

Performance on Advanced Systems Test Beds



Christian Trott, Simon Hammond, Sue Kelly, Jim Laros and Jim Ang

Scalable Algorithms (01426, NM) and Scalable Computer Architectures (01422, NM)

Stable

Volta

Cray XC30m

of Nodes: 56

Processor:

2x Intel E5-2695-V2

Ivy Bridge, 12 cores, 2.4GHz

Co-Processor:

None

Network:

Cray Aries

Usage:

Power profiling

System software

Baseline benchmarking



Shannon

NVIDIA Kepler

of Nodes: 32

Processor:

2x Intel E5-2670

Sandy Bridge, 8 cores, 2.6GHz

Co-Processor:

2x Nvidia K20xm

Network:

Mellanox QDR IB

Usage:

Programming model research

Portable applications

Heterogeneous computing



Experimental

Compton

Intel Xeon Phi

of Nodes: 42

Processor:

2x Intel E5-2670

Sandy Bridge, 8 cores, 2.6GHz

Co-Processor:

2x Intel Xeon Phi 57 cores

pre-production, 1.1GHz

Network:


Mellanox QDR IB

Usage:

Programming model research

Threading

System software



Teller

AMD Fusion

of Nodes: 104

Processors:

AMD A10-5800K

Piledriver, 4 cores, 3.8GHz

Co-Processor:

Radeon HD-7600D

Northern Islands on-die

Network:

QLogic QDR IB

Usage:

Hybrid Chip investigation

Power profiling



Future

Watson

IBM POWER

of Nodes: 10

Processor:

2x IBM POWER 7+

8 cores, 4 way SMT, 3.47 GHz

Co-Processor:

Altera FPGA


Network:

10Gb Ethernet

Usage:

Programming Models

Heterogenous computing



NVIDIA Kayla

32 bit ARM + GPU

of Nodes: 1

Processor:

32 bit Quad core A9

Co-Processor:

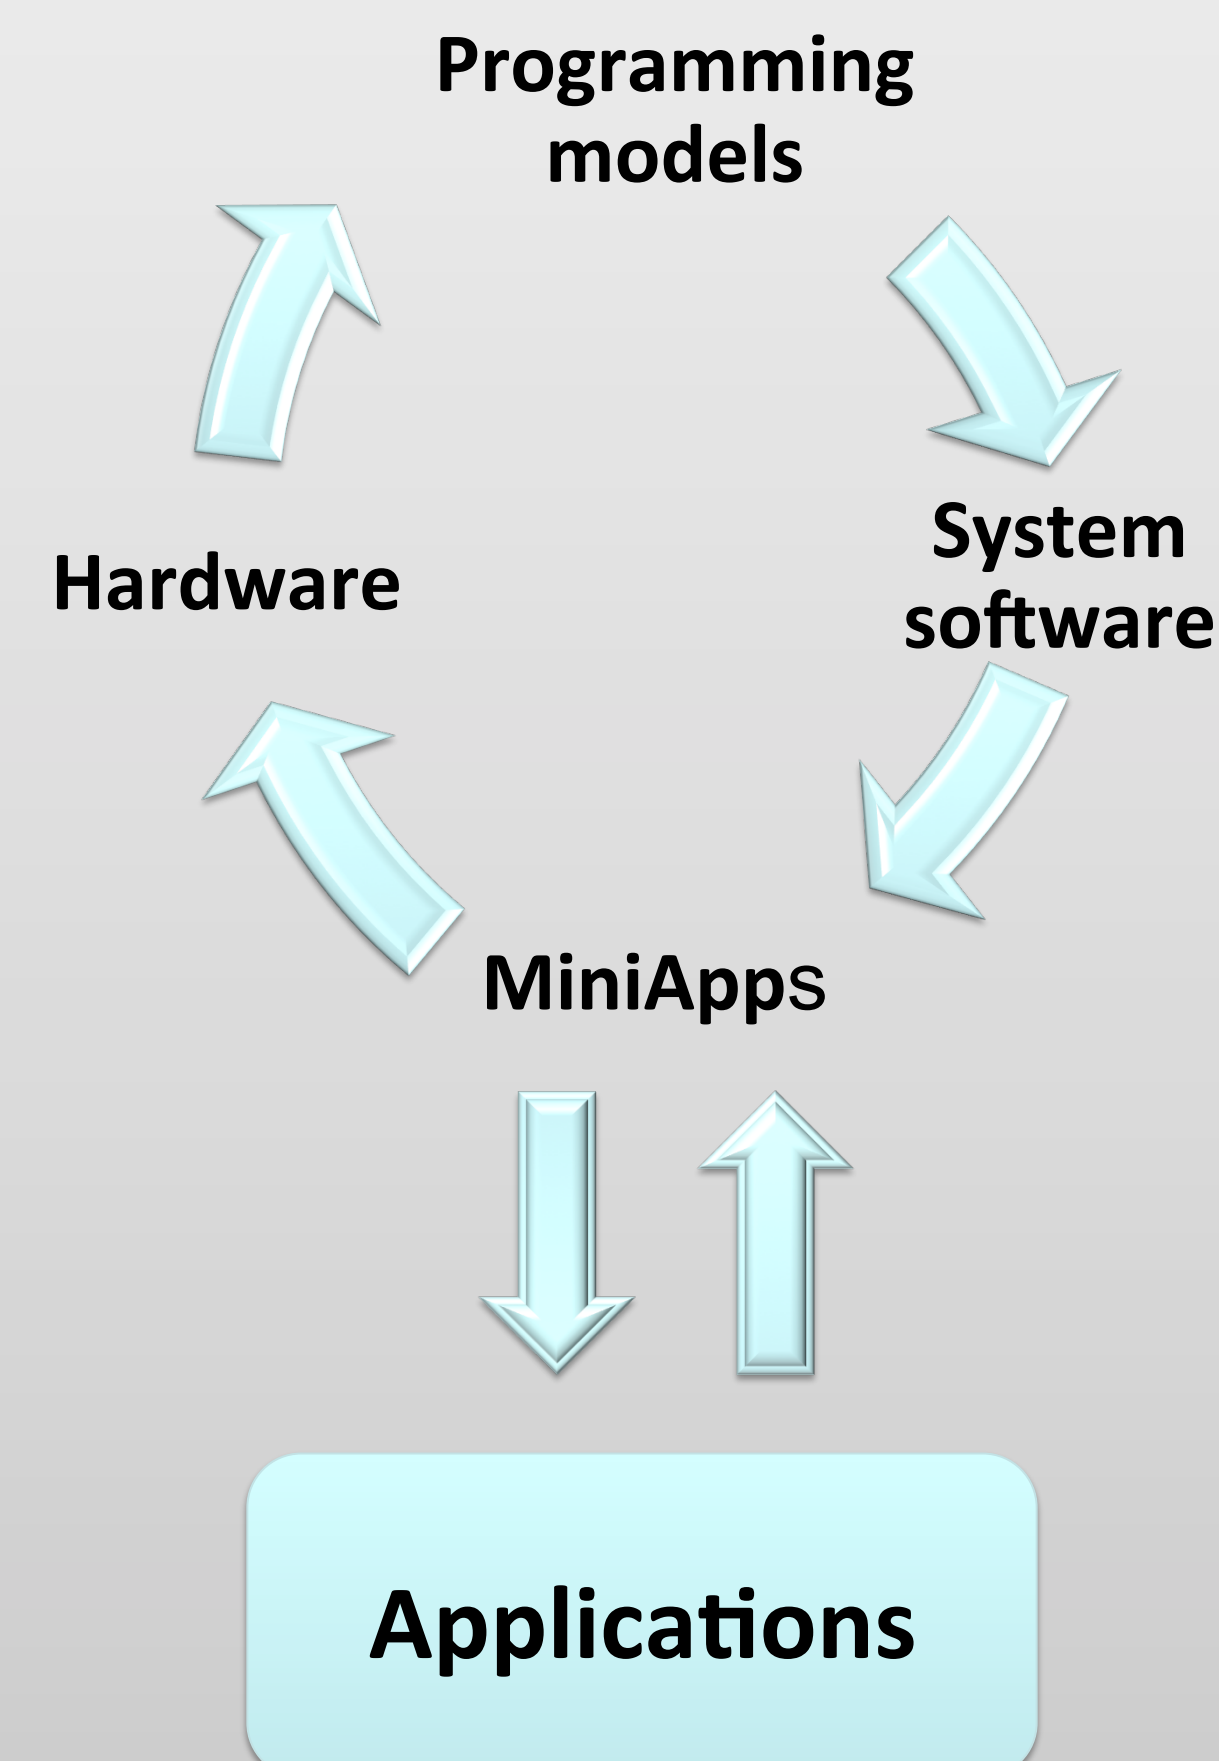
Nvidia GTX640

Coming Soon

64 bit ARM Cluster



Workflow



MiniFE

Finite element code miniApp in Mantevo (mantevo.org)

Heat conduction, Matrix assembly, CG solve

Most variants of any miniApp in Mantevo

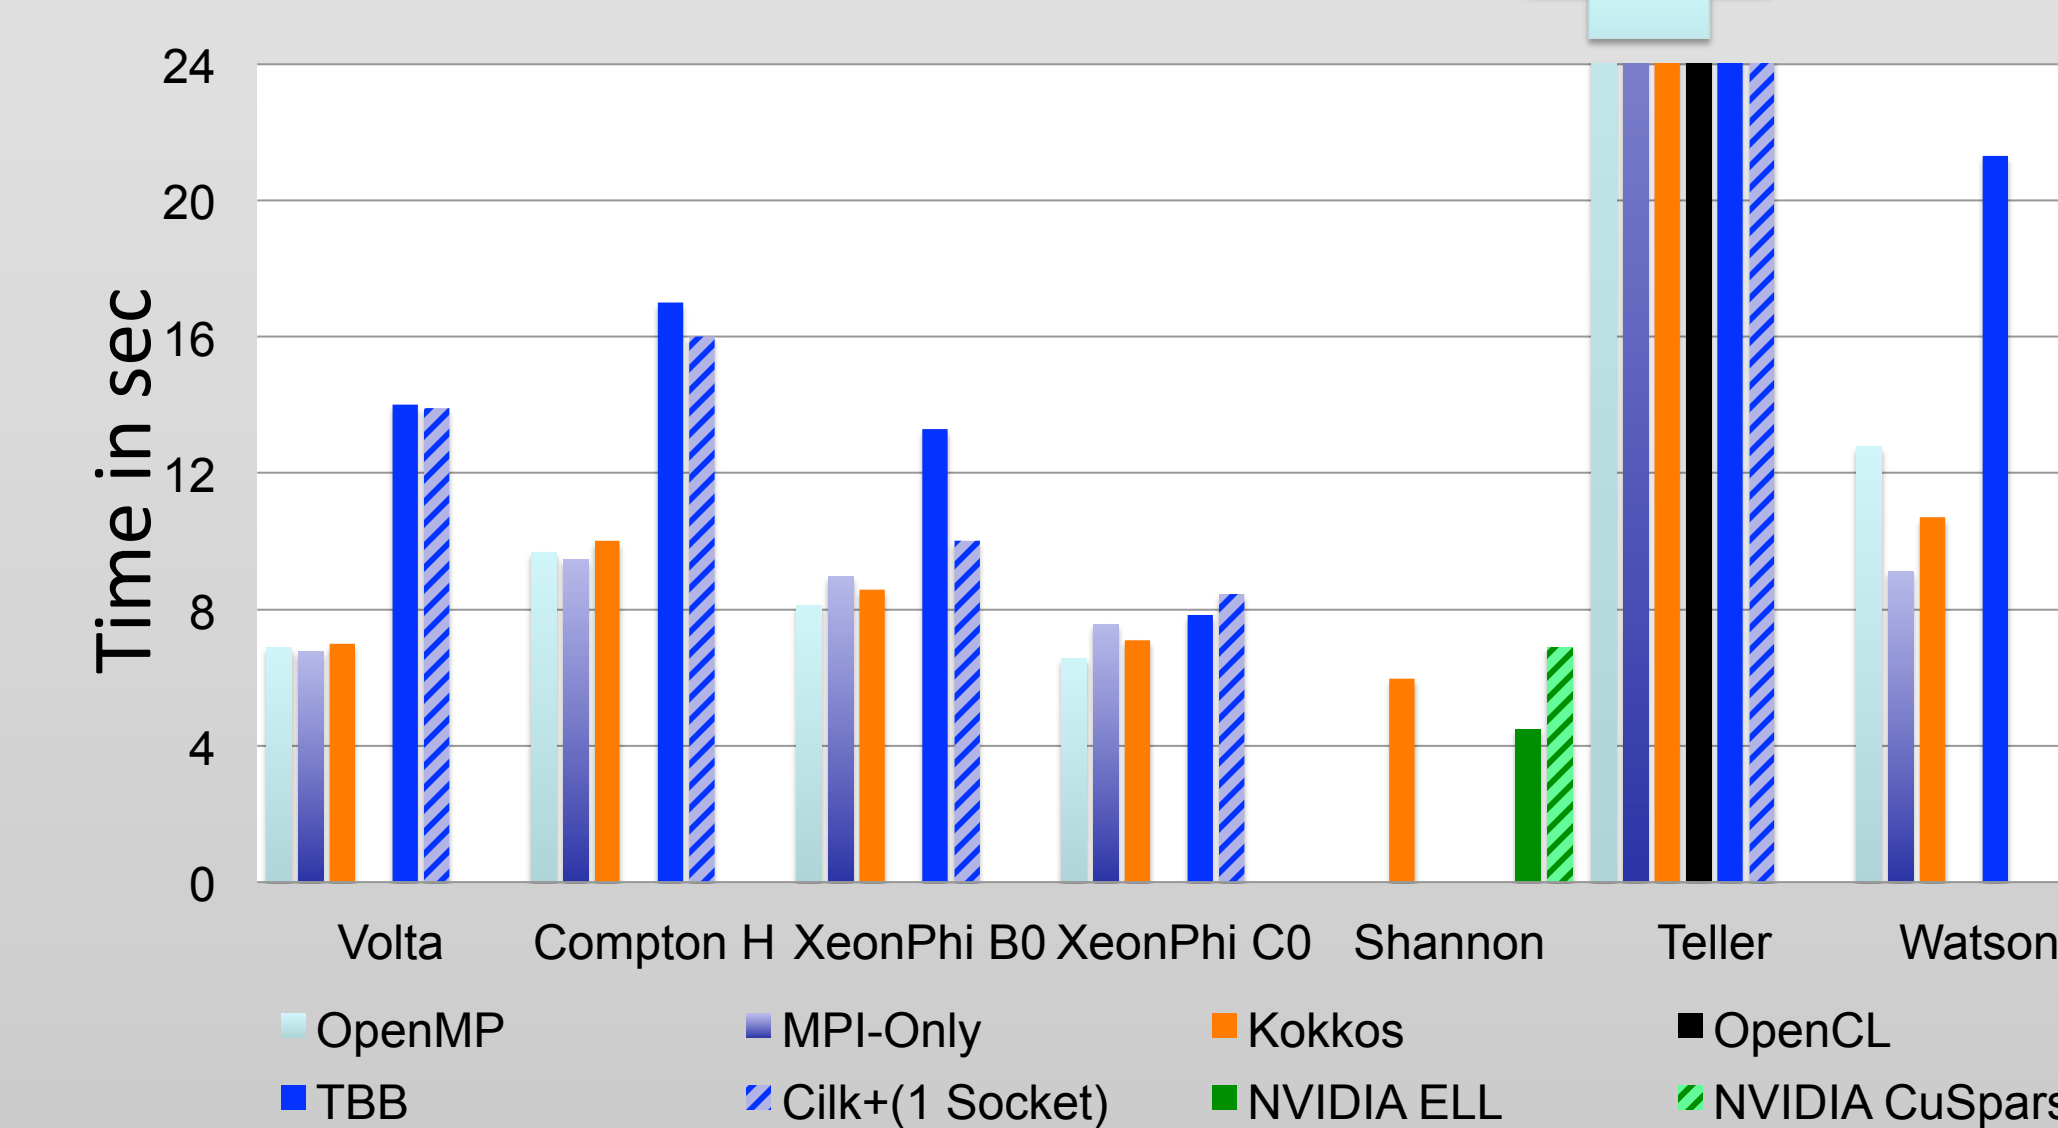
more than 20 implementations in Mantevo repository

Models aspects of Sandia's mechanical engineering codes

MiniFE CG-Solve time

200x200x200 cells, 200 iterations

44-57s



Trilinos

Large package based library for the solution of multi-physics engineering and scientific problems

Prototype using Kokkos under development

enable thread scalability on all platforms throughout software stack

Expected to be pushed in production during FY 2014

Goal: Preparing Trilinos and its dependend software (e.g. Sierra) for next generation super computers such as the Trinity machine at LANL and Sandia in 2015.

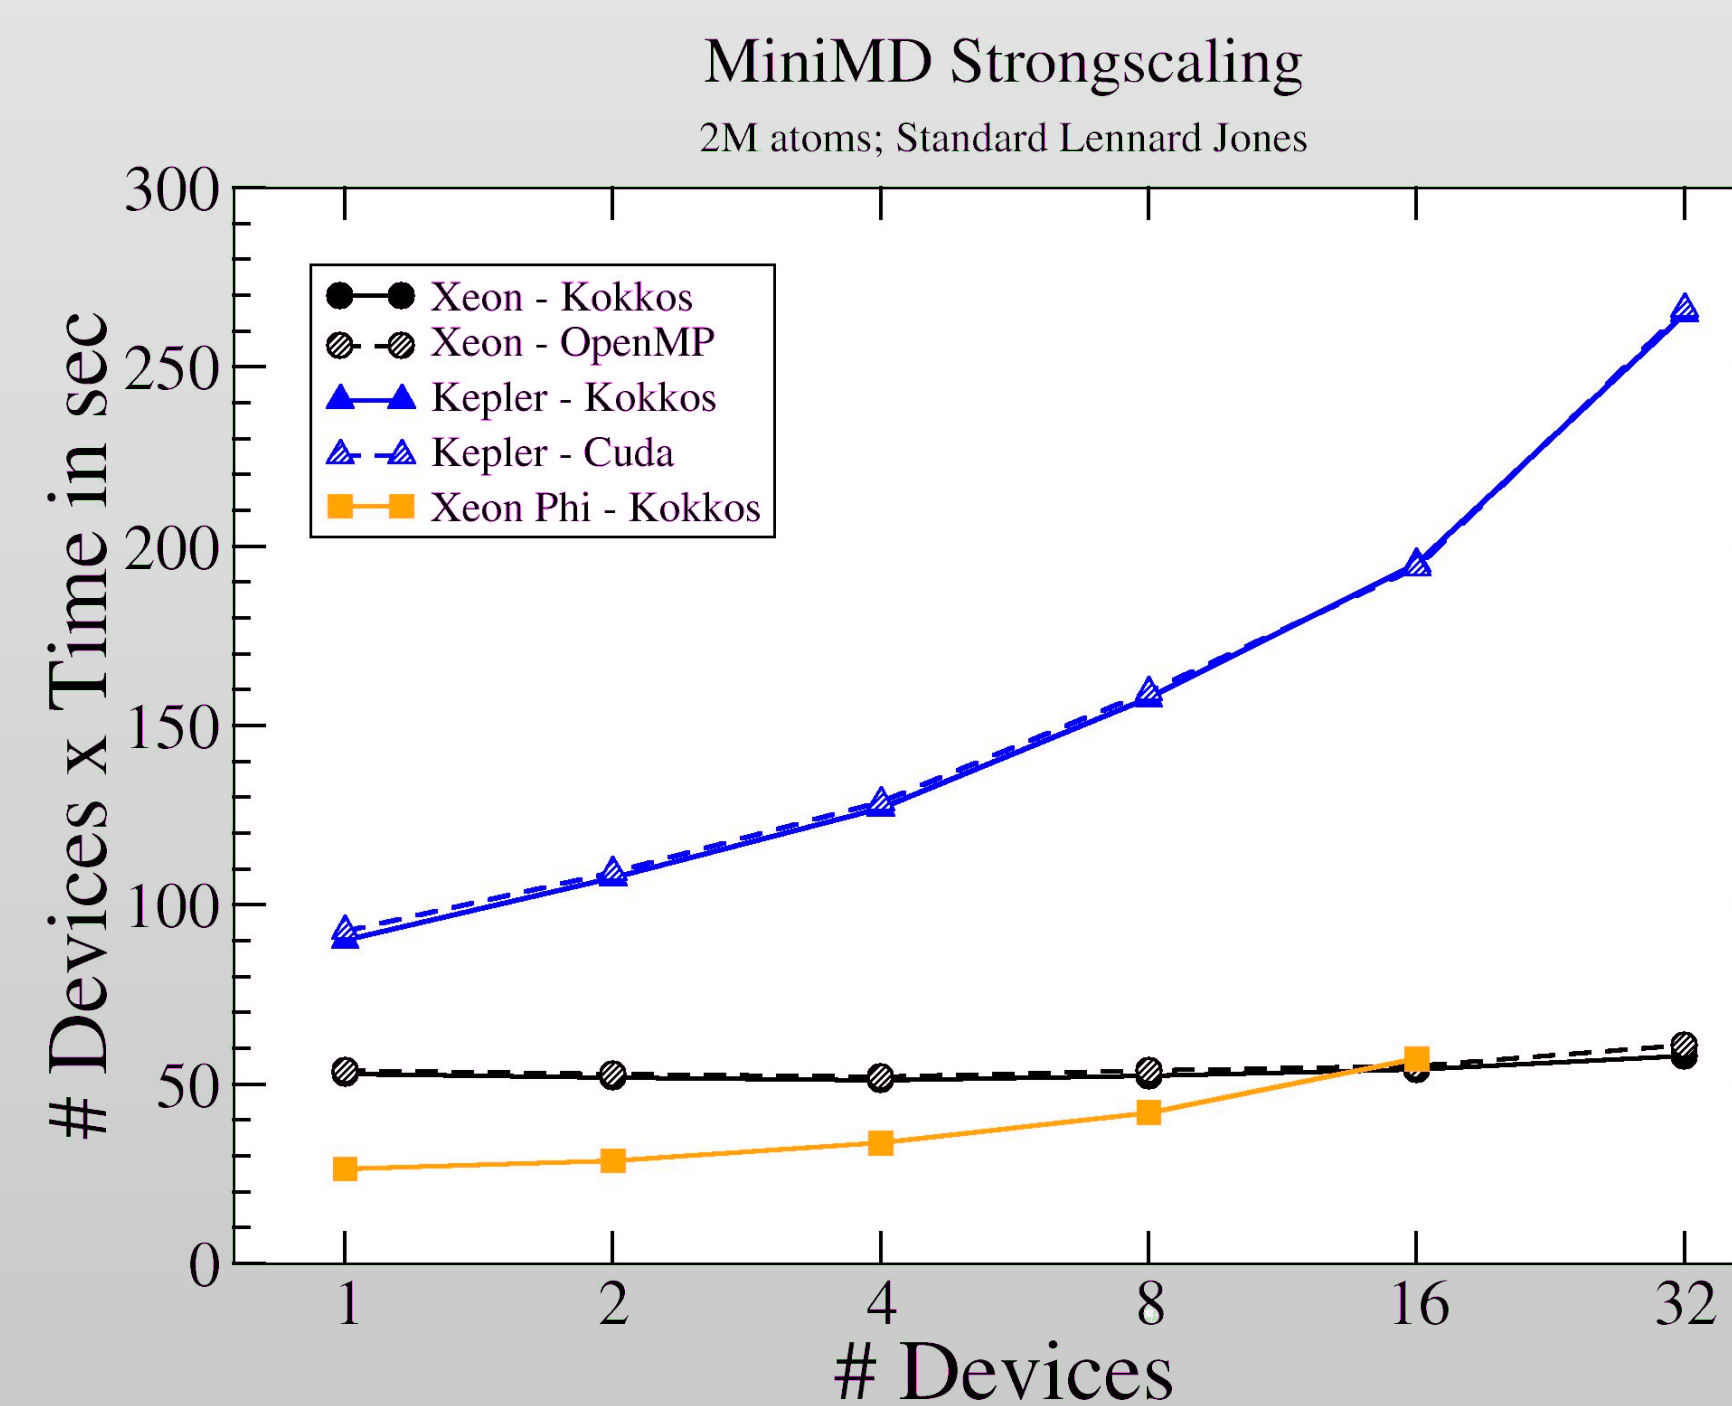
MiniMD

Molecular Dynamics application (mantevo.org):

simplified LAMMPS: single atom type, less table lookups

Variants:

Reference (MPI+OpenMP; SNL), OpenCL (SNL), Kokkos (SNL), Intel Xeon Phi intrinsics (Intel), OpenACC (AMD), Chapel (Cray), Intel intrinsics (Warwick/Intel), Qthreads (SNL),



LAMMPS-Kokkos Prototyp

Exploration of Kokkos for use in LAMMPS (lammps.sandia.gov)

replace specialized packages => reduce code redundancy 3x

enable thread scalability throughout code base

Leverage algorithmic exploration from MiniMD

transferring thread-scalable algorithms

