Sandia National Laboratories

Exceptional service in the national interest

LAMMPS-KOKKOS: Performance portable molecular dynamics at exascale

Anders Johansson

Evan Weinberg, Christian Trott, Meg McCarthy, Stan Moore

LAMMPS workshop 8/13/2025

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology and Engineering Solutions of Sandia LLC, a wholly owned subsidiary of Honeywell International Inc. for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525.
SAND NO. 2025-10109C.

U.S. DEPARTMENT OF ENERGY

NNSA National Nuclear Security Administration

MOTIVATION: COMPUTATIONAL POWER

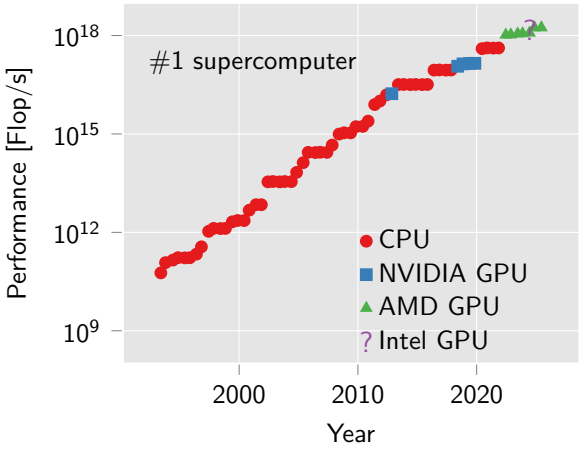

LAMMPS on GPUs:

- 2010: GPU package
  - Pair style on GPU, otherwise CPU
  - Frequent data transfers
  - Should use multiple MPI ranks per GPU
- 2011: USER-CUDA package
  - GPU resident
- 2014: KOKKOS package
  - Based on USER-CUDA package
  - Using Kokkos library
  - Performance *portable*

HOW DOES IT WORK?

Input script:

```
pair_style eam
```

```
class PairEAM
   virtual void compute(...)
       comm->reverse_comm(this);
   }
   int pack_reverse_comm(...)
```

HOW DOES IT WORK?

Input script:

`pair_style eam`

```
class PairEAM
    virtual void compute(...)
        comm->reverse_comm(this);
    }
    int pack_reverse_comm(...)
```

`pair_style eam/kk`

```
class PairEAMKokkos
    void compute(...){
        commKK->reverse_comm(this);
    }
    int pack_reverse_comm_kokkos(...)
```

HOW DOES IT WORK?

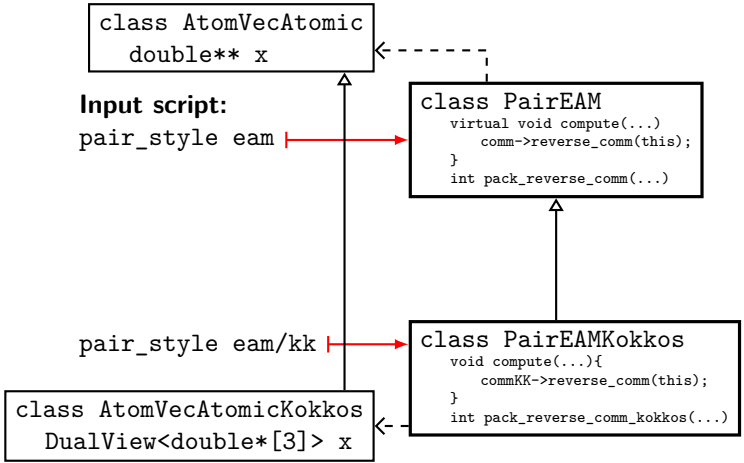

HOW DOES IT WORK?

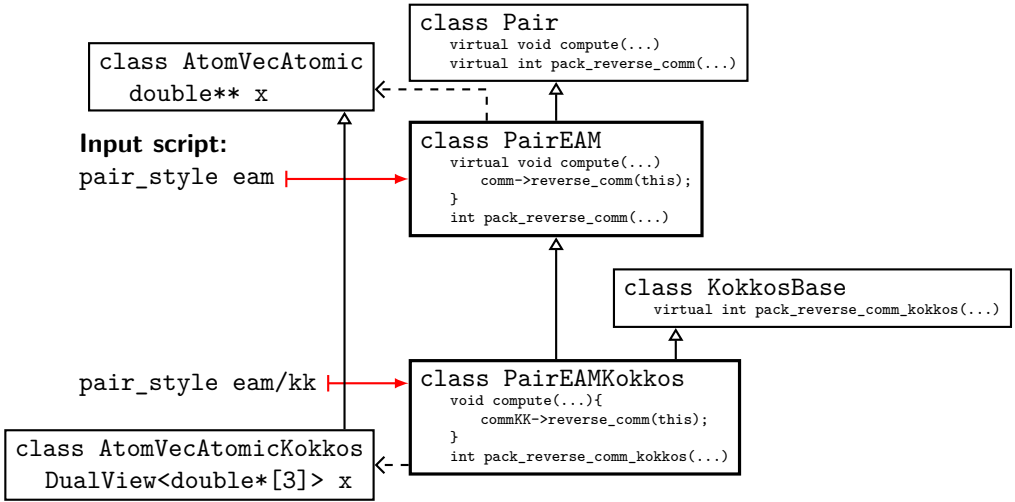

HOW DOES IT WORK?

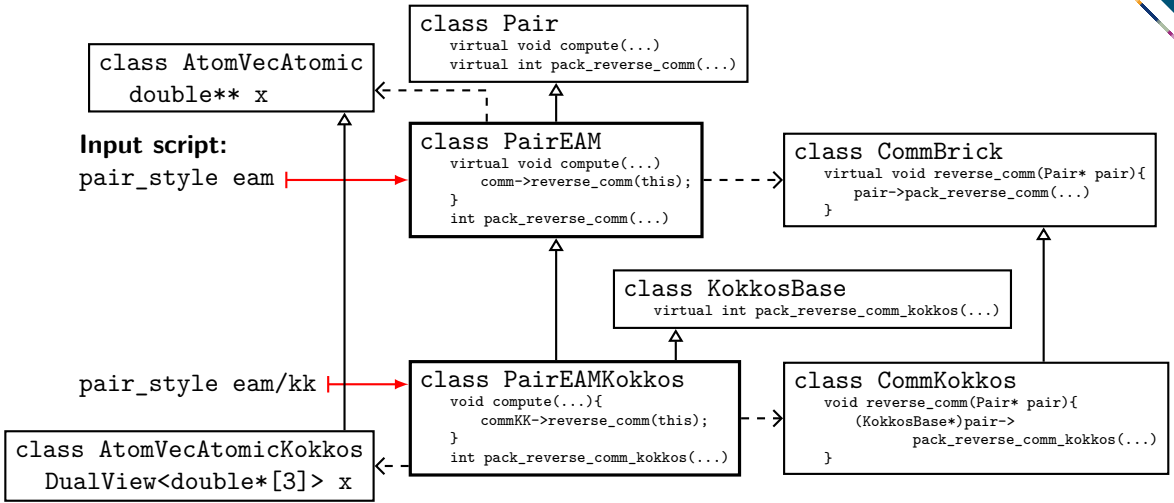

HOW FAST IS IT?

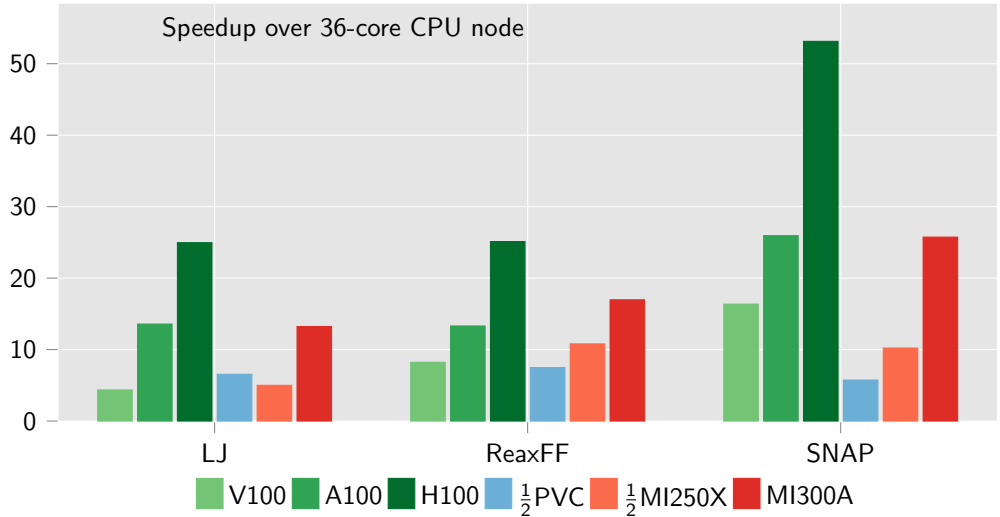

GPU CHALLENGE: SATURATION

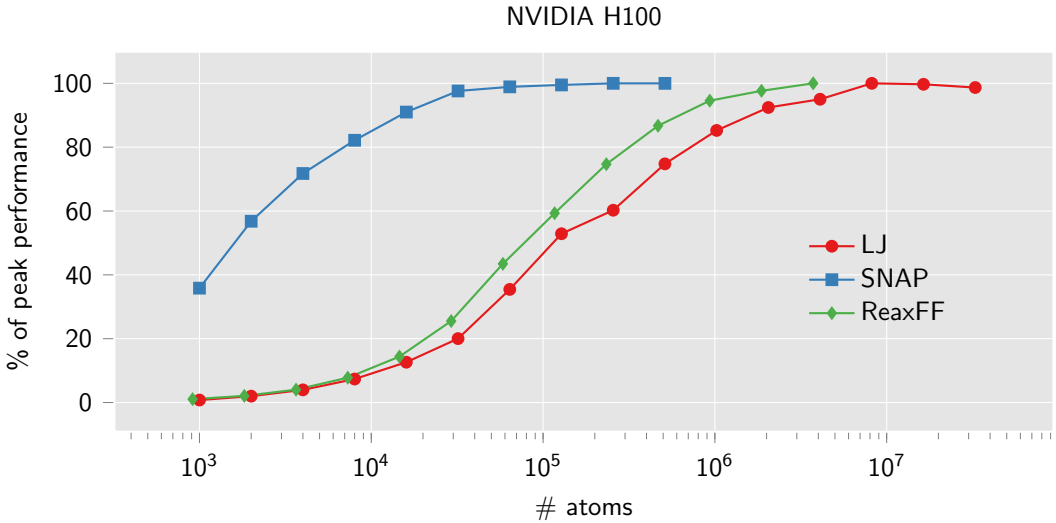

OPTIONS MATTER: PAIR-WISE POTENTIALS

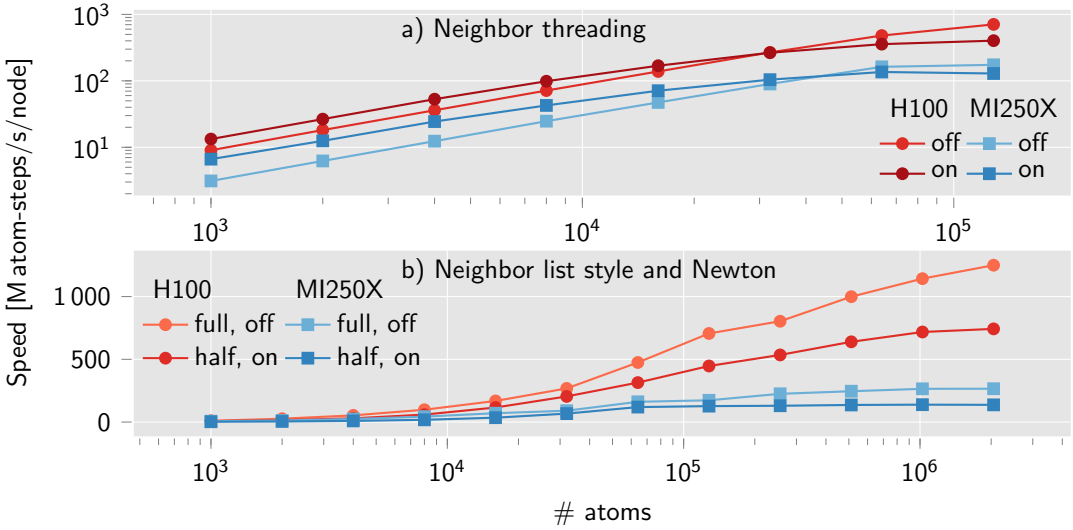

LAMMPS AT EXASCALE

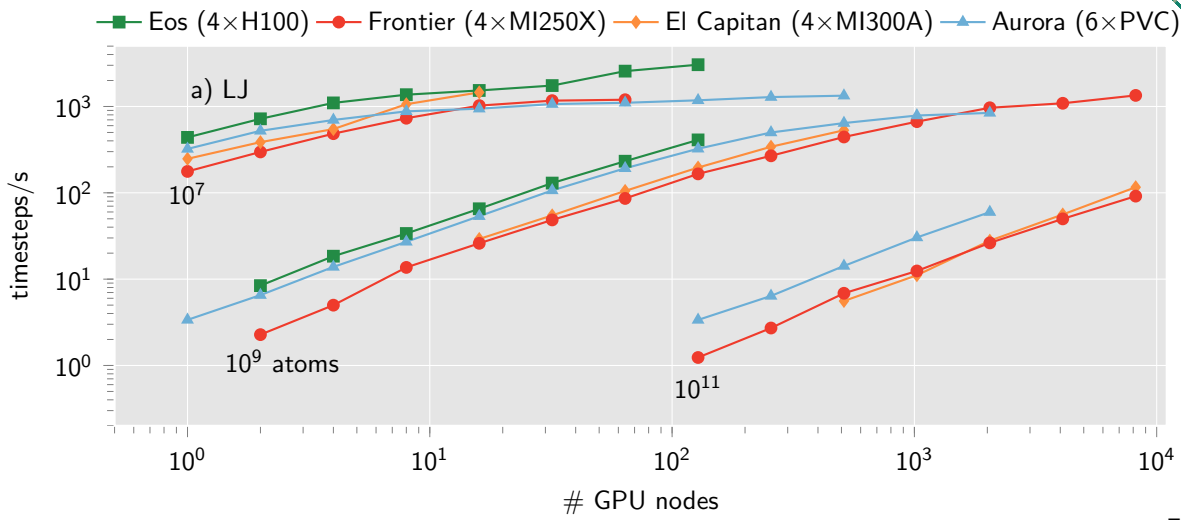

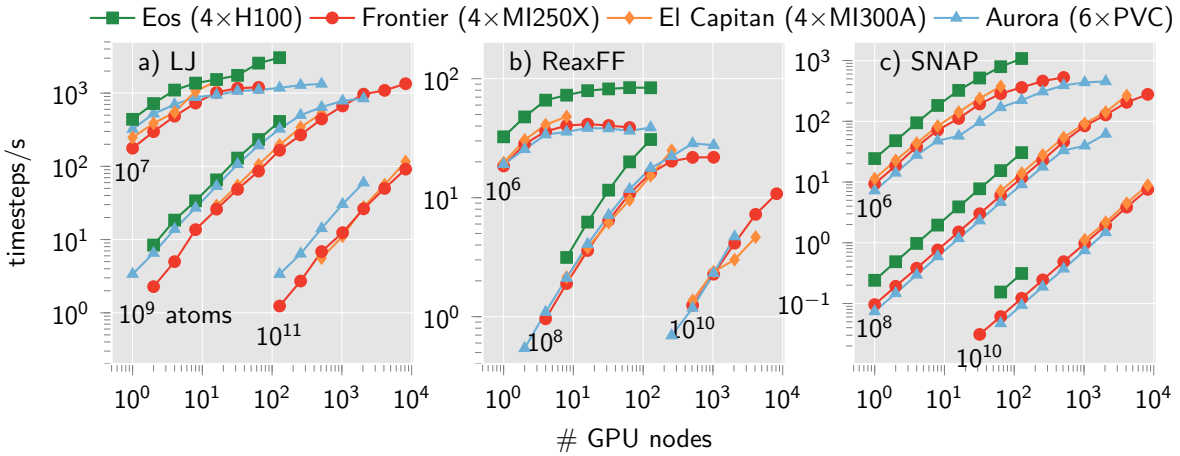
Eos (4×H100)
Frontier (4×MI250X)
El Capitan (4×MI300A)
Aurora (6×PVC)
a) LJ
b) ReaxFF
c) SNAP
timesteps/s
GPU nodes
10^7
10^9 atoms
10^{11}
10^6
10^8
10^{10}
10^6
10^8
10^{10}

WHY ARE H100S SO FAST?

GPU	BW	FP64	L1 + Shared
NVIDIA H100	3.3 TB/s	34 TF	**256 kB**
AMD MI250x/2	1.6 TB/s	24 TF	16 + 64 kB
AMD MI300A	5.3 TB/s	61 TF	32 + 64 kB
Intel PVC stack	1.6 TB/s	26 TF	n/a + 128 kB

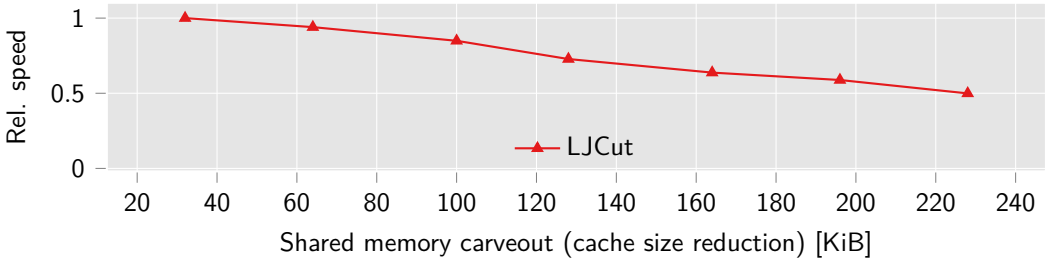

WHY ARE H100S SO FAST?

GPU	BW	FP64	L1 + Shared
NVIDIA H100	3.3 TB/s	34 TF	**256 kB**
AMD MI250x/2	1.6 TB/s	24 TF	16 + 64 kB
AMD MI300A	5.3 TB/s	61 TF	32 + 64 kB
Intel PVC stack	1.6 TB/s	26 TF	n/a + 128 kB

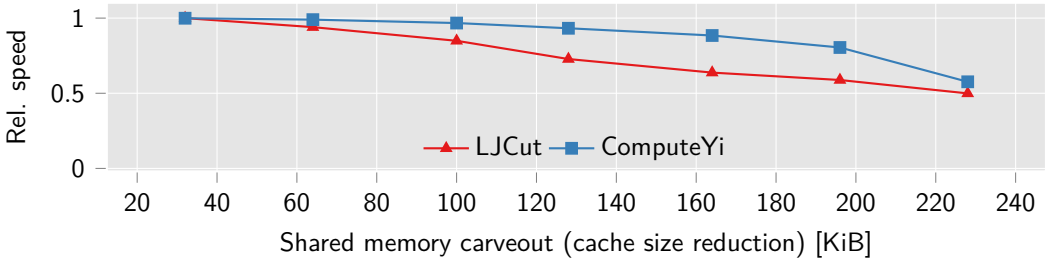

WHY ARE H100S SO FAST?

GPU	BW	FP64	L1 + Shared
NVIDIA H100	3.3 TB/s	34 TF	**256 kB**
AMD MI250x/2	1.6 TB/s	24 TF	16 + 64 kB
AMD MI300A	5.3 TB/s	61 TF	32 + 64 kB
Intel PVC stack	1.6 TB/s	26 TF	n/a + 128 kB

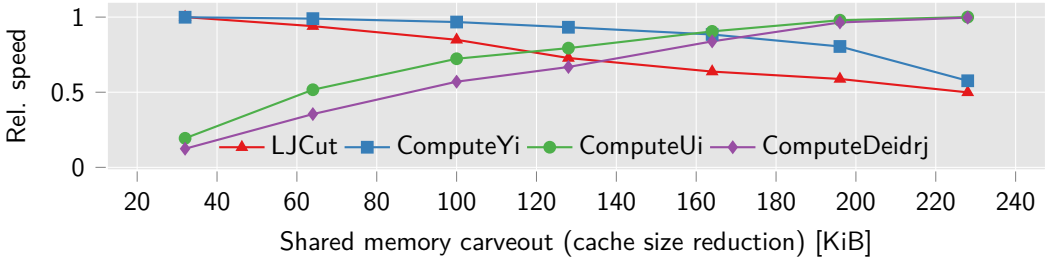

TRILLION-ATOM MD WITH FLARE

May 2024: Got access to Frontier, then #1 in the world

First trillion-atom MD on GPUs!

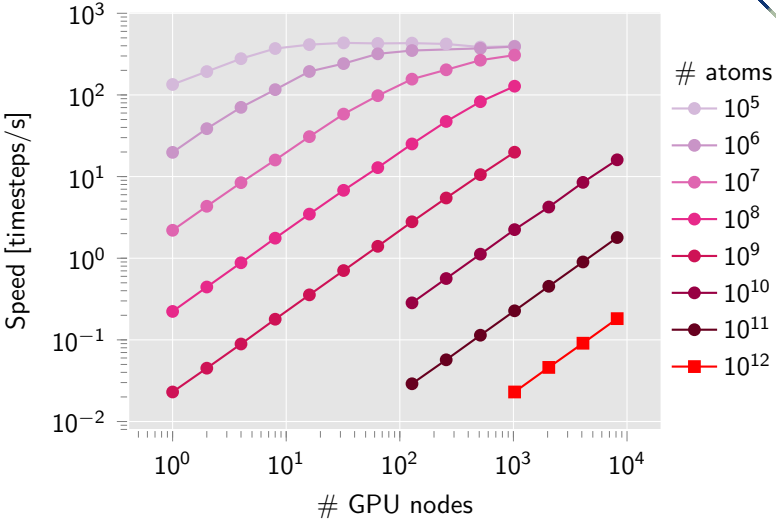

TRILLIONS