 53% less energy than CPU only

Energy Expended = Power x Time

Power calculated by combining the component's TDPs

Blue node uses 2x E5-2687W (8 Cores and 150W per CPU) and CUDA 4.2.9.
Green nodes have 2x E5-2687W and 1 or 2 NVIDIA K20X GPUs (235W) running CUDA 5.0.35.

Try GPU accelerated LAMMPS for free – www.nvidia.com/GPUTestDrive

Accelerate LAMMPS Simulations with GPUs

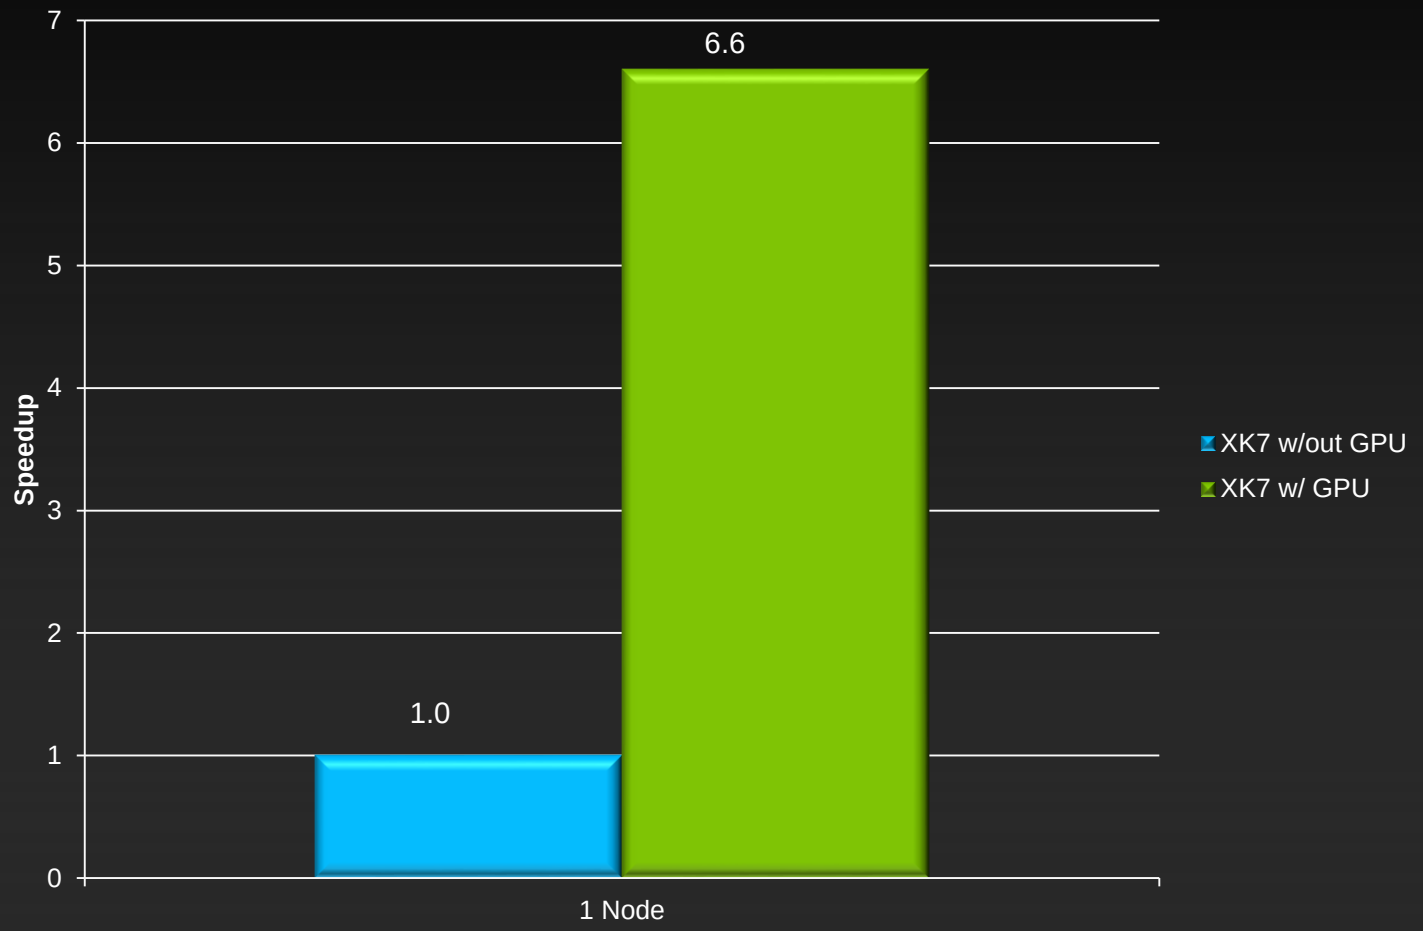

"Summary of best speedups versus running on a single XK7 CPU for CPU-only and accelerated runs. Simulation is 400 timesteps for a 1 million molecule droplet. The speedups are calculated based on the single node loop time of 440.3 seconds."

*Brown, W.M., Yamada, M., "Implementing Molecule Dynamics on Hybrid High Performance Computers – Three-Body Potentials," Computer Physics Communications (2013, submitted)

Accelerate LAMMPS Simulations with GPUs

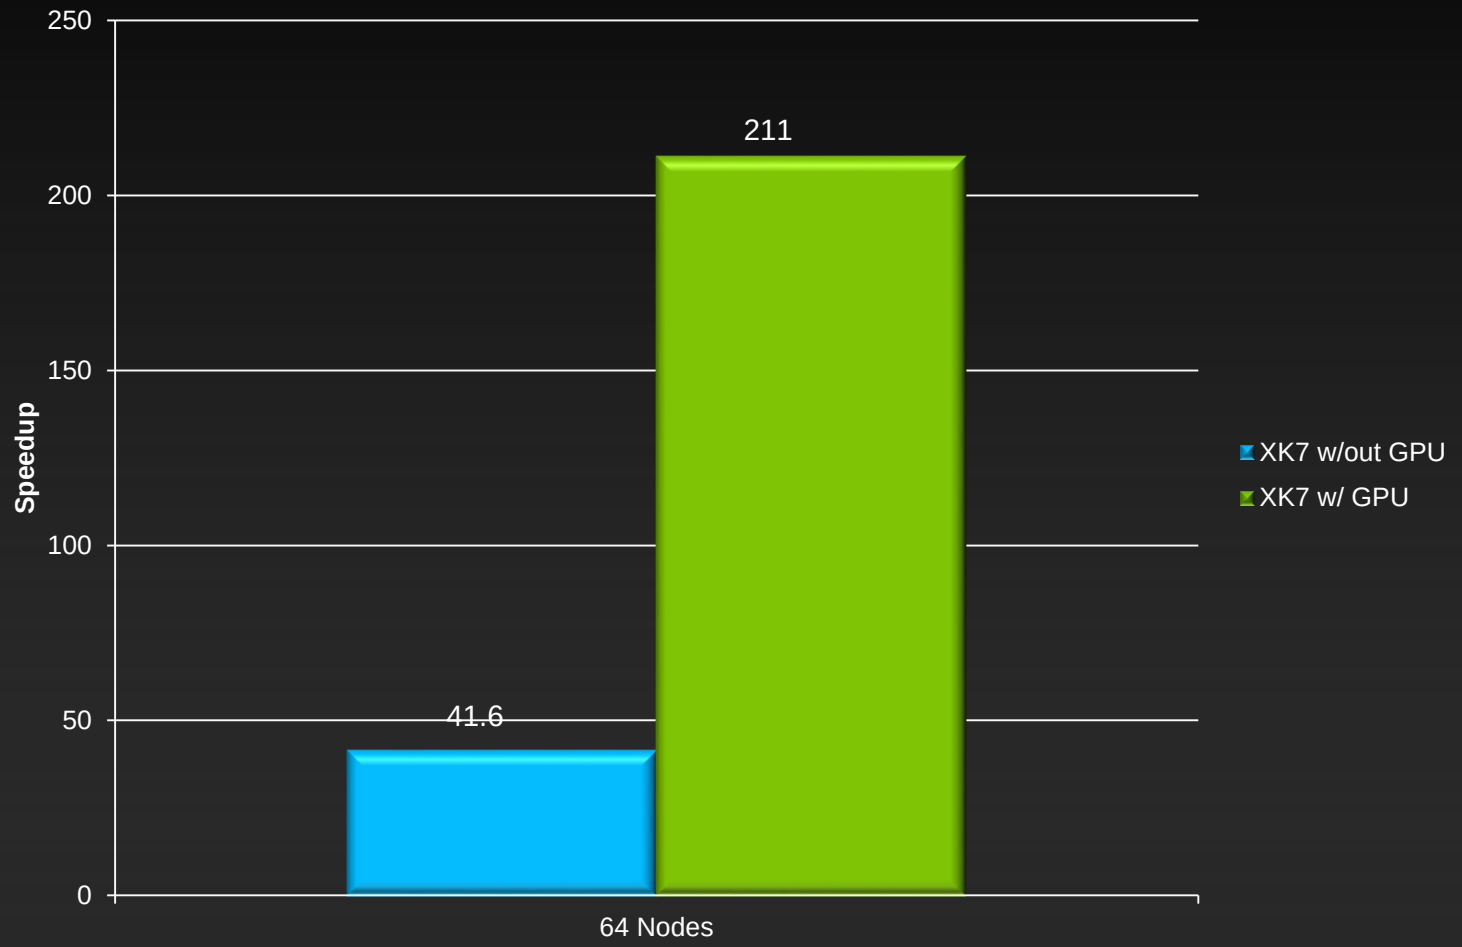

"Summary of best speedups versus running on a single XK7 CPU for CPU-only and accelerated runs. Simulation is 400 timesteps for a 1 million molecule droplet. The speedups are calculated based on the single node loop time of 440.3 seconds."

*Brown, W.M., Yamada, M., "Implementing Molecule Dynamics on Hybrid High Performance Computers – Three-Body Potentials," Computer Physics Communications (2013, submitted)

Recommended GPU Node Configuration for LAMMPS Computational Chemistry

Workstation or Single Node Configuration	
# of CPU sockets	2
Cores per CPU socket	6+
CPU speed (Ghz)	2.66+
System memory per socket (GB)	32
GPUs	GTX Titan X, Kepler K20, K40, K80, M40
# of GPUs per CPU socket	1-2
GPU memory preference (GB)	6+
GPU to CPU connection	PCIe 3.0 or higher
Server storage	500 GB or higher
Network configuration	Gemini, InfiniBand

Scale to thousands of nodes with same single node configuration

NAMD 2.11 – Up to 2X Faster

New GPU features in NAMD 2.11

Selected Text from the NAMD website

- **GPU-accelerated simulations up to twice as fast as NAMD 2.10**
- **Pressure calculation with fixed atoms on GPU works as on CPU**
- **Improved scaling for GPU-accelerated particle-mesh Ewald calculation**
  - CPU-side operations overlap better and are parallelized across cores.
- **Improved scaling for GPU-accelerated simulations**
  - Nonbonded force calculation results are streamed from the GPU for better overlap.
- **NVIDIA CUDA GPU-acceleration binaries for Mac OS X**

NAMD 2.11 is up to 2x faster

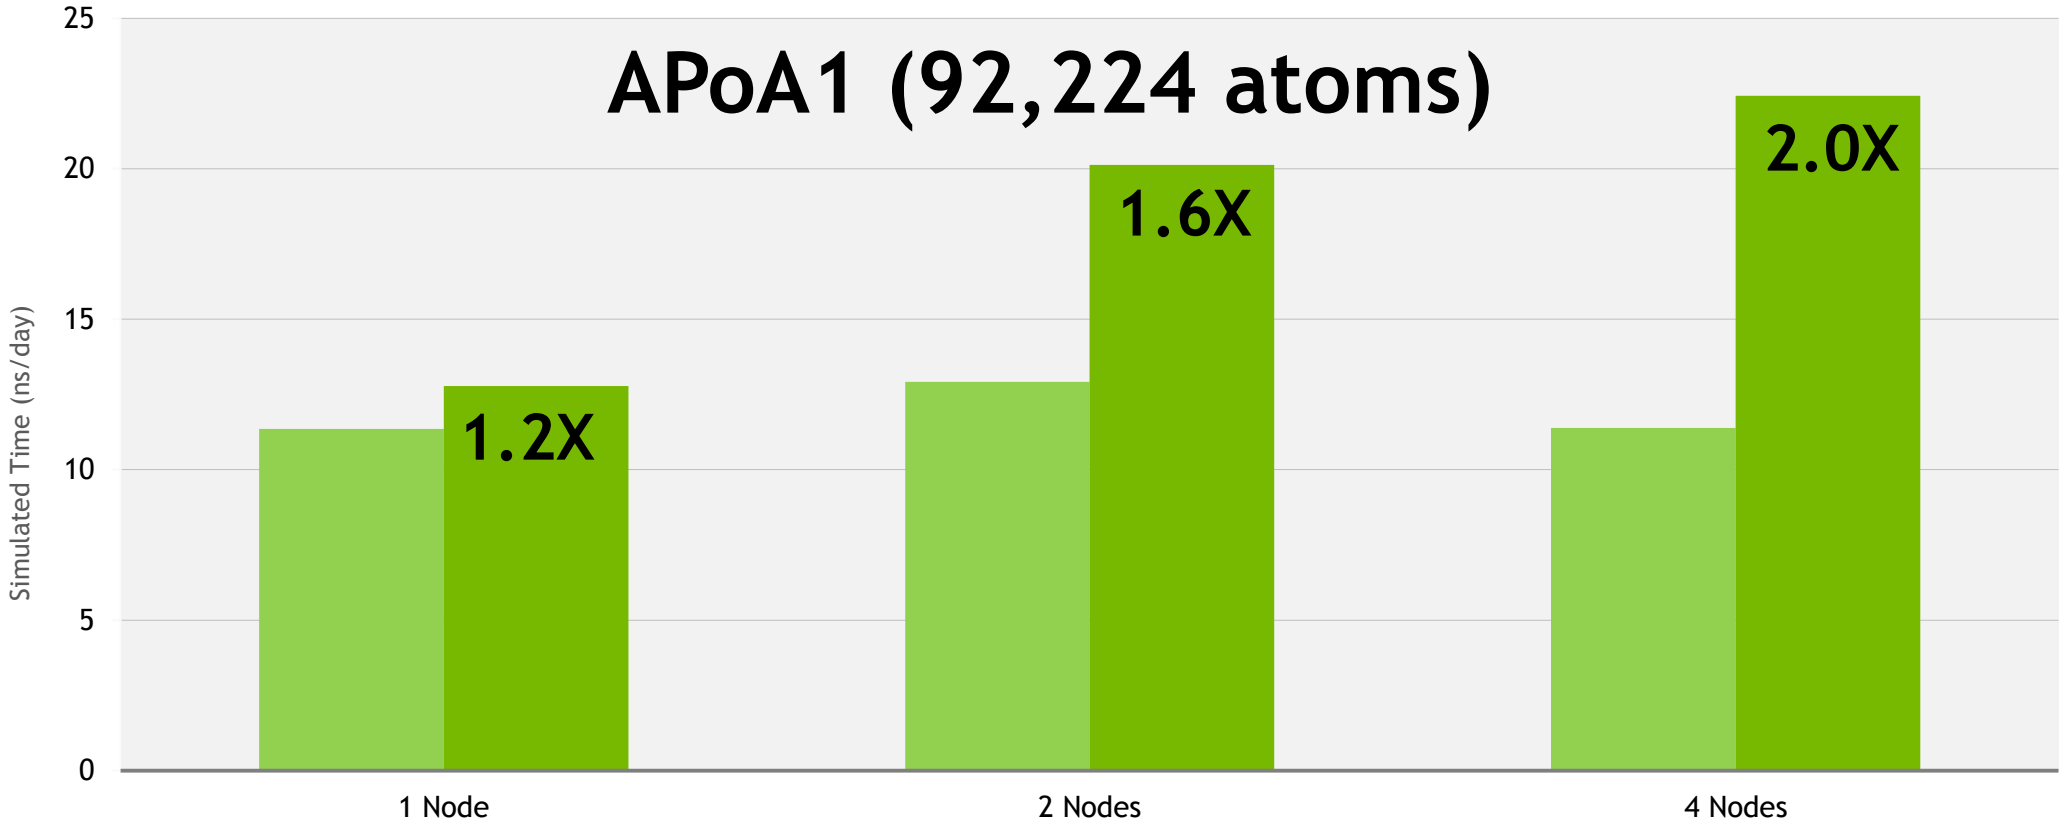

NAMD 2.10 & NAMD 2.11 contain Dual Intel E5-2697 v2@2.7GHz (IvyBridge) CPUs + 2 Tesla K80 (autoboost) GPUs

NAMD 2.11 APoA1 on 1 and 2 nodes

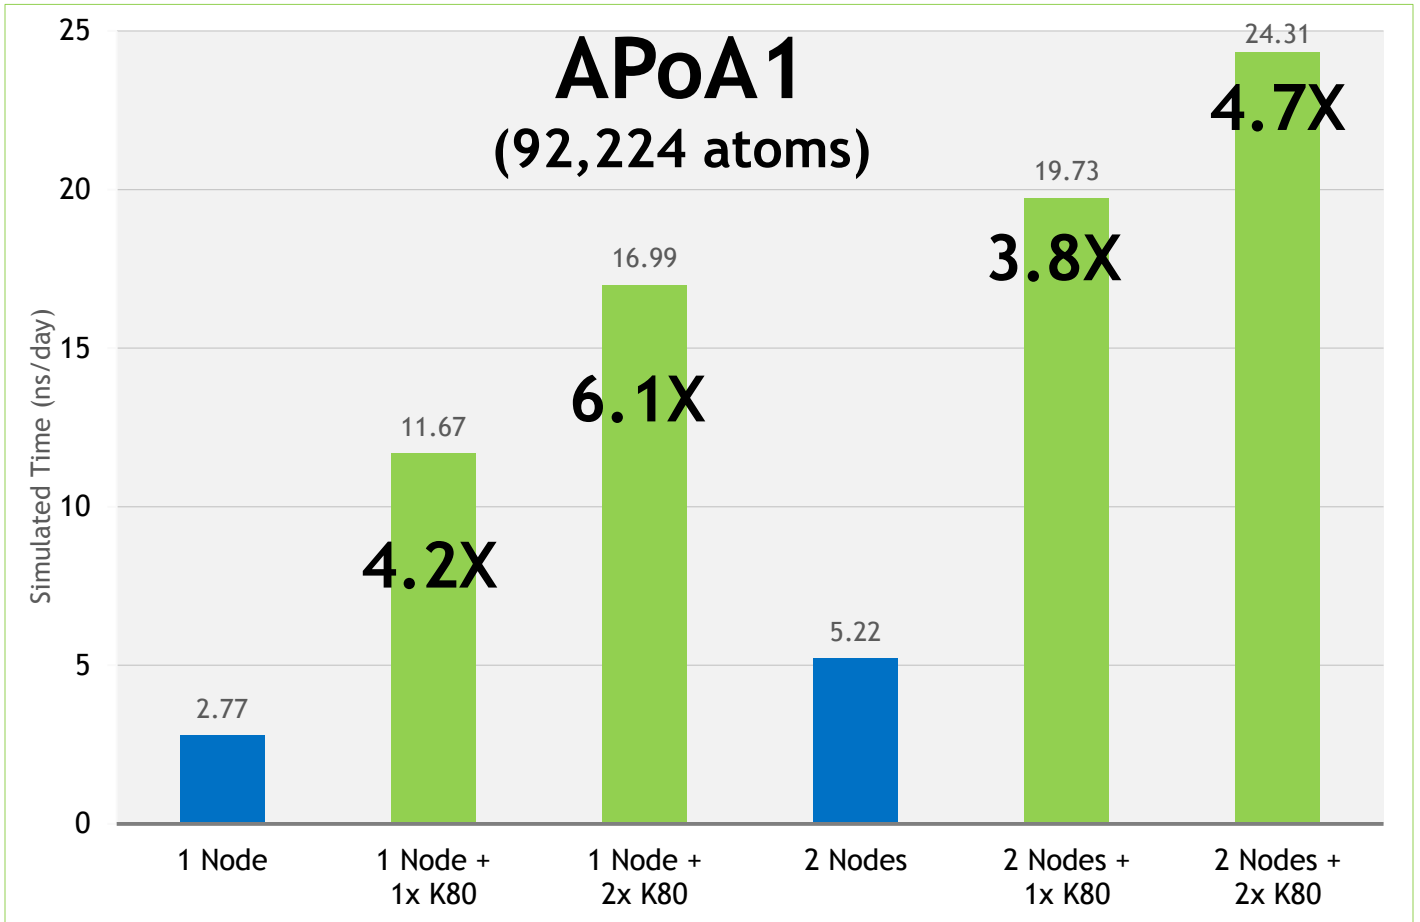

Running NAMD version 2.11

The blue nodes contain Dual Intel E5-2698 v3@2.3GHz (Haswell) CPUs

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

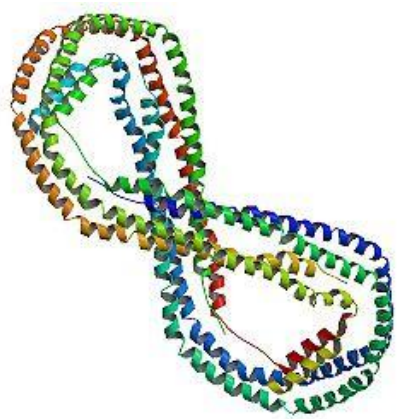

NAMD 2.11 APoA1 on 4 and 8 nodes

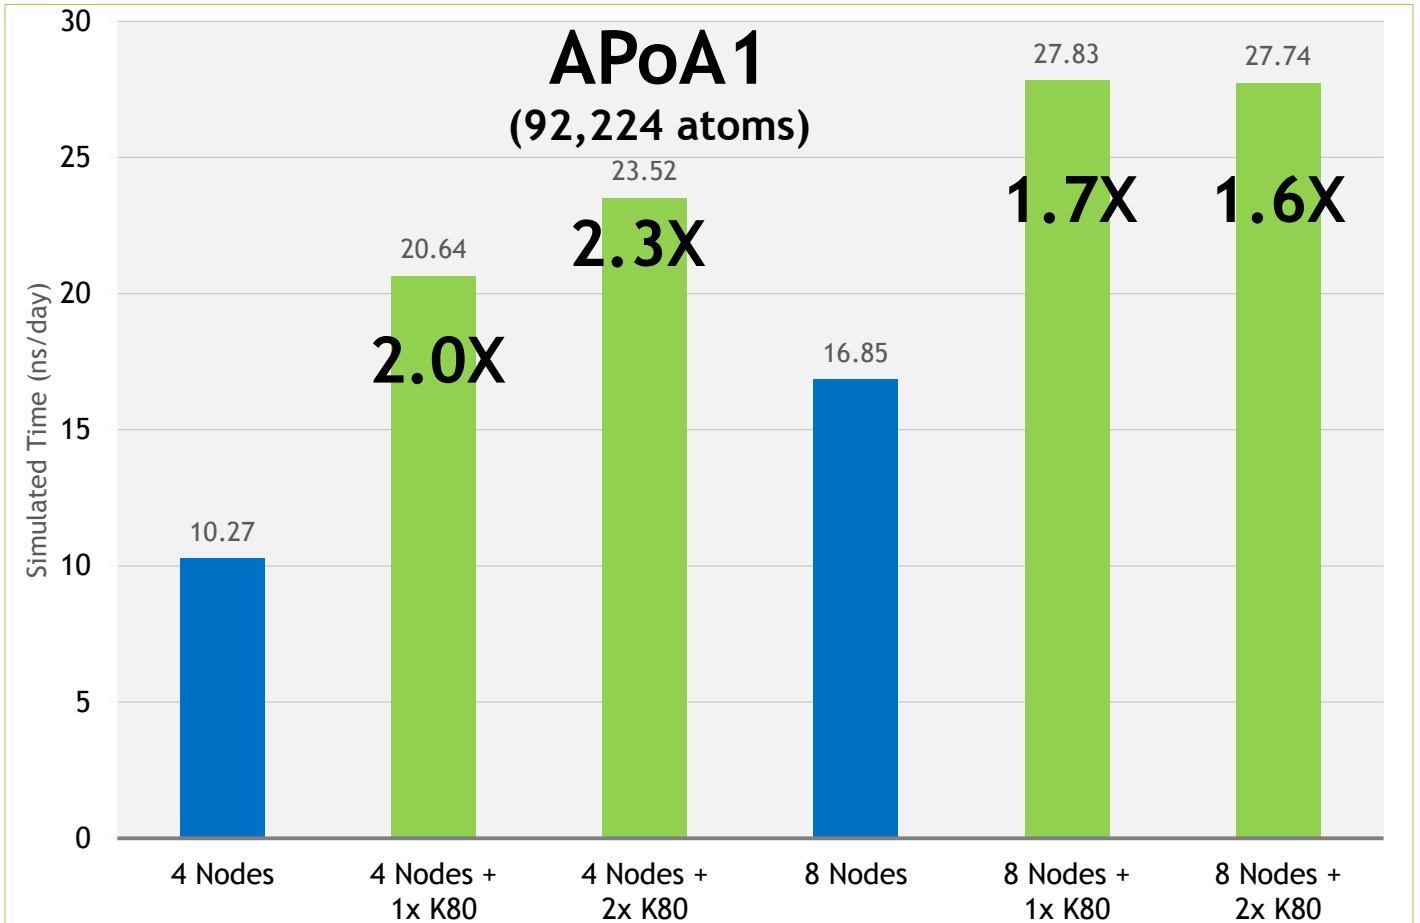

Running NAMD version 2.11

The blue nodes contain Dual Intel E5-2698 v3@2.3GHz (Haswell) CPUs

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

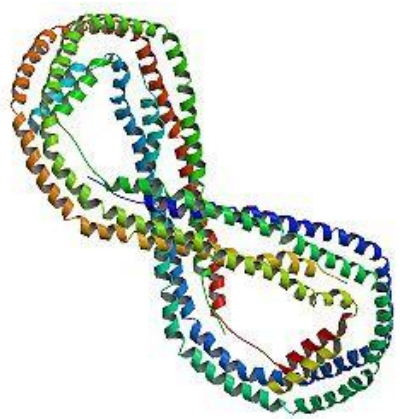

NAMD 2.11 is up to 1.8x faster

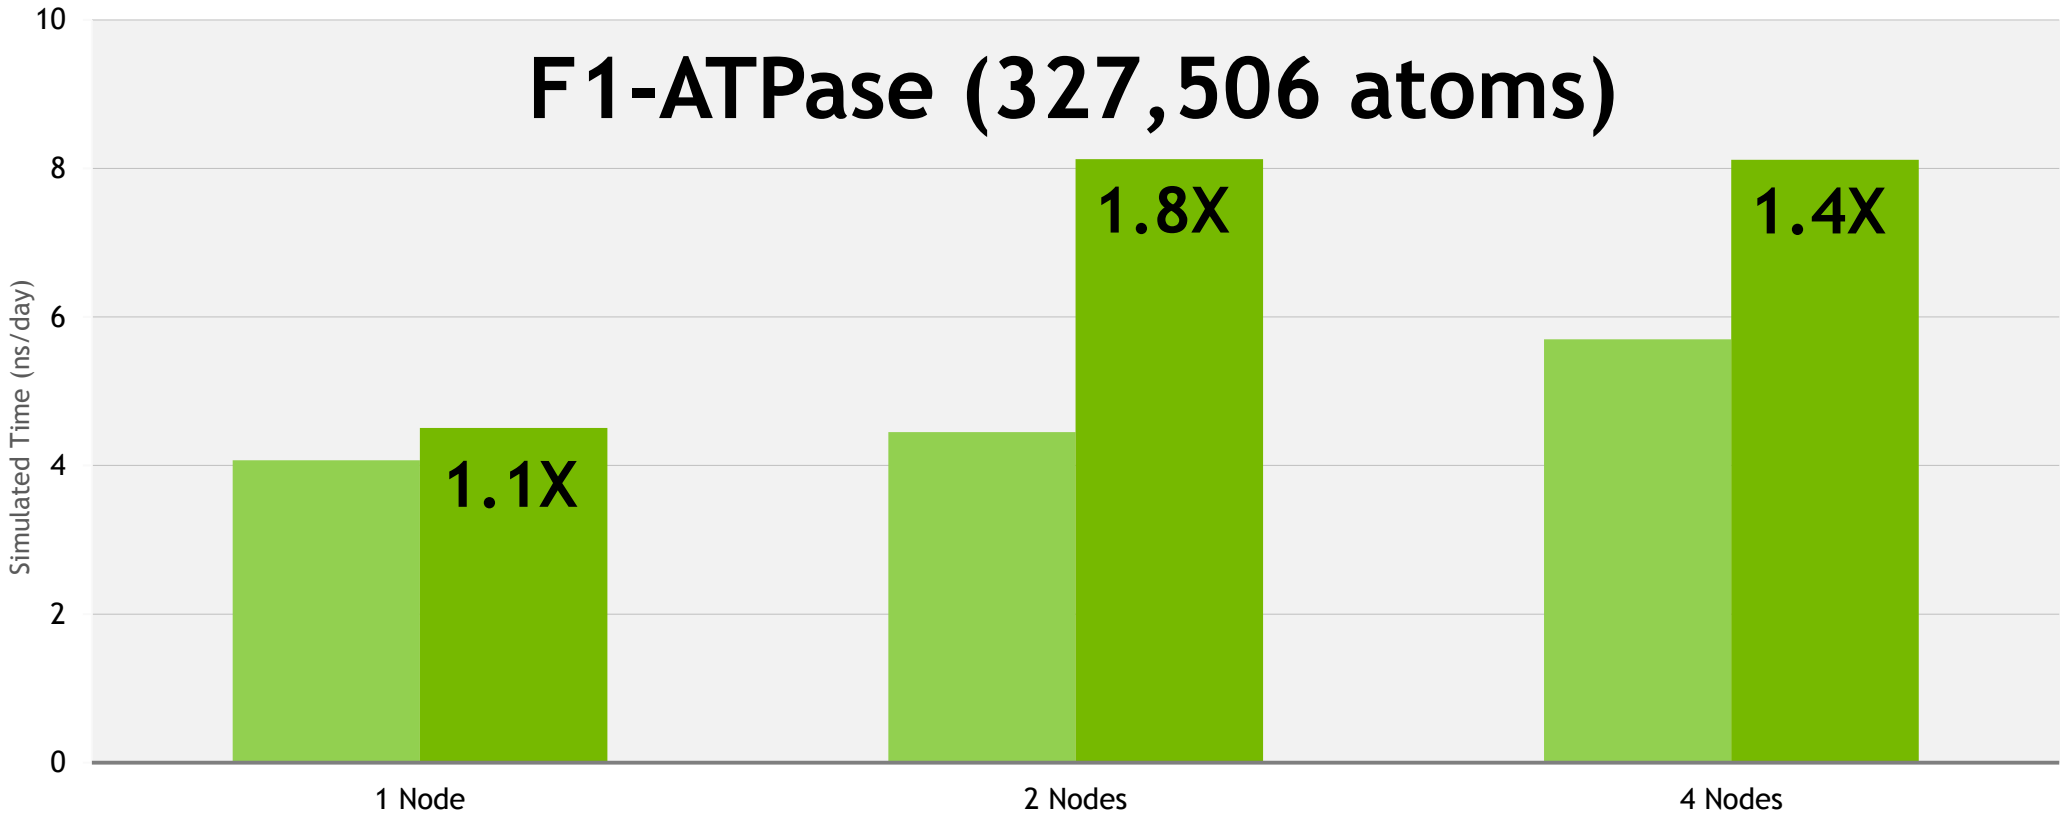

NAMD 2.10 & NAMD 2.11 contain Dual Intel E5-2697 v2@2.7GHz (IvyBridge) CPUs + 2 Tesla K80 (autoboost) GPUs

NAMD 2.11 F1-ATPase on 1 and 2 nodes

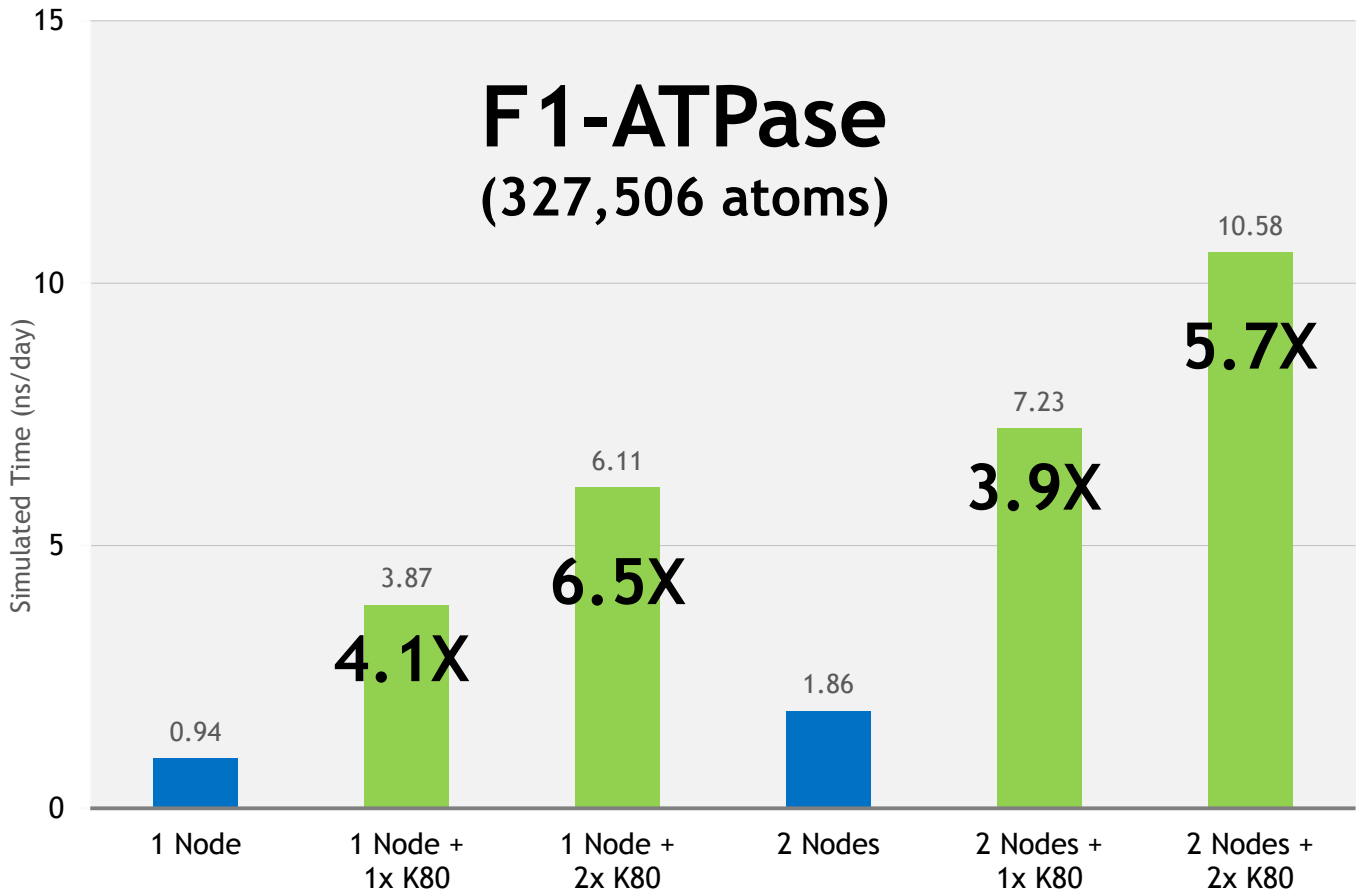

Running NAMD version 2.11

The blue nodes contain Dual Intel E5-2698 v3@2.3GHz (Haswell) CPUs

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

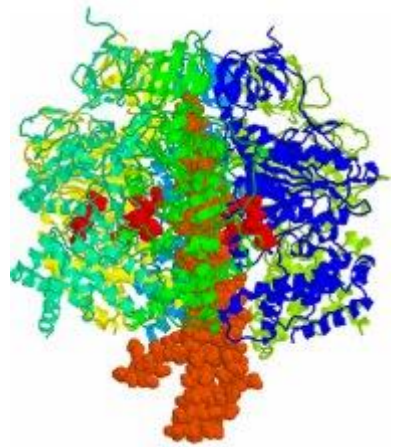

NAMD 2.11 F1-ATPase on 4 and 8 nodes

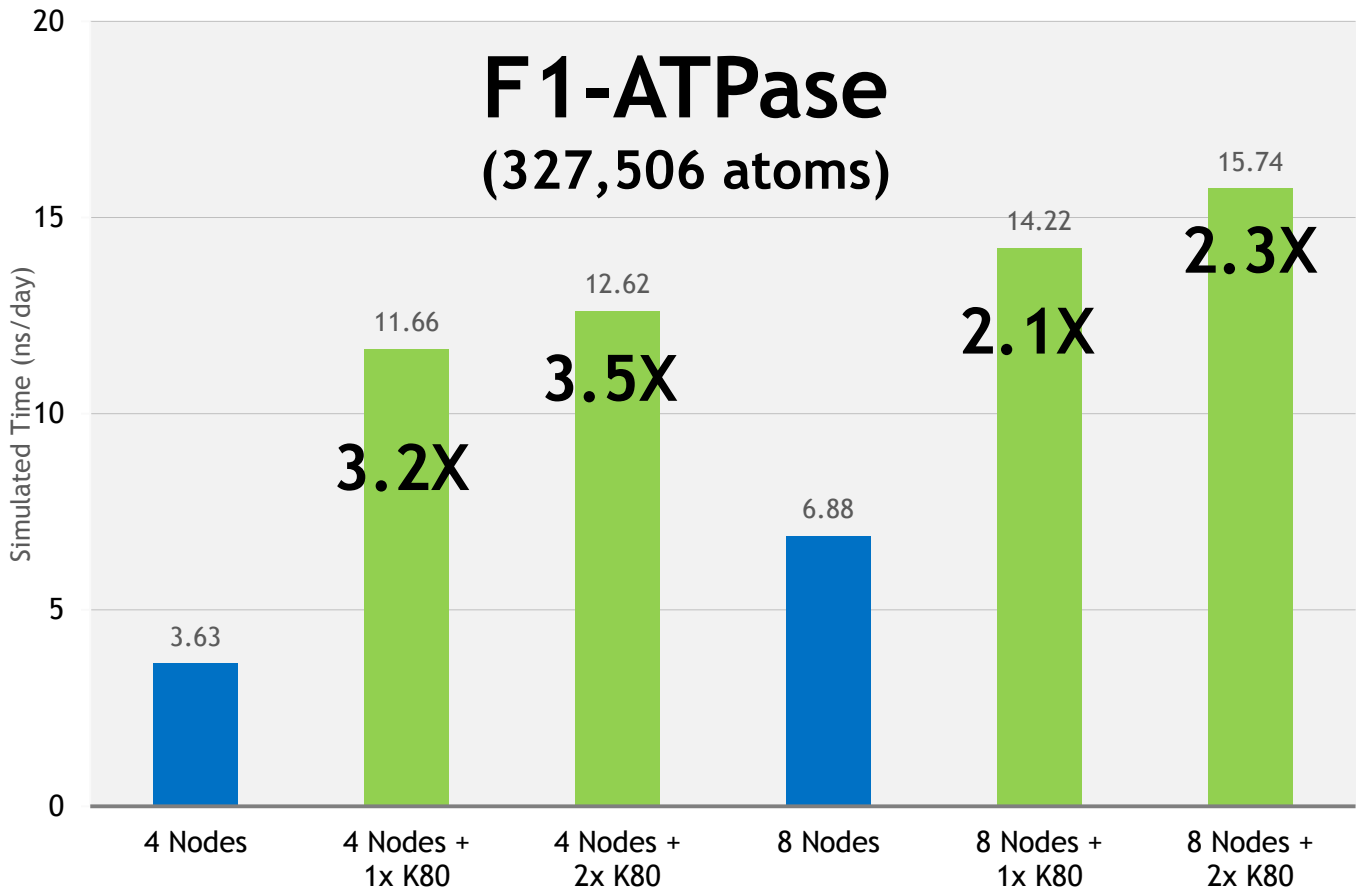

Running NAMD version 2.11

The blue nodes contain Dual Intel E5-2698 v3@2.3GHz (Haswell) CPUs

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

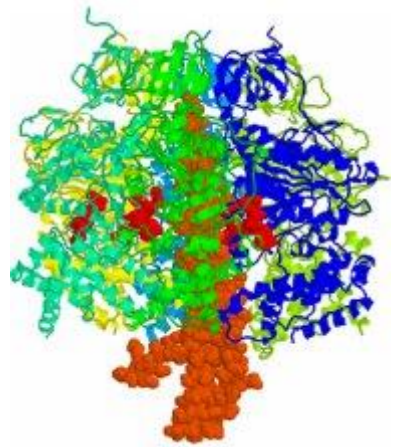

NAMD 2.11 is up to 1.5x faster

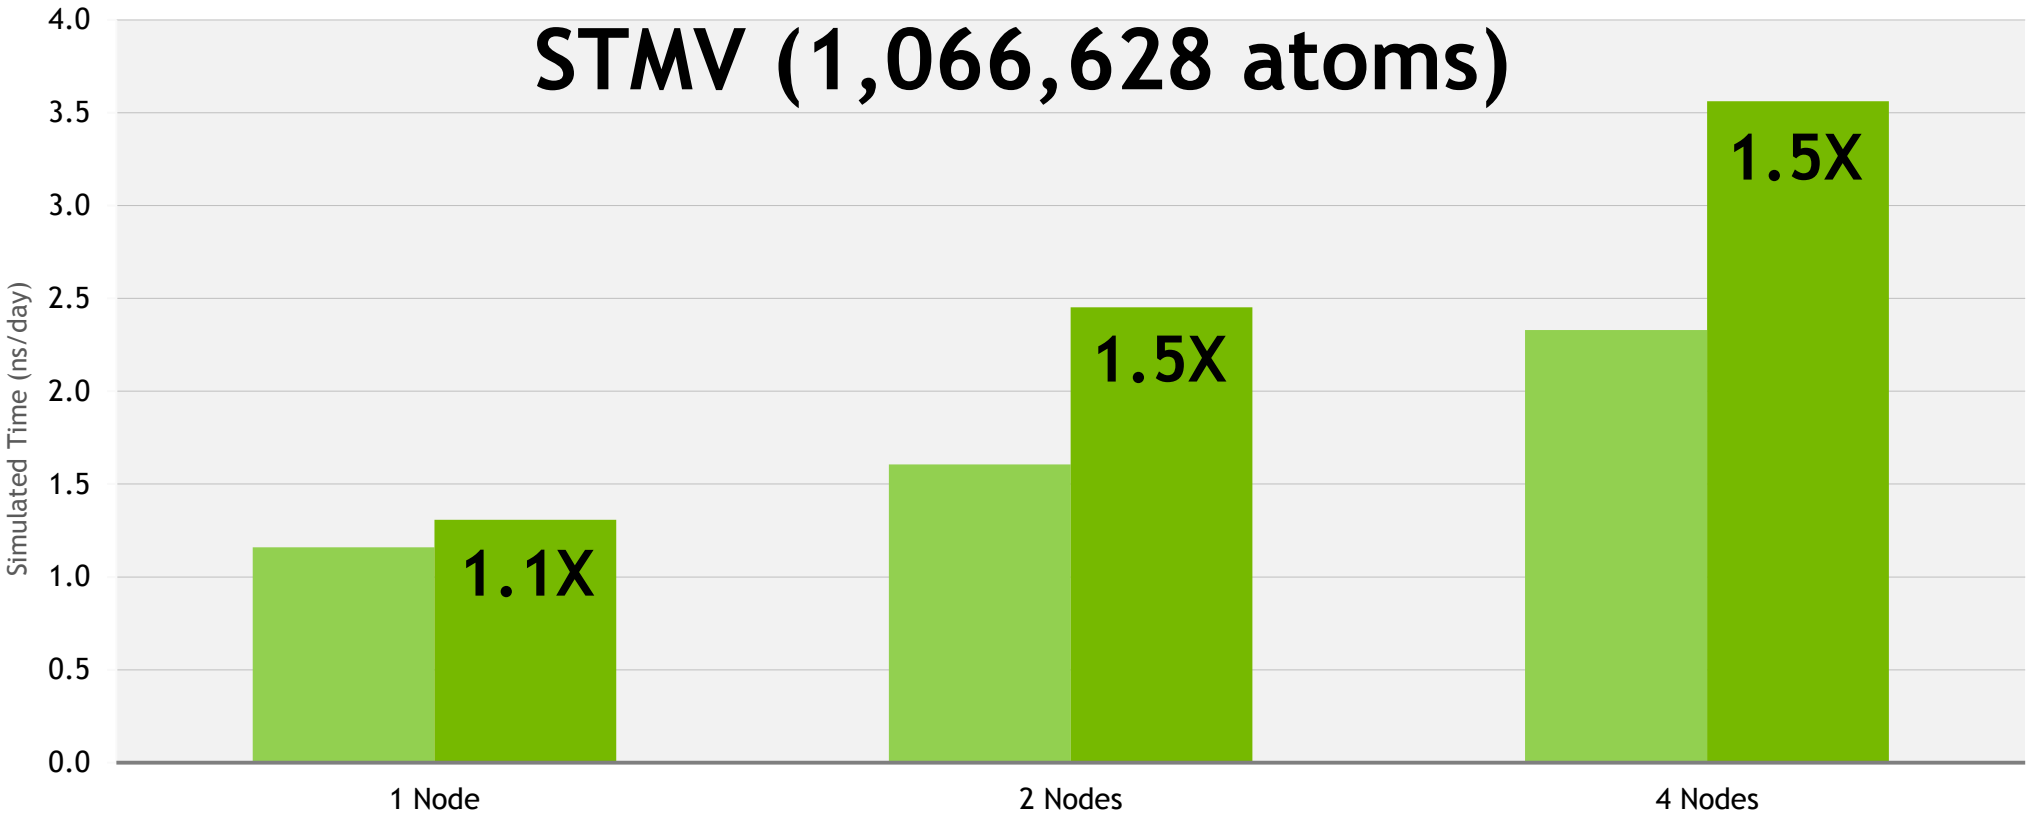

NAMD 2.10 & NAMD 2.11 contain Dual Intel E5-2697 v2@2.7GHz (IvyBridge) CPUs + 2 Tesla K80 (autoboost) GPUs

NAMD 2.11 STMV on 1 and 2 nodes

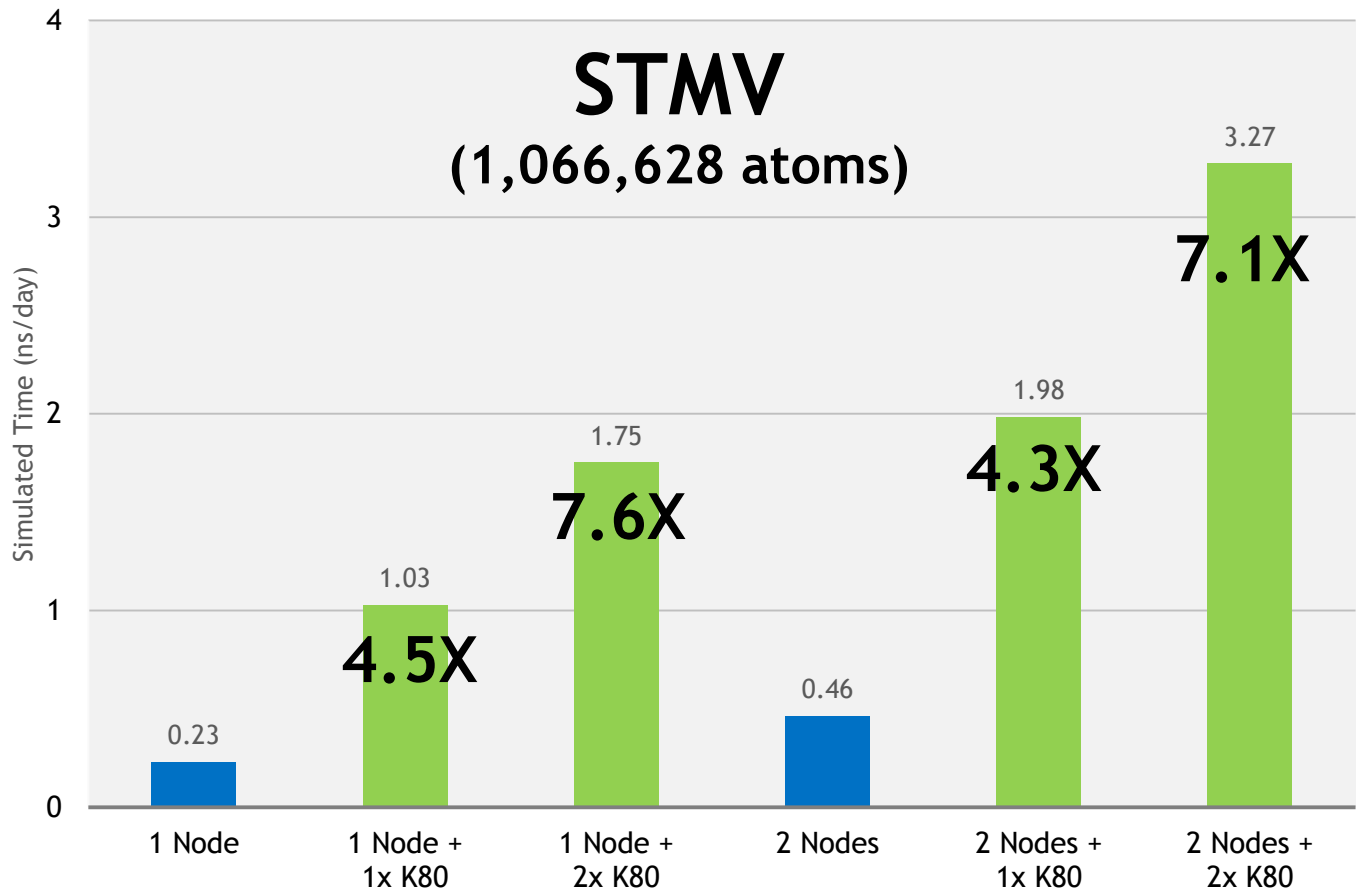

Running NAMD version 2.11

The blue nodes contain Dual Intel E5-2698 v3@2.3GHz (Haswell) CPUs

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

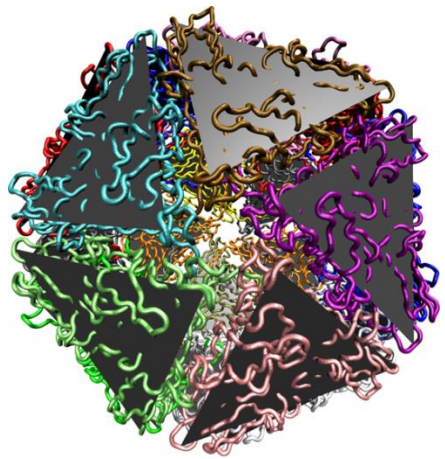

NAMD 2.11 STMV on 4 and 8 nodes

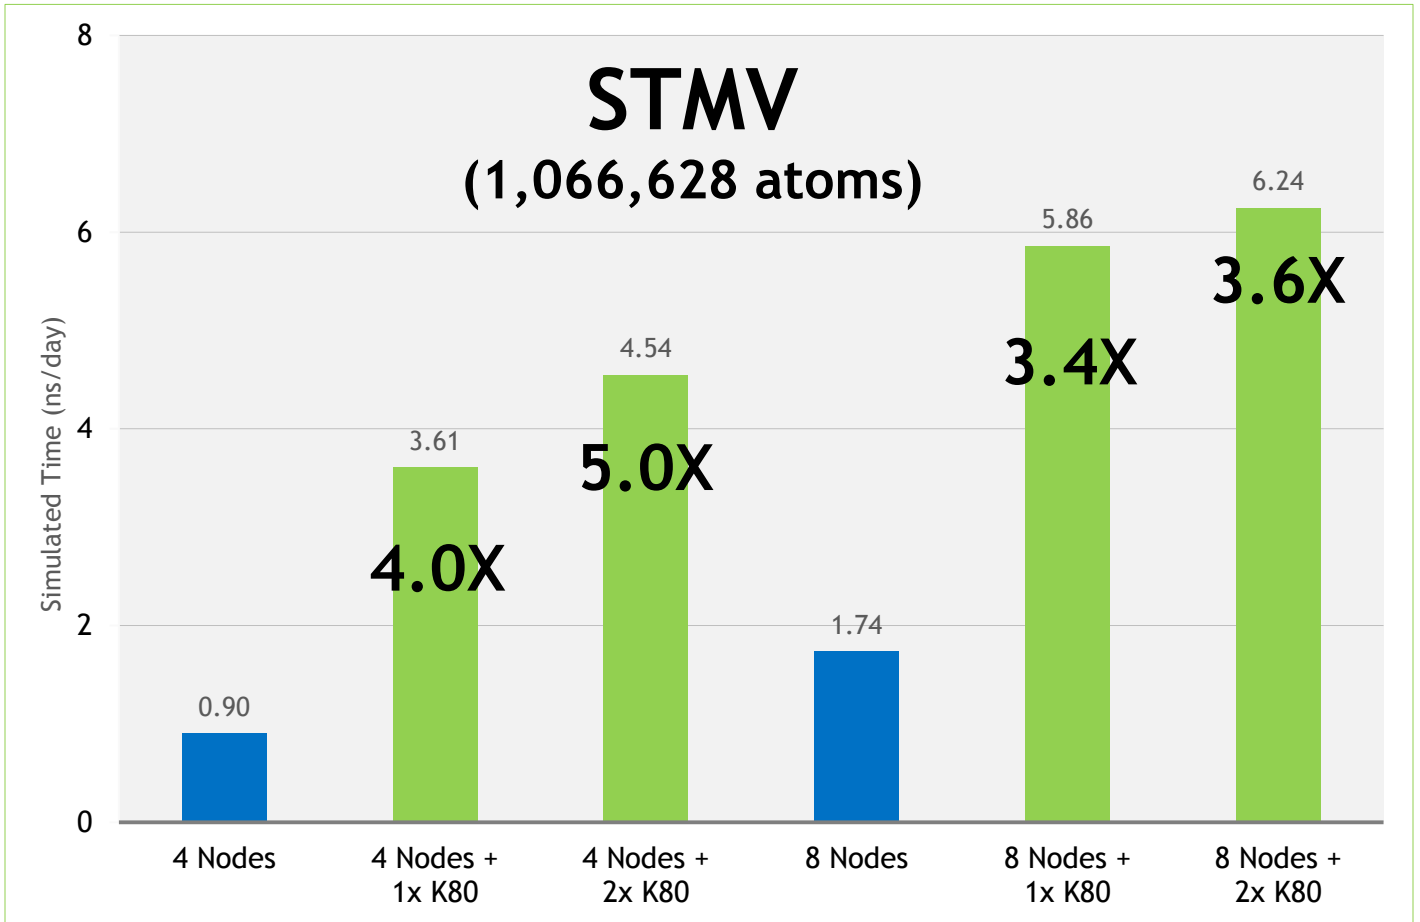

Running NAMD version 2.11

The blue nodes contain Dual Intel E5-2698 v3@2.3GHz (Haswell) CPUs

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

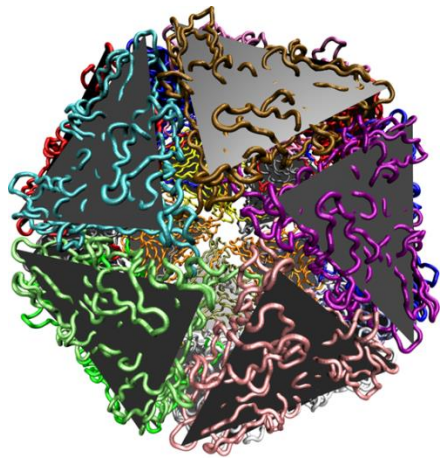

Recommended GPU Node Configuration for NAMD Computational Chemistry

Workstation or Single Node Configuration	
# of CPU sockets	2
Cores per CPU socket	6+
CPU speed (Ghz)	2.66+
System memory per socket (GB)	32
GPUs	Kepler K20, K40, K80
# of GPUs per CPU socket	1-2
GPU memory preference (GB)	6-12
GPU to CPU connection	PCIe 3.0 or higher
Server storage	500 GB or higher
Network configuration	Gemini, InfiniBand

Scale to thousands of nodes with same single node configuration

Benefits of MD GPU-Accelerated Computing

Why wouldn't you want to turbocharge your research?

- 3x-8x Faster than CPU only systems in all tests (on average)
- Most major compute intensive aspects of classical MD ported
- Large performance boost with marginal price increase
- Energy usage cut by more than half
- GPUs scale well within a node and/or over multiple nodes
- K80 GPU is our fastest and lowest power high performance GPU yet

Try GPU accelerated MD apps for free – www.nvidia.com/GPUTestDrive

Molecular Dynamics (MD) on GPUs

February 11, 2016