




LIFE AND MATERIAL SCIENCE APPLICATION PERFORMANCE



May 2020



Critical COVID-19 Research Accelerated on NVIDIA GPUs


NVIDIA GPUs accelerate today's most critical global health research efforts. In just under two weeks, researchers determined the first high resolution structure of the 2019-nCoV spike protein using Structura Bio application, cryoSPARC. The turnaround time came by ongoing improvements in computational throughput and more accurate methods. Over the past 5 years, today's cryo-EM pipelines take days rather than months.

From real-time processing as data is acquired to advanced molecular simulations using HPC and AI, future vaccines and therapies will be ready to fight global health threats.



END TO END GPU ACCELERATED CRYO-EM

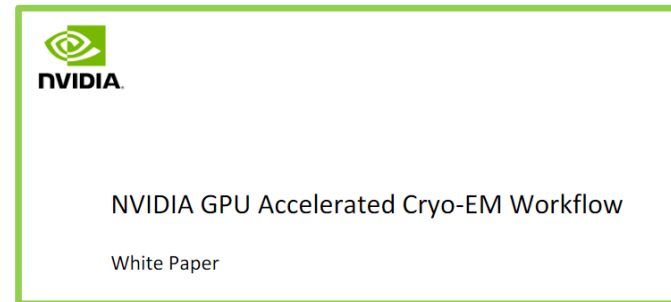
Single Compute Platform for AI, HPC, Visualization




GPU-accelerated apps

- AMIRA
- BioEM
- cryoSPARC
- cyYOLO
- Dynamo
- EMAN2
- emClarity
- GCTF
- IMOD
- MotionCor2
- RELION
- Tomviz
- Topaz
- VMD
- Warp

Cryo-EM Performance White Paper



Overview of Life & Material Accelerated Apps




MD

All key codes are GPU-accelerated

Great multi-GPU, multi-node (dense) performance

GPU-accelerated apps

ACEMD, AMBER, 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

GPU-accelerated math libraries, OpenACC directives

GPU-accelerated apps

ABINIT, ACES III, ADF, BigDFT, CP2K, GAMESS, GAMESS-UK, GPAW, LATTE, LSDalton, LSMS,
MOLCAS, MOPAC2012, NWChem, NWChemEx, OCTOPUS, PEtot, QUICK, Q-Chem,
QMCPack, Quantum Espresso, QUICK, TeraChem, VASP


Active acceleration projects

CASTEP, GAMESS, Gaussian, ONETEP, Quantum Supercharger Library, & more




AMBERMD 2020

May 2020




AmberMD 2020

Cellulose
408,609 atoms




Performance Comparison - Cellulose 4fs




AmberMD 2020

Factor IX
90,906 atoms




Performance Comparison - FactorIX 4fs




AmberMD 2020

DHFR 23,558
atoms




Performance Comparison - JAC 4fs




AmberMD 2020

Performance Comparison - STMV 4fs



AmberMD

Benefit of NVSWITCH and NCCL - Average benchmark performance




AmberMD recommended usage

Motherboard and CPU	Dual-socket CPU with server
System memory	$\geq 16\text{GB}$
GPUs	V100 or A100 NVSWITCH
GPUs per socket	4 or 8
GPUs per task	1 - 4



GROMACS 2020

May 2020




GROMACS - 2020



GROMACS - 2020

Performance Comparison




GROMACS recommended usage

Motherboard and CPU	Dual-socket CPU server
System memory	$\geq 16\text{GB}$
GPUs	V100 / A100 with SXM2/3/4
GPUs per socket	1 to 4
GPUs per task	1 - 4




LAMMPS PATCH_20NOV2019


May 2020



LAMMPS - patch_20Nov2019




Performance Comparison - Atomic-Fluid Lennard-Jones 2.5 Cutoff




LAMMPS - patch_20Nov2019

Bulk Cu lattice

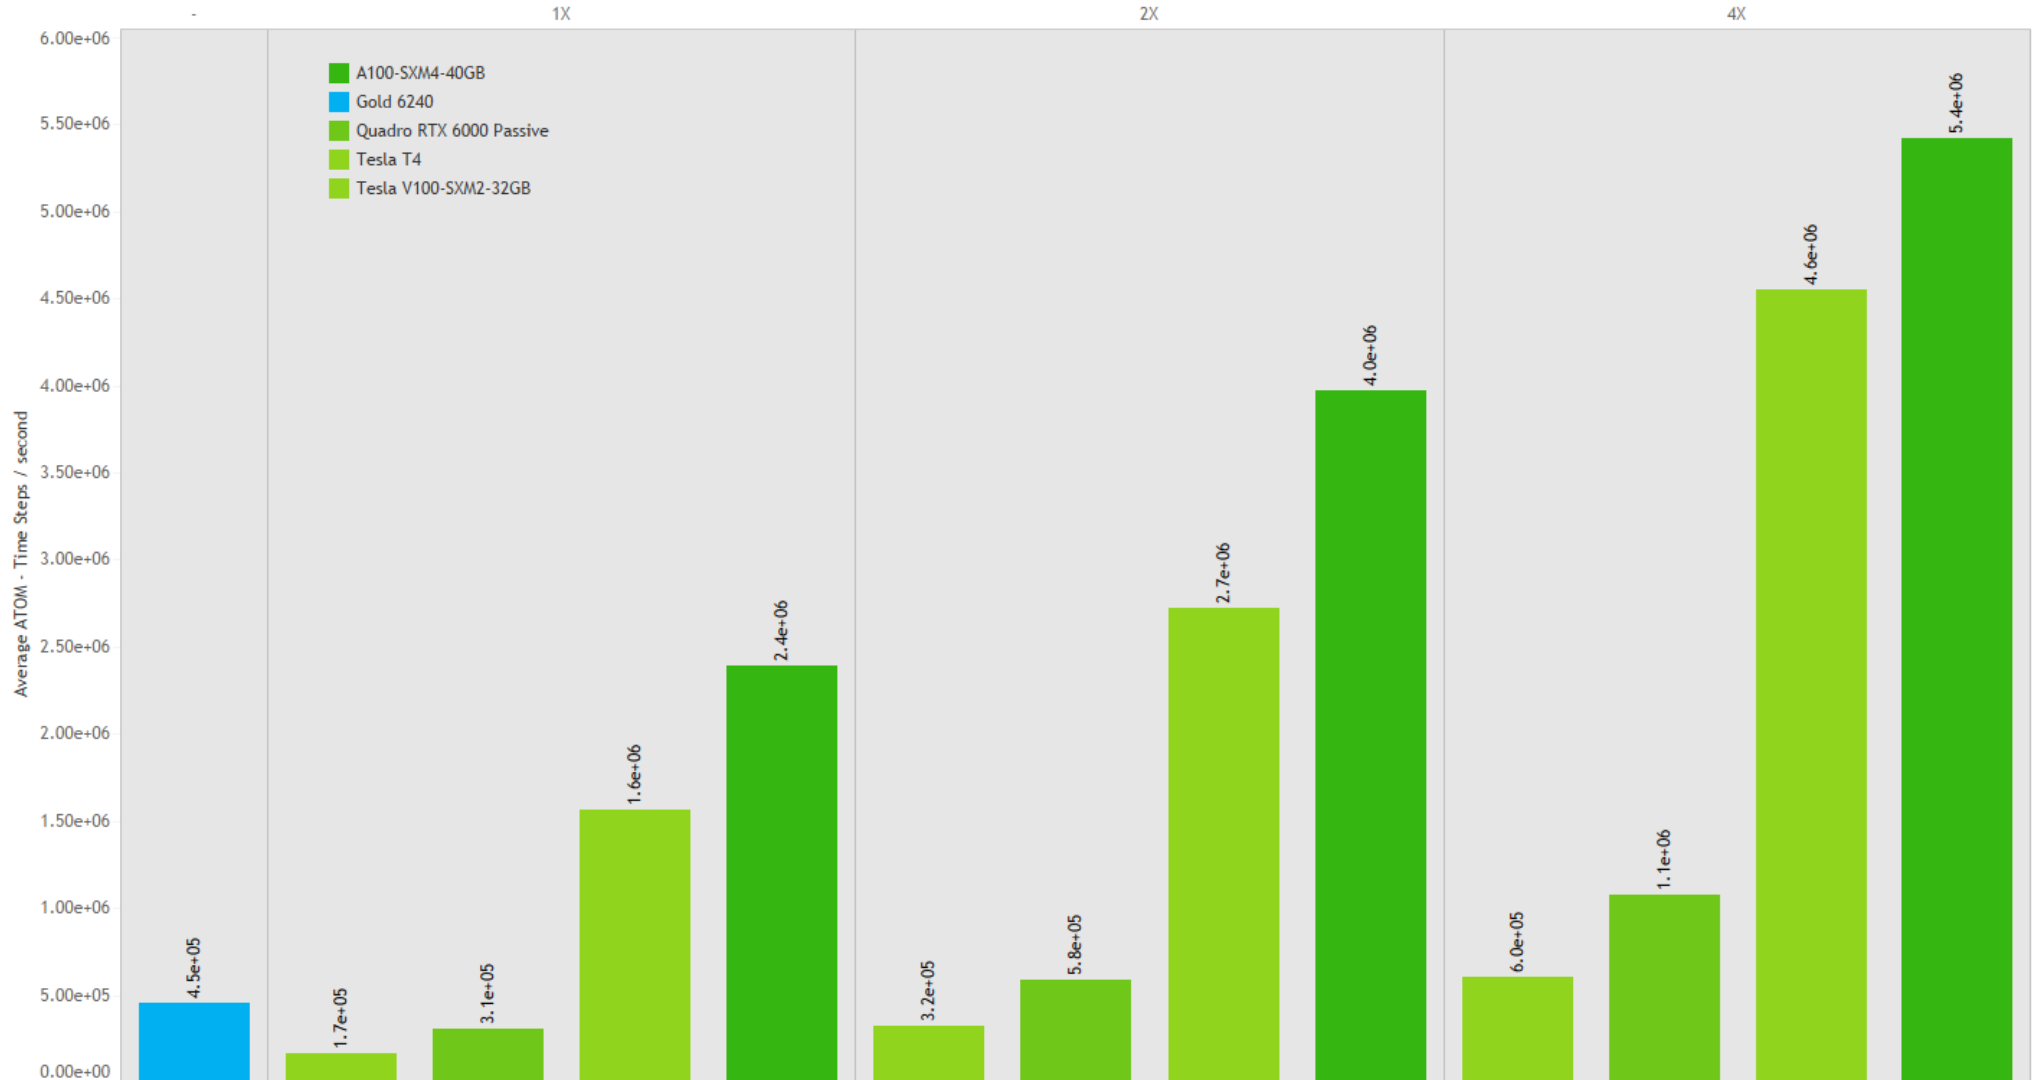


Performance Comparison - EAM




LAMMPS - patch_20Nov2019

Performance Comparison - ReaxFF/C




LAMMPS - patch_20Nov2019

Performance Comparison - Rhodopsin




LAMMPS - patch_20Nov2019

Si crystallization




Performance Comparison - Tersoff



LAMMPS - patch_20Nov2019

Performance Comparison - SNAP




LAMMPS recommended usage

Motherboard and CPU	Dual-socket CPU
System memory	$\geq 32\text{GB}$
GPUs	V100 / A100
GPUs per socket	1 to 4
GPUs per task	4




NAMD 3.0A

May 2020




NAMD - 3.0a1

Performance Comparison - Average Performance




NAMD - 3.0a1

ApoA1
92,224 atoms




Performance Comparison of APOA1 benchmarks




NAMD - 3.0a1

STMV
1,067,095 atoms



Performance Comparison of STMV benchmarks



NAMD recommended usage


v3.0a has strong single GPU performance

Motherboard and CPU	Dual-socket CPU server
System memory	$\geq 16\text{GB}$
GPUs	V100 / A100
GPUs per socket	1 to 4
GPUs per task	1




RELION 3.0.8

May 2020



RELION - 3.0.8

Performance Comparison - 3D Refinement



RELION RECOMMENDED USAGE

Motherboard and CPU	Dual-socket CPU
System memory	$\geq 32\text{GB}$
GPUs	RTX 6000 / RTX 8000 / V100 / A100
GPUs per socket	1 to 4
GPUs per task	1 to 4

