IPCC @ RWTH Aachen University
Showcase: KNL results

Paolo Bientinesi Rodrigo Canales
Markus Höhnerbach William McDoniel
Paul Springer Ahmed E. Ismail

ISC'16

June 19, Frankfurt

Road to Performance on KNL

- Part 1: **Performance via compiler**
- Part 2: **Performance via intrinsics**
- Part 3: **Performance via search**

LAMMPS

Large-scale **A**tomic–**M**olecular **M**assively **P**arallel **S**imulator

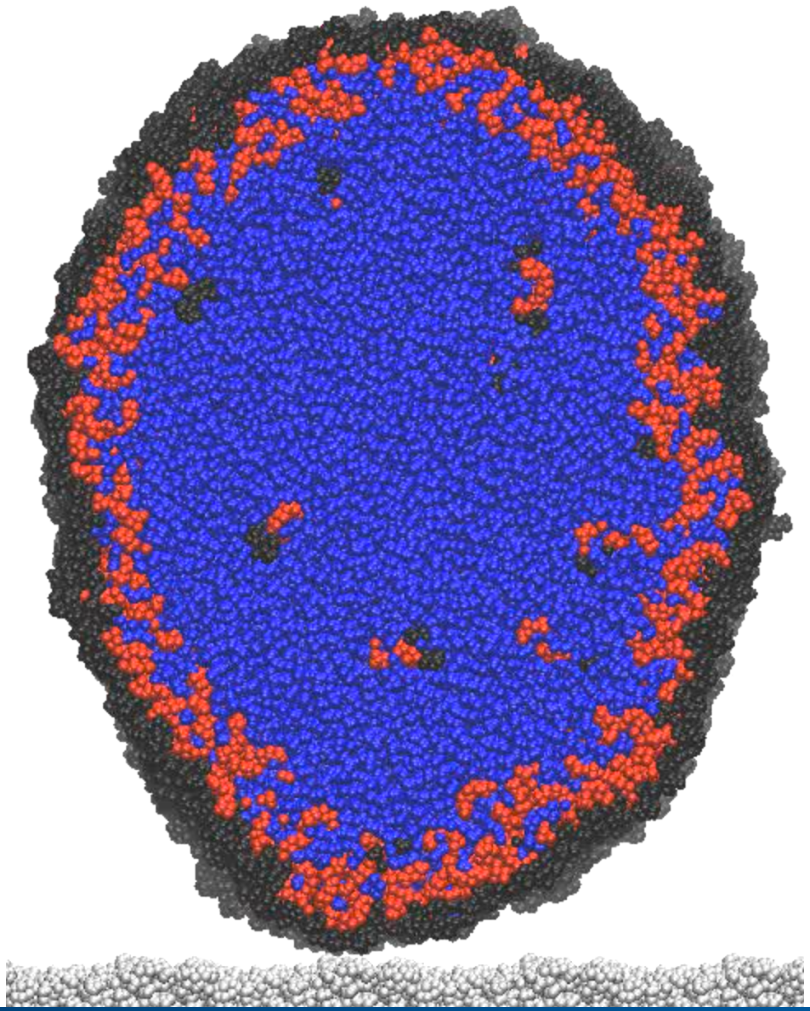

- Sandia National Labs
  `http://lammps.sandia.gov`
- Wide collection of potentials
- Open source
  Support for OpenMP, Xeon Phi, and GPU (CUDA and OpenCL)

Part 1: Buckingham pair potential

Rodrigo Canales

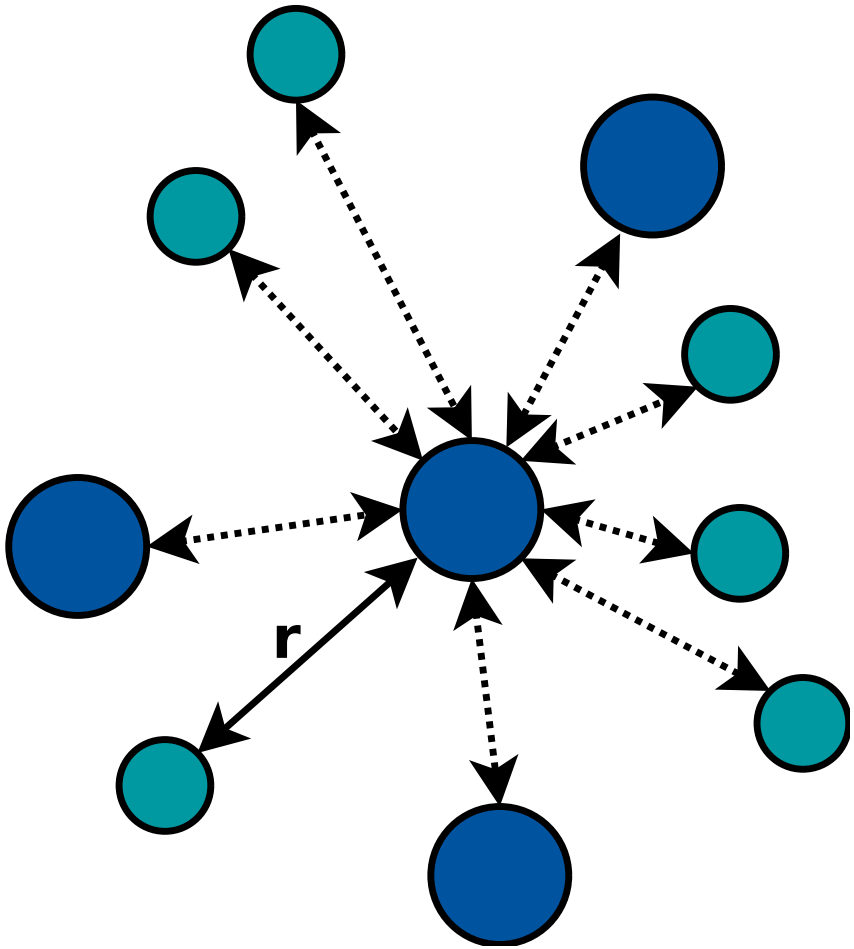

$$\Phi_{buck} := Ae^{-r/\rho} - \frac{C_b}{r^6}$$

$$\Phi_{buck/coul} := \Phi_{buck} + \frac{C_c q_i q_j}{\epsilon\, r_{ij}}$$

Buckingham: vectorization, single thread

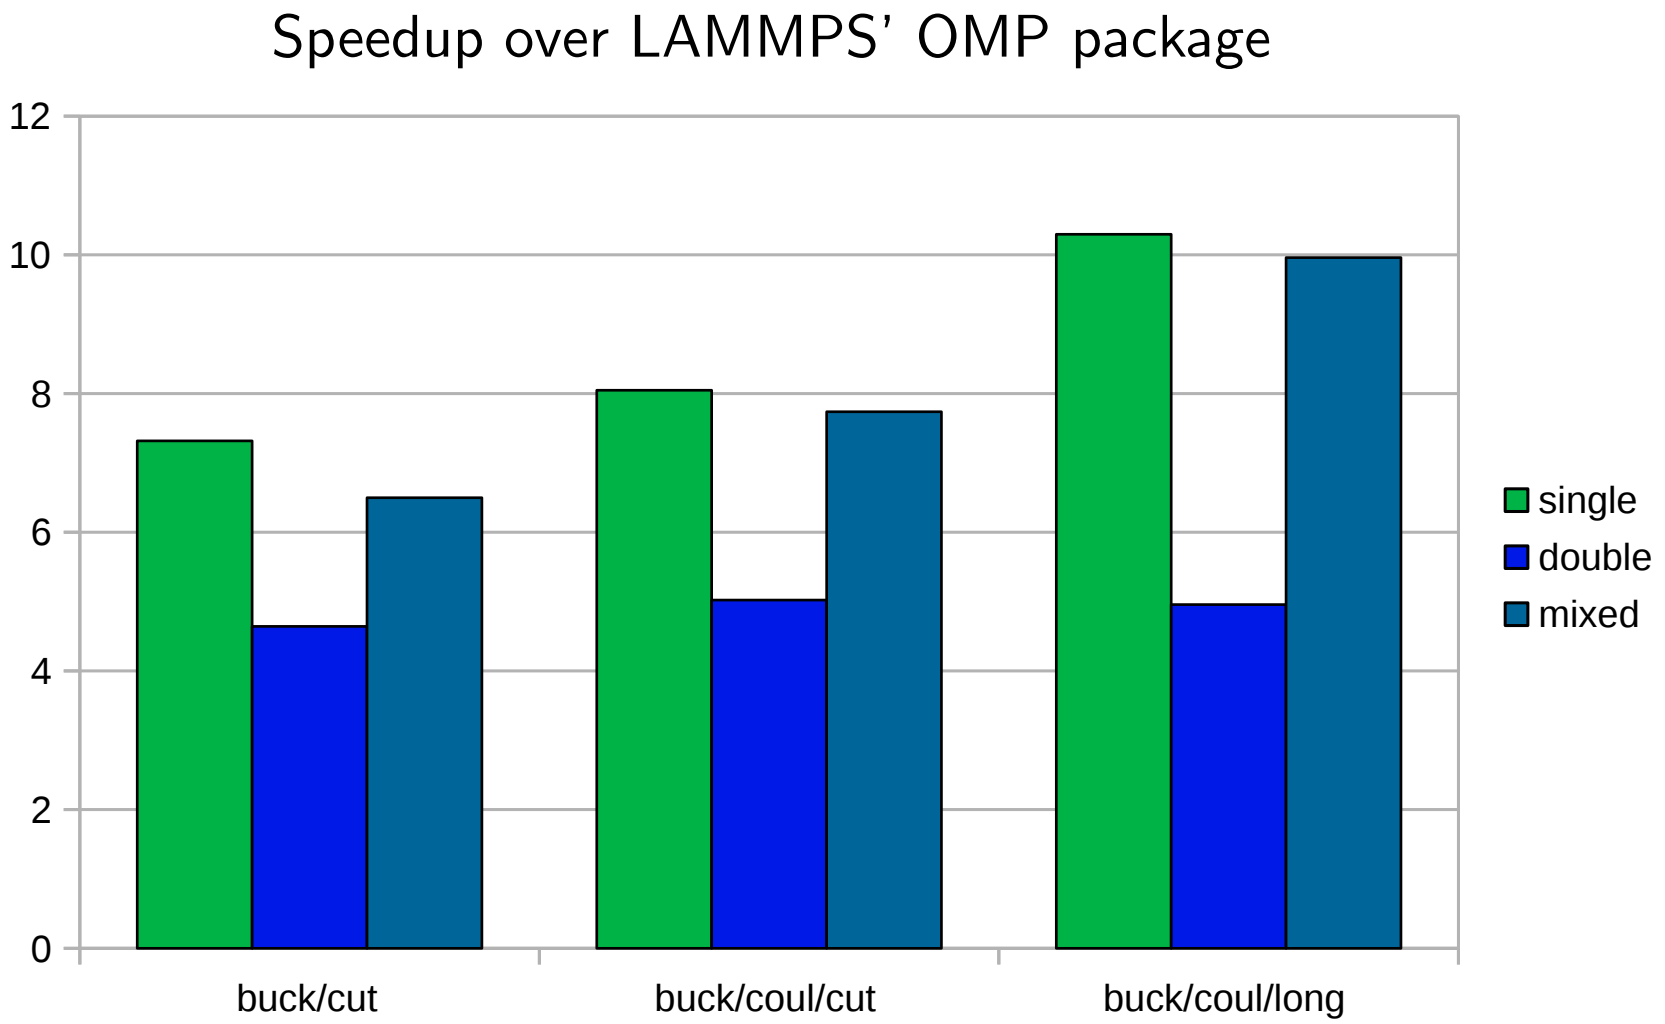

Buckingham: vectorization, full node

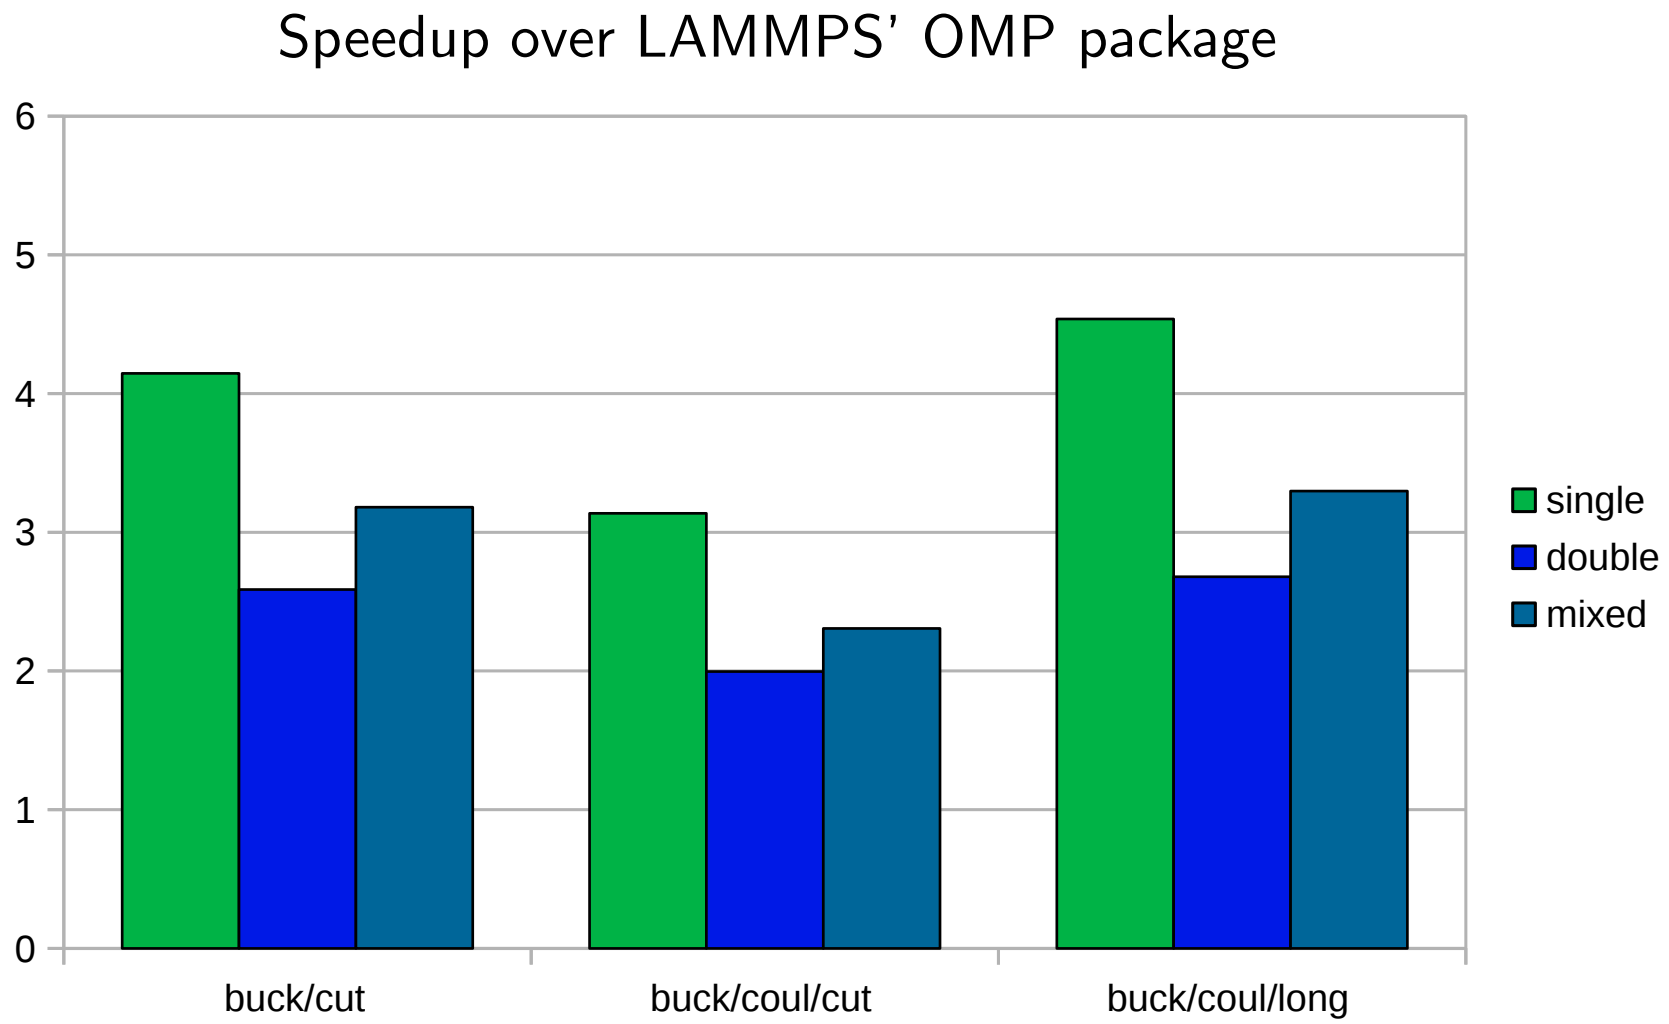

Buckingham: KNC vs. KNL

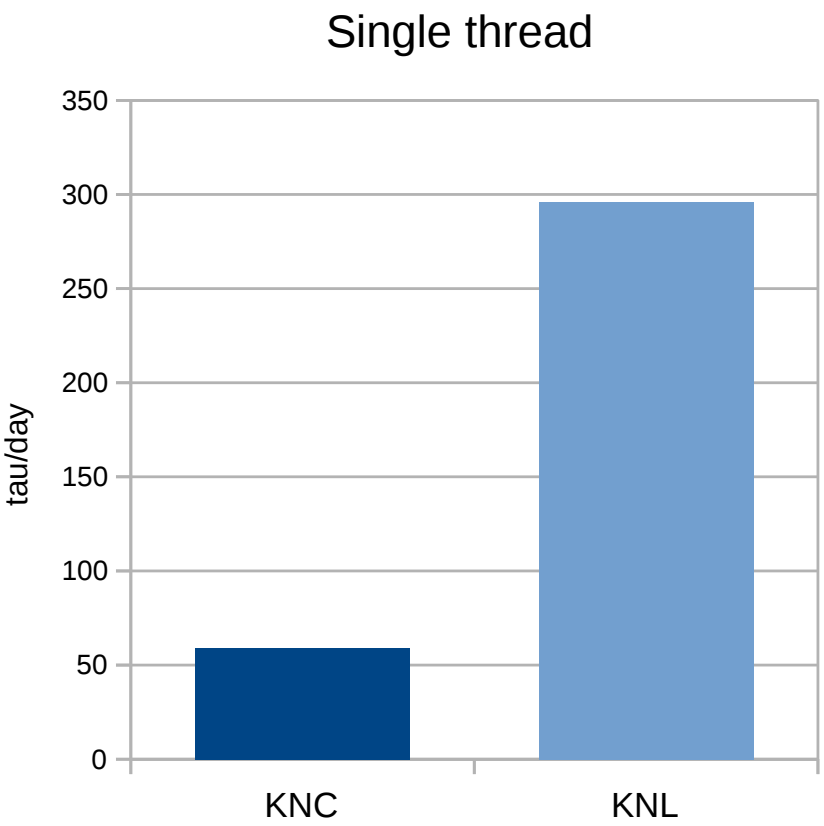

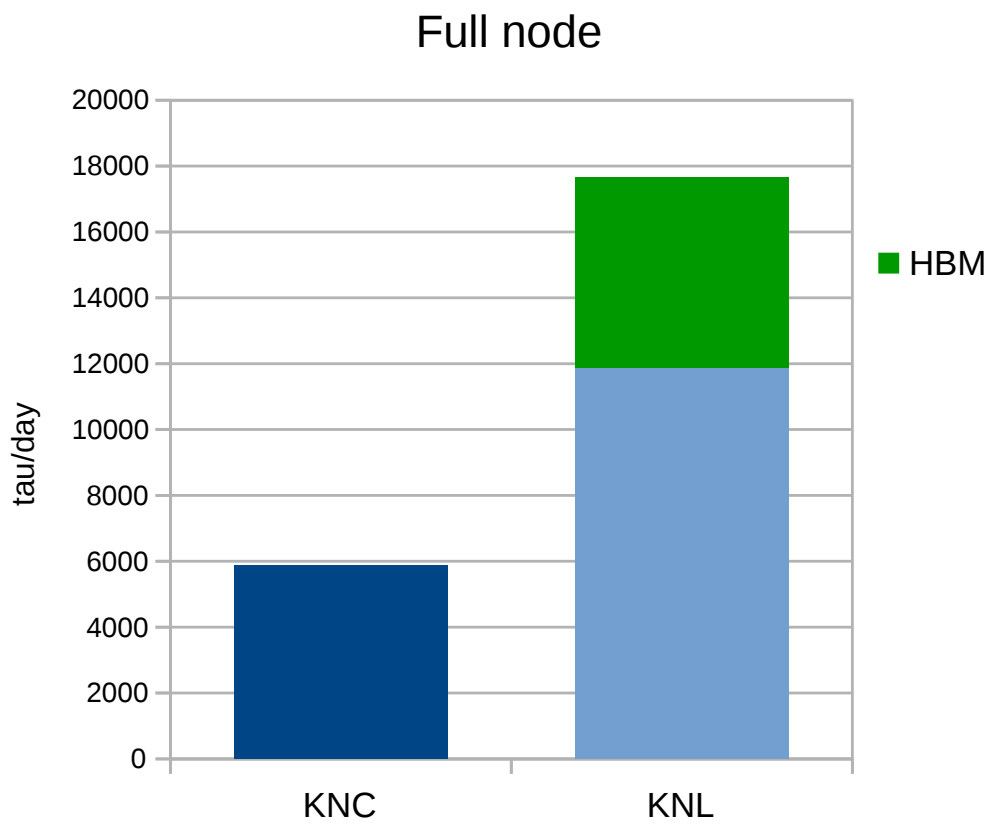

Part 2: Tersoff multibody potential

Markus Höhnerbach

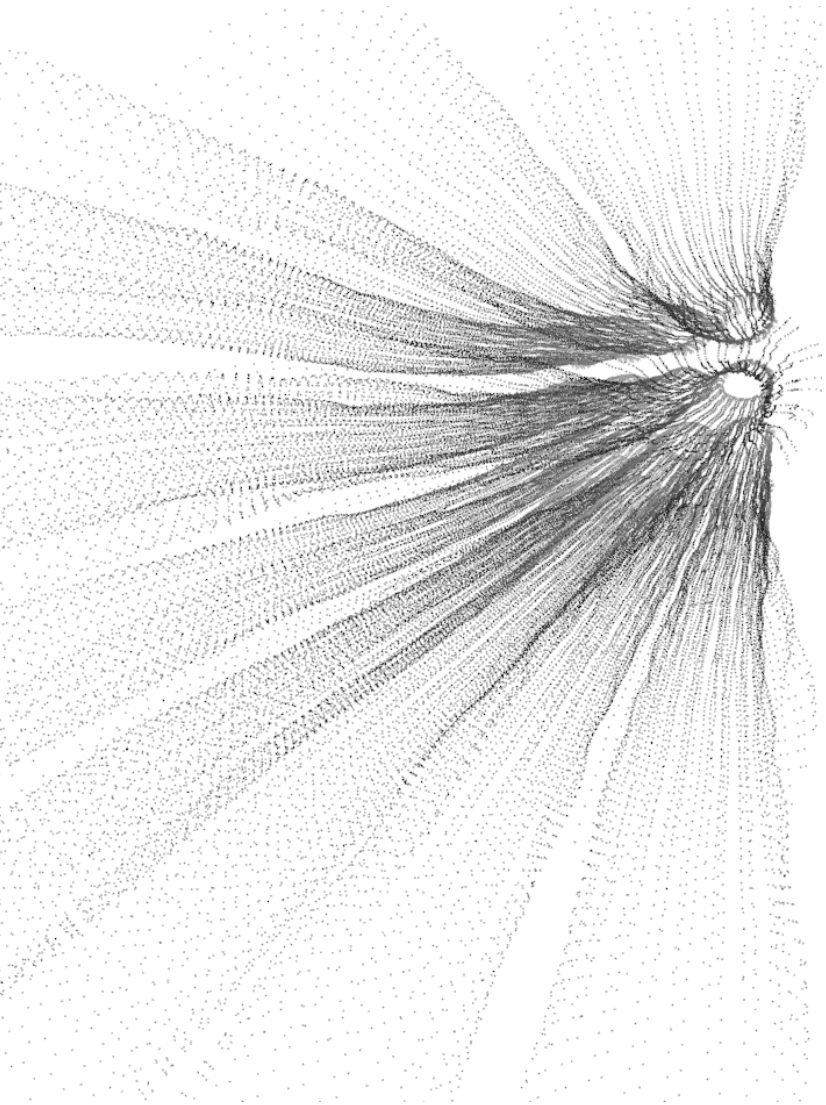

for *i in local atoms of the current thread* **do**
  for *j in atoms neighboring i* **do**
    $\zeta_{ij} \leftarrow 0;$
    for *k in atoms neighboring i* **do**
      $\zeta_{ij} \leftarrow \zeta_{ij} + \zeta(i,j,k);$
    $E \leftarrow E + V(i,j,\zeta_{ij});$
    $F_i \leftarrow F_i - \partial_{x_i} V(i,j,\zeta_{ij});$
    $F_j \leftarrow F_j - \partial_{x_j} V(i,j,\zeta_{ij});$
    $\delta\zeta \leftarrow \partial_\zeta V(i,j,\zeta_{ij});$
    for *k in atoms neighboring i* **do**
      $F_i \leftarrow F_i - \delta\zeta \cdot \partial_{x_i} \zeta(i,j,k);$
      $F_j \leftarrow F_j - \delta\zeta \cdot \partial_{x_j} \zeta(i,j,k);$
      $F_k \leftarrow F_k - \delta\zeta \cdot \partial_{x_k} \zeta(i,j,k);$

Tersoff: KNC

Thread	
Cores	1
SMT	1
Atoms	32.000

Core	
Cores	1
SMT	4
Atoms	32.000

Full	
Cores	64
SMT	4
Atoms	512.000

Measurements in 1000 atom-ns/day/core.

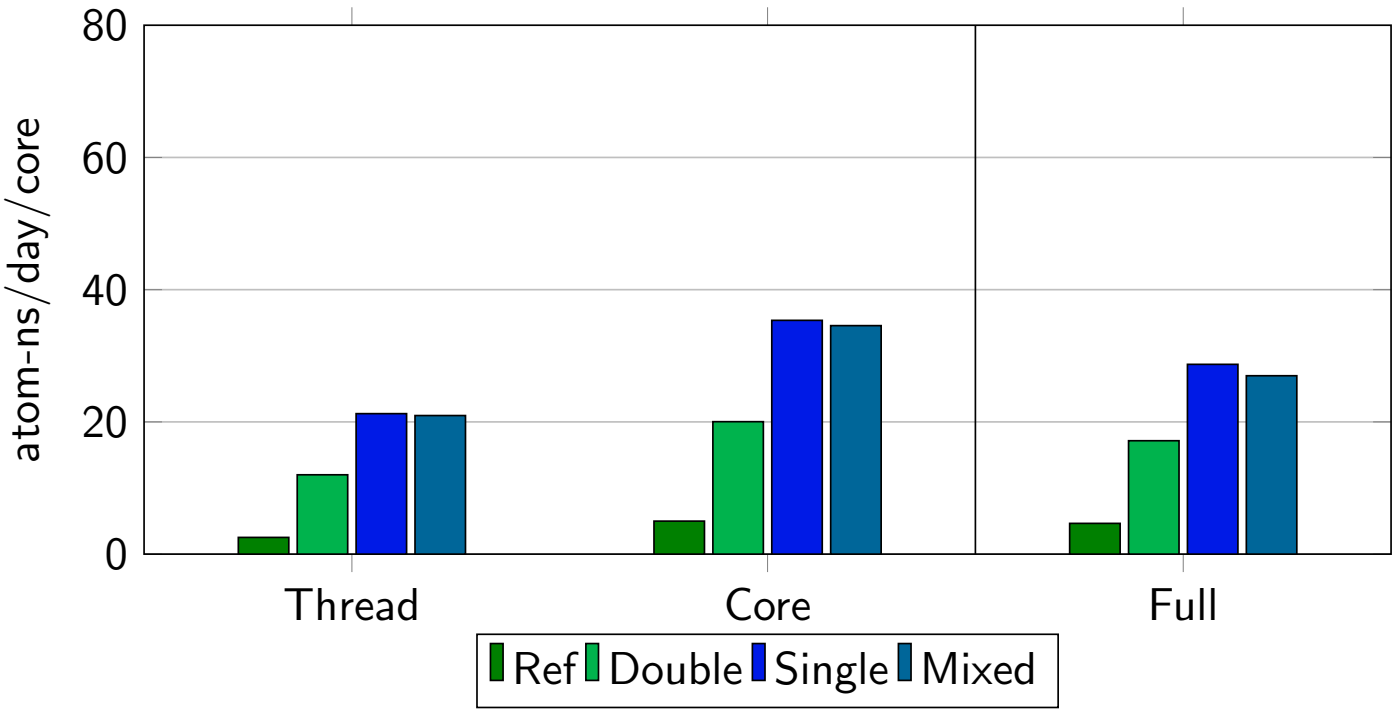

Tersoff: KNL

Thread	
Cores	1
SMT	1
Atoms	32.000
HBM	Yes

Core	
Cores	1
SMT	4
Atoms	32.000
HBM	Yes

Full	
Cores	64
SMT	4
Atoms	512.000
HBM	Yes

Measurements in 1000 atom-ns/day/core.

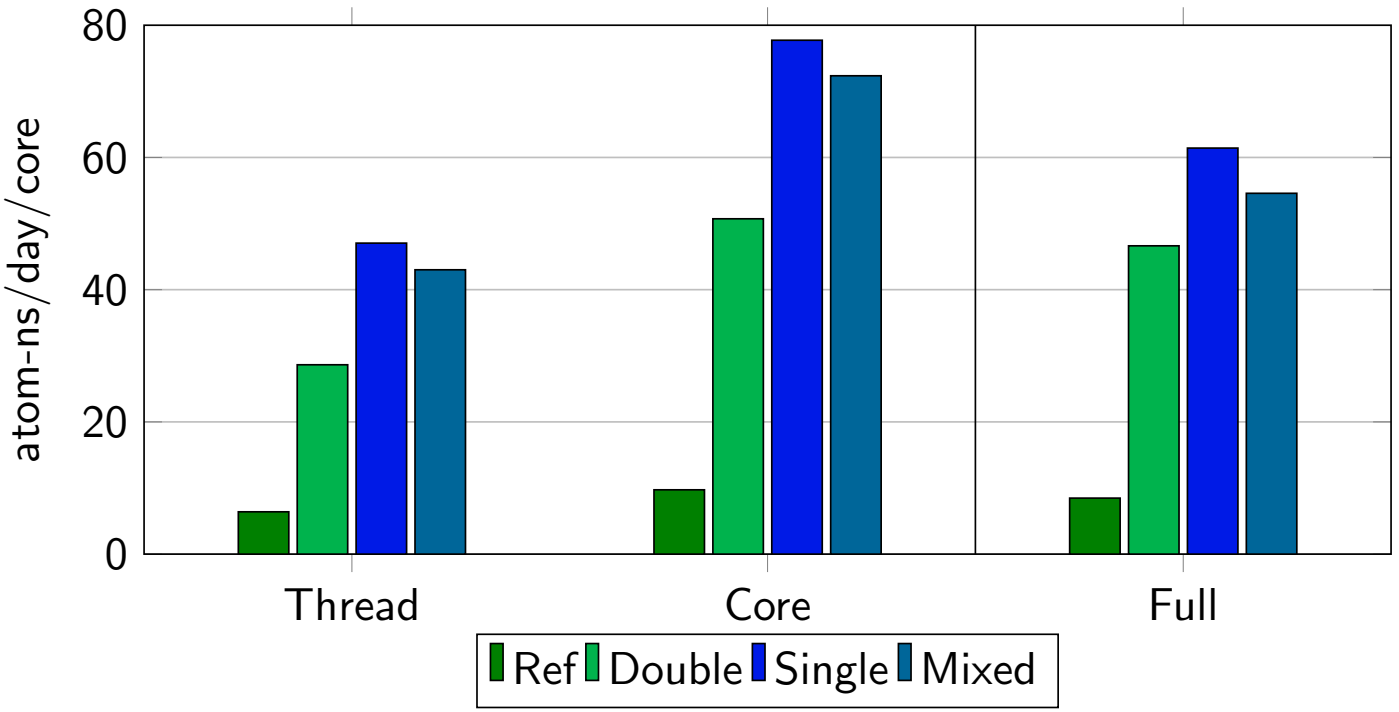

Tersoff: strong scalability

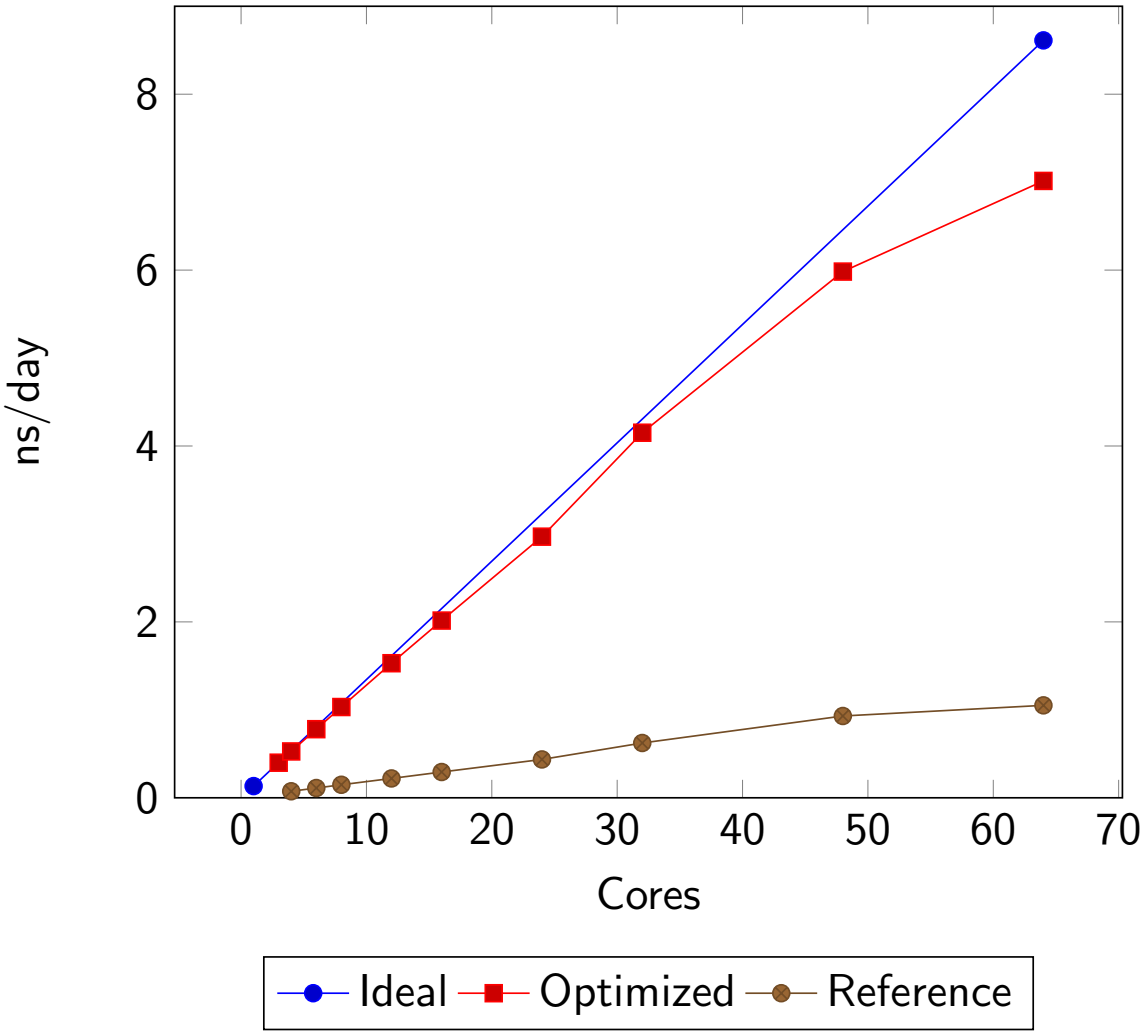

Part 3: Tensor transposition

Paul Springer

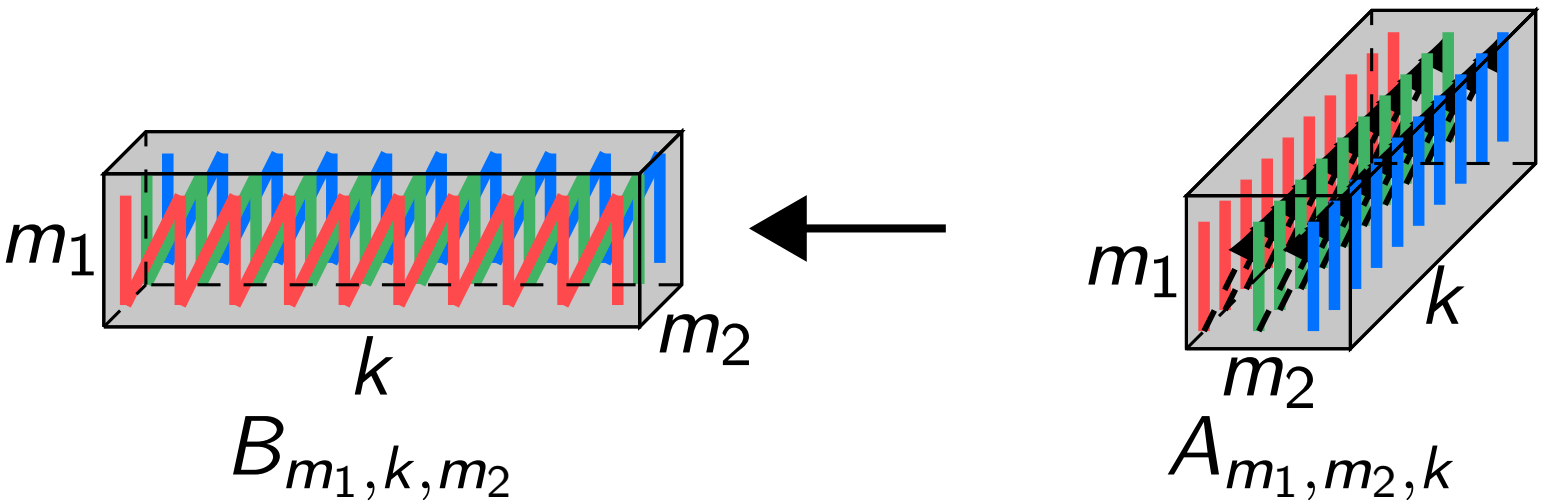

Tensor Contraction Compiler (TCC)

KNC

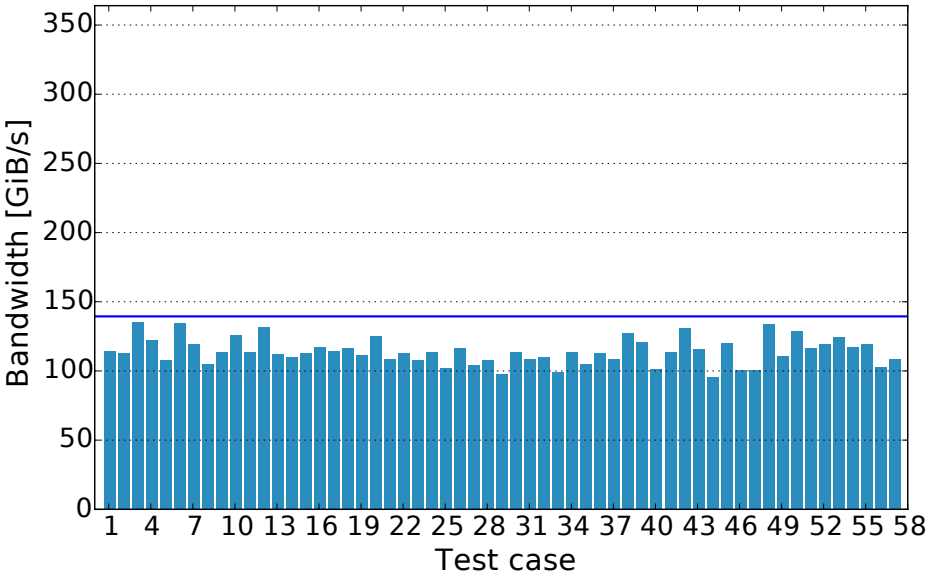

KNL

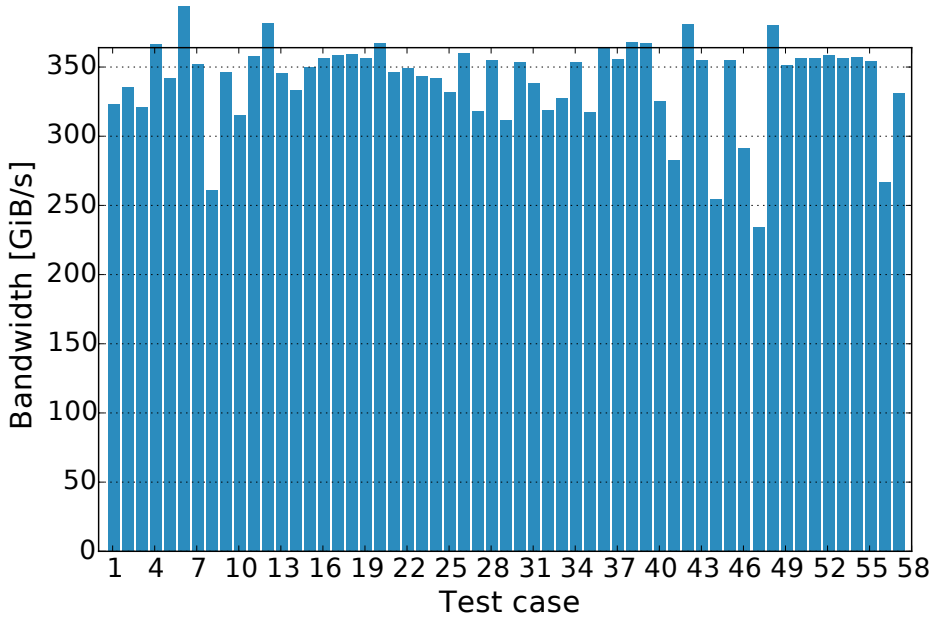

Benchmark: 57 traspositions, 2D to 6D tensors, $\geq$ 200MBs

Tensor Contraction Compiler (TCC)

KNC

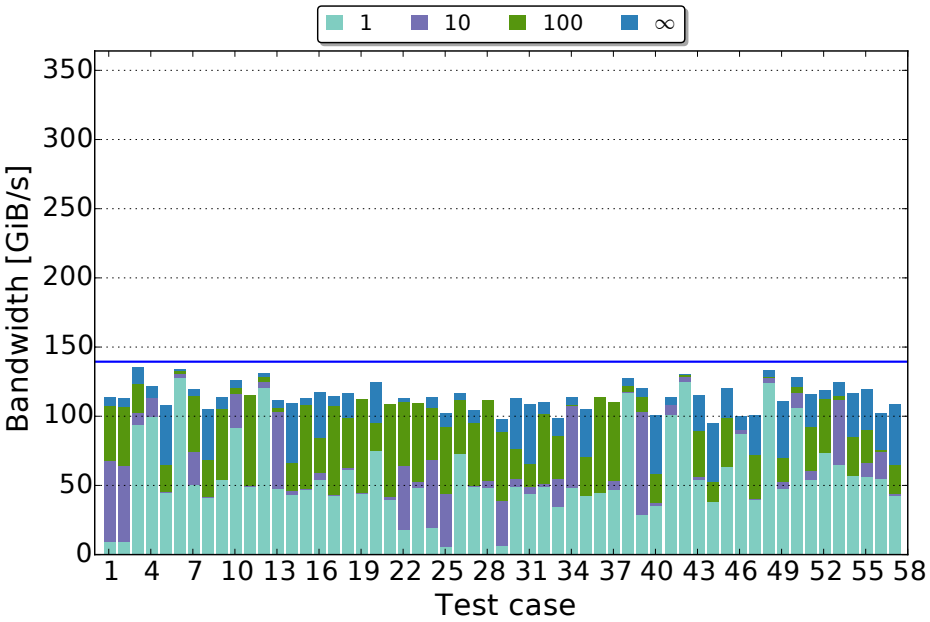

KNL

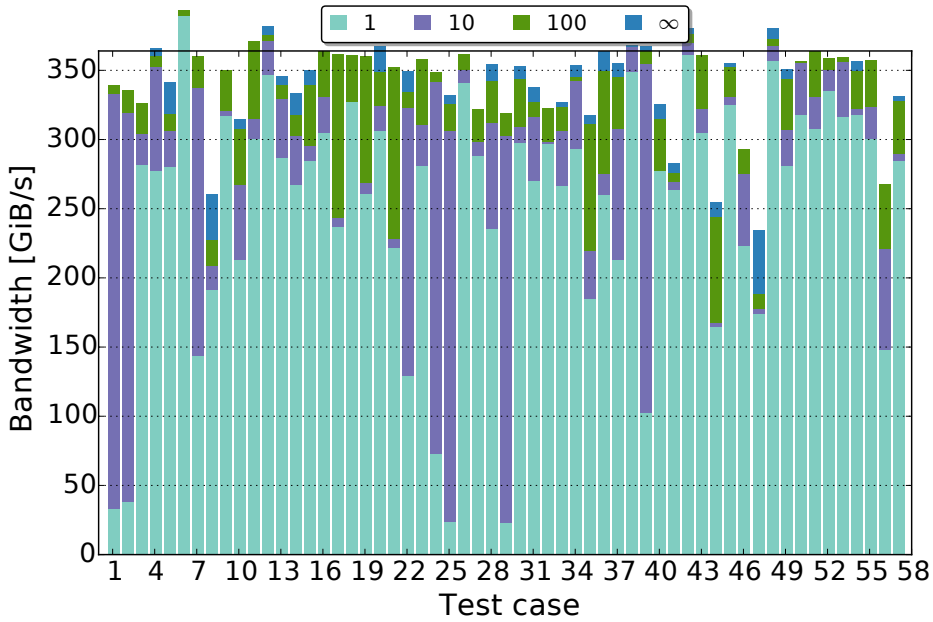

Bandwidth for the top 1, 10, 100 candidates generated by TCC.

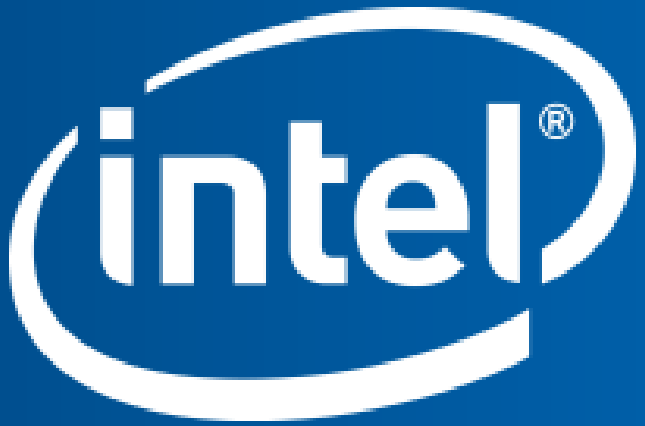
intel

TTC: KNC vs. KNL