

## LA-UR-17-22676

Approved for public release; distribution is unlimited.

Title: A Case Study on Software Modernization using CoMD - A Molecular Dynamics Proxy Application

Author(s): Adedoyin, Adetokunbo

Intended for: colfax research webinar, 2017-04-04 (los alamos, New Mexico, United States)

Issued: 2017-04-03

---

**Disclaimer:**

Los Alamos National Laboratory, an affirmative action/equal opportunity employer, is operated by the Los Alamos National Security, LLC for the National Nuclear Security Administration of the U.S. Department of Energy under contract DE-AC52-06NA25396. By approving this article, the publisher recognizes that the U.S. Government retains nonexclusive, royalty-free license to publish or reproduce the published form of this contribution, or to allow others to do so, for U.S. Government purposes. Los Alamos National Laboratory requests that the publisher identify this article as work performed under the auspices of the U.S. Department of Energy. Los Alamos National Laboratory strongly supports academic freedom and a researcher's right to publish; as an institution, however, the Laboratory does not endorse the viewpoint of a publication or guarantee its technical correctness.

**A Case Study on  
Software Modernization  
using CoMD -  
A Molecular Dynamics  
Proxy Application**

by

Dr. Adedoyin  
toks@lanl.gov

<https://colfaxresearch.com/mc2-series/>

# About Me :

- Current :
  - Specializing in Software Modernization
  - Future Application and Architectures (FAA)
  - Application Performance Team (APT)
- Past :
  - Post doctoral Fellow (ND)
    - Developed crystal plasticity framework for c-BN Synthesis via shockwave processing
  - Computational Solid Mechanics (CSM)
    - Modeling phase transformation in solids
  - Computational Fluid Dynamics (CFD)
    - Turbulence modeling (real/spectral space)

# Gratitude:

- DoE
- LANL Institutional Computing
- Additional Contributions:
  - Cray (Hackaton/Boot-Camp)
  - Intel (Hackaton/Boot-Camp)
- LANL Folks:
  - CoMD Development Team
  - Colleagues

# This Talk I:

Performance Lingo

Pre Software Modernization

- In thinking about software modernization
- Tools - "The Necessary Evil"
- A Modernization Mindset - The Cutting Rod Approach
- Roofline Analysis - In a Nutshell
- Roofline Analysis - The Outcome

# This Talk II:

Software Modernization Case Study 2 - CoMD

CoMD:

- Algorithm Description
- Driver Blue Print
- Time Marching
- Force Kernels
- Modernization Goals
- L-J Force Kernel Optimization
- EAM Force Kernel Optimization

Post Software Modernization:

- Comparing: HSW vs. KNL, How?

# Performance Lingo:

- TDP - Thermal Design Power
- HSW - Haswell (Intel processor)
- KNL - Knights Landing (Intel Xeon Phi 2<sup>nd</sup> Generation)
- SNB - Sandy Bridge (Intel processor)
- IVY - IVY Bridge (Intel processor)
- BWD - Broadwell Bridge (Intel processor)
- SIMD - Single Instruction Multiple Data
- MPI - Message Passing Interface
- OpenMP - Open Multi-Processing (Thread Parallel Paradigm)
- Vectorization - Vector representation of similar  
scalar operations (Data Parallelization)
- Directive/Primitive -



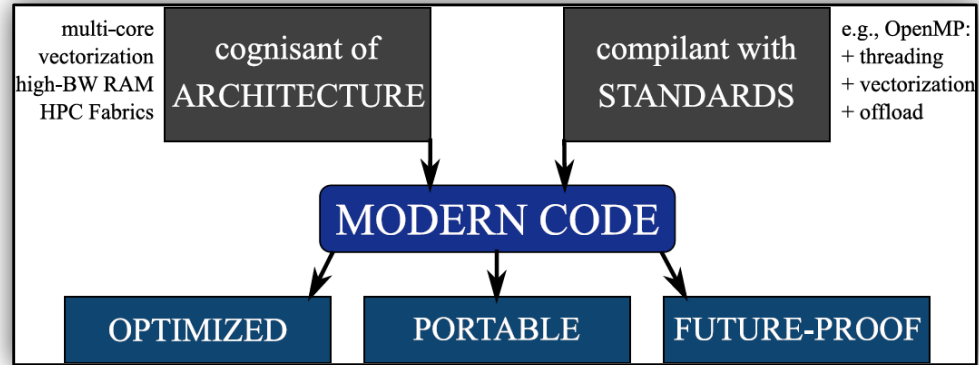
This Section discusses ...

**PRE**

**SOFTWARE MODERNIZATION**

# In Thinking about Software Modernization:

- Why? Hardware is a rate quantity:
  - $H' = f(\text{Technology})$
- Hiding Communication Primitives
  - c++: communication classes
  - c: communication source files
  - Fortran: communication modules
- Hiding Threading Primitives
  - Pragma&/Directives
  - Challenging!
- Hiding Vectorization Primitives
  - Pragma&/Directives
  - Challenging!
- Hiding Alignment/Allocation Primitives
  - aligned malloc:
    - `__aligned_malloc(,)`
    - `posix_align(,)` (still in dev.)
  - align compiler hint: `__mm_aligned(,)`
  - align size: `#ifdef <MachineType> ALIGNMENT_SIZE 64`



# Tools -

## The Necessary Evil:

- Understand
- Intel® Advisor
  - Analysis kernels
  - Roofline
- Intel® Vtune™ Amplifier
  - Analysis kernels
- Allinea (ddt/map)

# A Modernization Mindset - The Cutting Rod Approach:

- Goal:
  - Determine optimal substructure.
  - Combined optimal substructures will probably be global optimum.
- Software modernization is:
  - Iterative (pick your poison)
  - Exhaustive(work smart):
    - Use Proxy/Mini App (~20K lines)
    - Try Multiples approaches
- Note:
  - Optimal code for KNL does not translate to optimal code either HSW/SNB/BWD

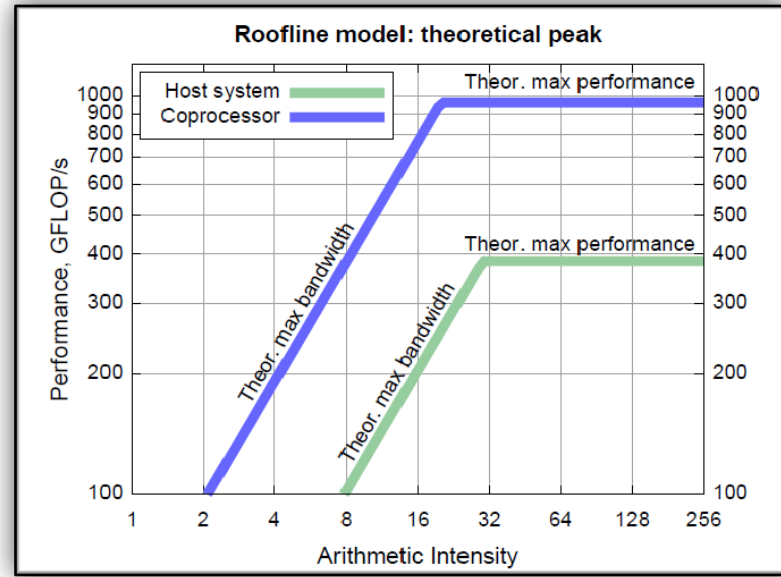
# Roofline Analysis

## - In a Nutshell:

Family of functions

Axis:

- Abscissa:
  - Flops/DRAM-Byte
  - Relative CPU to Memory utilization.
- Ordinate:
  - Flops/s



roofline model graphic - colfax research <<https://colfaxresearch.com/>>

# Roofline Analysis -

## The Outcome:

Roofline results:

- Arithmetic Intensity (AI)
  - What is thread affinity for max performance?
  - What is max number of threads for max performance?
  - Do you need to incorporate:
    - Vectorization (Relative Quality/Quantity)?
    - Prefetching (primitives/compiler flags)...?
    - Data Contiguity...?
- Does your code need:
  - *Re-timing?*
  - Dynamic thread affinity?
- Does your code need heavy rewrite to change the AI?

This Section discusses ...

# **CASE STUDY: CoMD SOFTWARE MODERNIZATION**

# CoMD – Algorithm Description:

- **CoMD** Open-Source Molecular Dynamics Proxy Application:
  - Molecular dynamics is a low level “higher resolution” material modeling approach.
  - <https://github.com/exmatex/CoMD>
- Types of Force Kernels:
  - Lenard Jones (L-J)
  - Embedded Atom Model (EAM)
- Code Branches:
  - Serial (focus for data-parallel)
  - OpenMP (focus for data & thread parallel)
    - Loop level implementation
  - MPI / MPI+OpenMP
- Problem Type:
  - N-Body
- Decomposition Type:
  - Cartesian
- Complexity:
  - $O(\sim n^2)$



# CoMD –

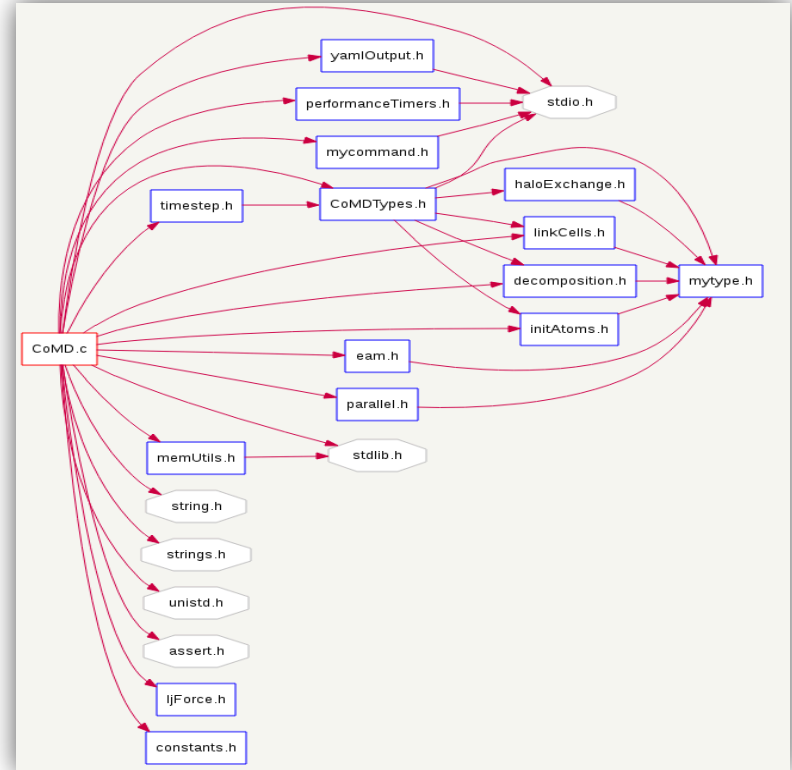
## Example Application:

- An approach to computationally fabricating c-BN via shock-wave processing will require modeling a material domain at impact.
- Upon impact, the material domain will categorically contain multiple activity regions:
  - Shock Zone – Higher activity region
  - Transition Zone – Mid activity region
  - Inert Zone – Low activity region
    - (chemically/mechanically relatively inactive)
- Molecular dynamics is a theoretically feasible approach to capturing the underlying physics (the essence) in the material's shock zone.
  - A relatively higher level mid resolution may be sufficient!
    - i.e. Crystal plasticity

# CoMD - Driver Blue-Print

## Understand View:

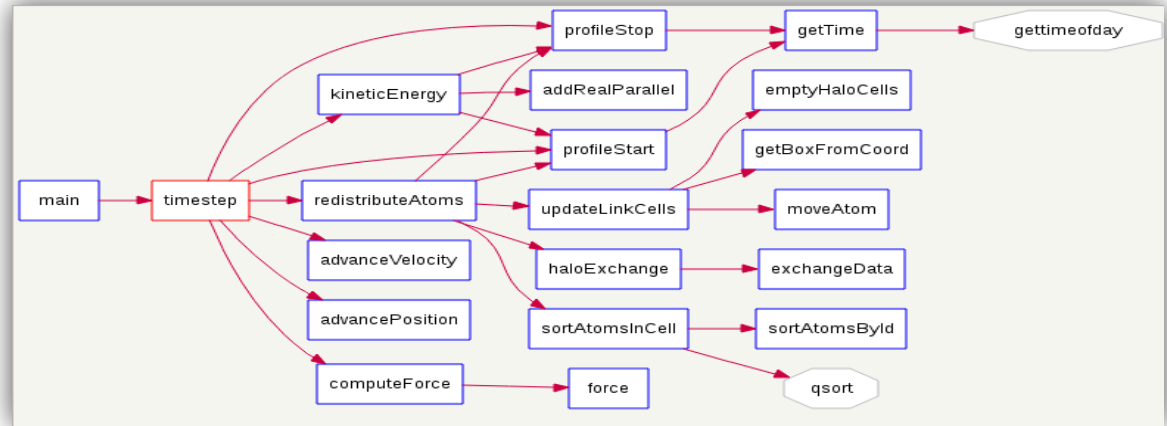
- Infinite call depth:
  - Connectivity of the code
- Blue-square imply:
  - Subroutines/function/headers
- Gray-hexagon imply:
  - c-libs



# CoMD - Time Marching:

## Understand View:

- Infinite call depth
- Blue Square imply:
  - subroutines/function
- Gray-hexagon imply:
  - c-lib



# CoMD - Force Kernels:

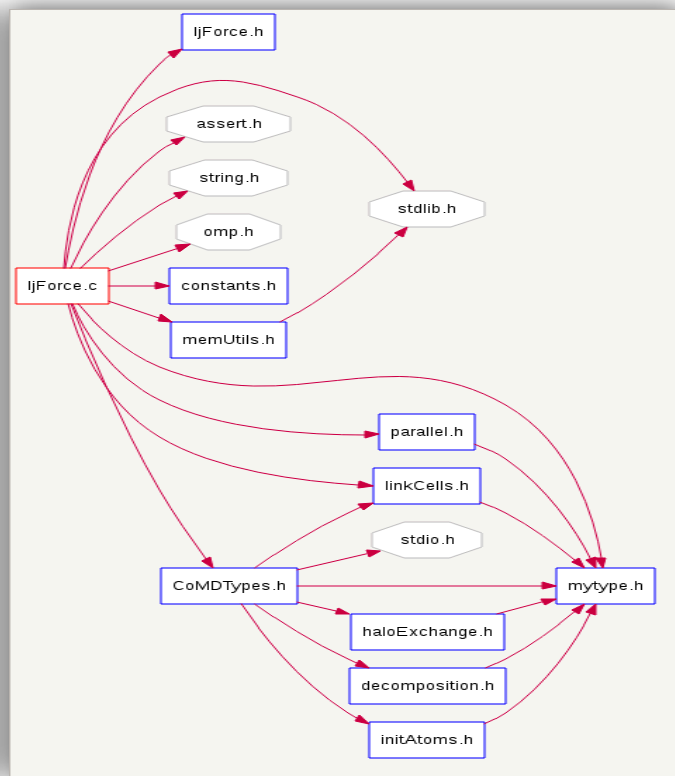
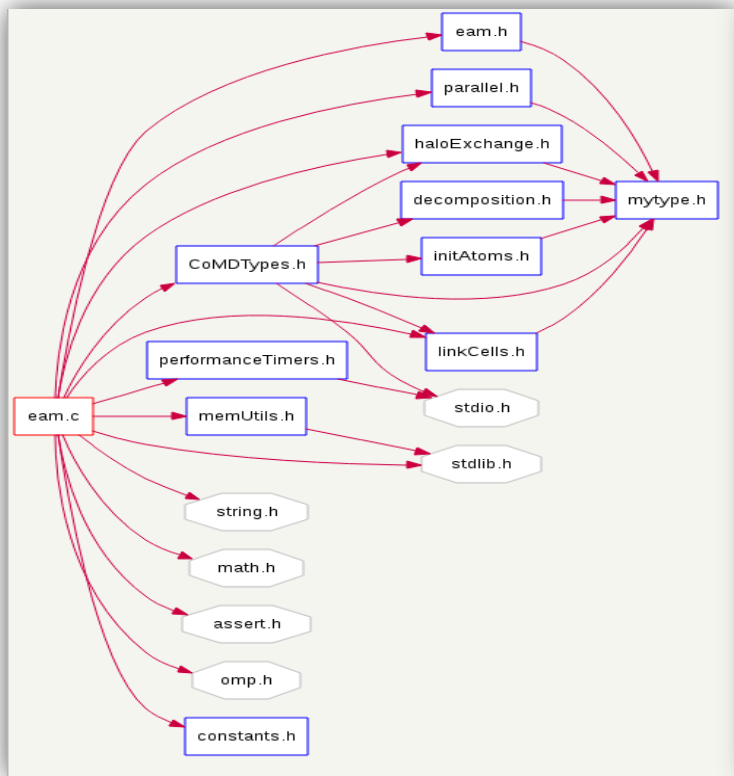
Force kernels:

•EAM:

•3-passes

•LJ:

•1-pass



# CoMD – Modernization

- Modernization Mindset:
  - “Cutting Rod Approach”
  - Optimal substructures
- Machine Choice (KNL/64-68cores/4-HT):
  - Based on Baseline runs
    - NUMA decomposition(SNC4/SNC2/QUAD)
    - Memory hierarchy(Flat/Cache)
- Modernization Exercise Goals:
  - Improve Vectorization
  - Improve Threading

# App Runtime Comparison -

## KNL vs. HSW:

How?:


- TDP is a "useful" metric
- It is "Not" an exact science!
- Given:
  - Single-Node KNL
  - ~200W CPU TDP
- Choose HSW with comparable TDP:
  - Single-Node, Dual-Socket HSW
  - 2x(Intel® Xeon® processor E5-2697v3)
  - ~2x(145W CPU TDP)
- Intel Processor Specs:
  - [ark.intel.com](http://ark.intel.com)

# CoMD - Baseline Runs I

## CoMD Simulation Details:

### KNL:

- Intel® Xeon Phi™ Processor 7210
- CoMD on Node performance
- Multiple MPI Ranks
- Complementary OpenMP Threads
- Best Results:
  - QUAD-Cache
  - 2-Threads
  - 32-Ranks
- Comparable Results:
  - QUAD-Cache
  - 1TH-64R
  - 4TH-16R

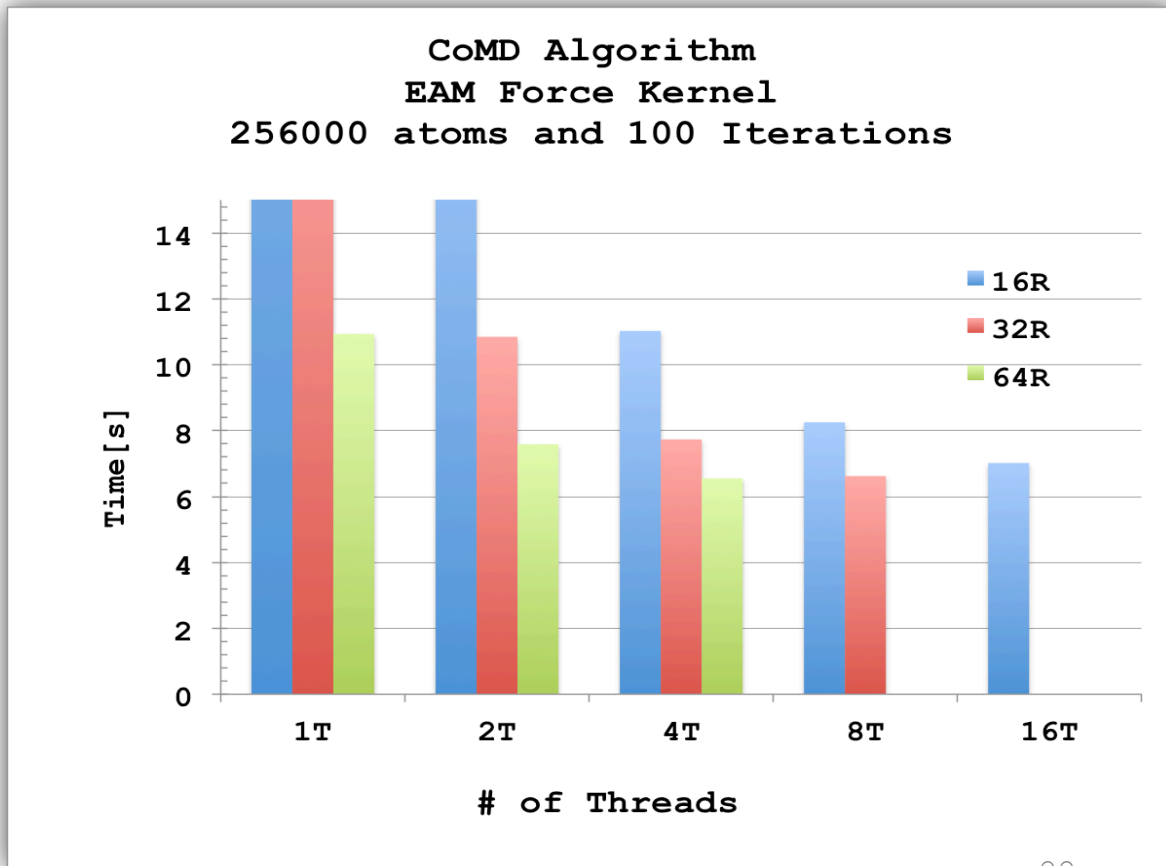


# CoMD - Baseline Runs II

## CoMD Simulation Details:

KNL:

- Intel® Xeon Phi™ Processor 7210
- CoMD on Node performance
- MPI Ranks (16/32/64)
- Complementary OpenMP Threads
- Best Results:
  - 64-Ranks
  - 32-Ranks
- Comparable Results:
  - QUAD-Cache
  - 1TH-64R
  - 4TH-16R





# CoMD - Baseline Runs III

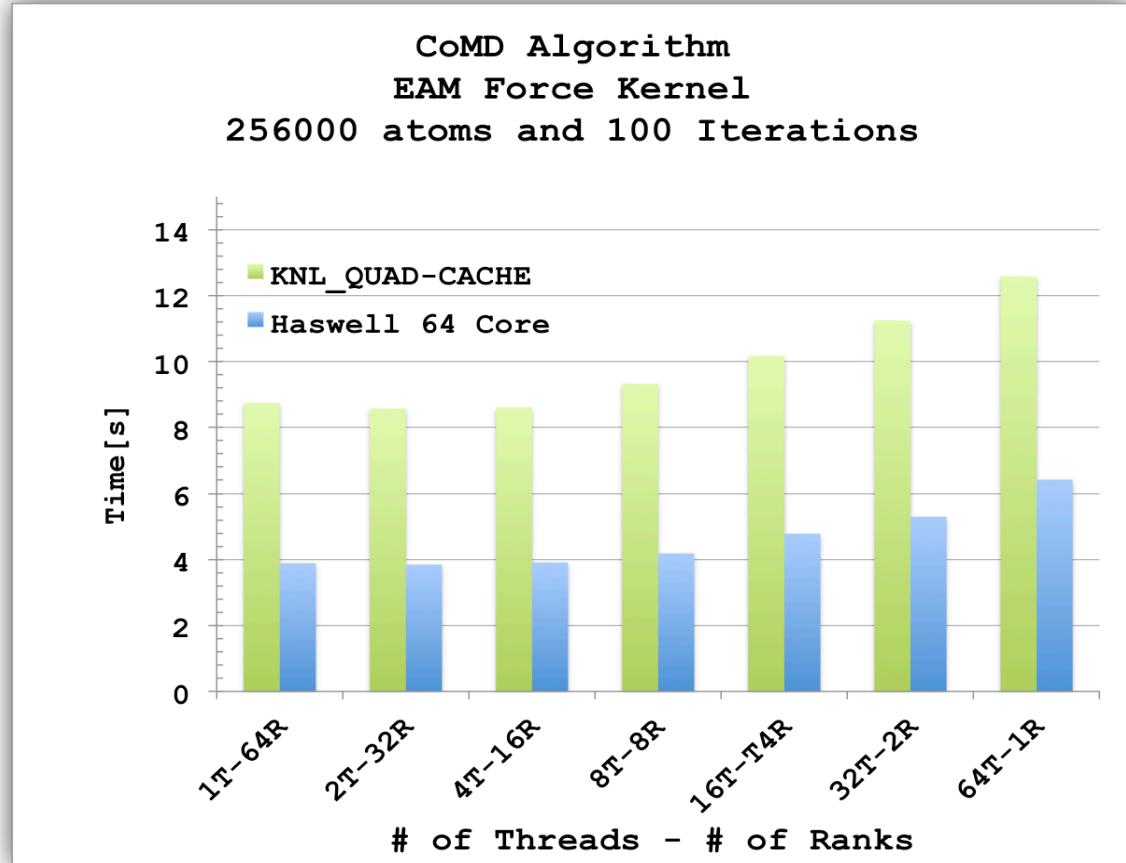
## CoMD Simulation Details:

### KNL:

- Intel® Xeon Phi™ Processor 7210

### HSW - Dual Socket:

- Intel® Xeon® Processor E5-2697v3
- CoMD on Node performance
- MPI Ranks
- Complementary OpenMP Threads
- Best Results:
  - 32-Ranks/2TH
- Note:
  - Oversubscribing
    - HSW > 2TH
    - KNL > 4TH



# CoMD - Hotspots

Intel® Vtune™ Amplifier:  
hotspots analysis

## CoMD Performance Profiling:

Intel® Vtune™ Amplifier:

- Hotspots kernel

- Summary:

CoMD with LJ Force:

- LJ-Force(,,)
- putAtomInBox(,,)

CoMD with EAM Force:

- EAM-Force(,,)
- sortAtomInCell(,,)

Summary:	L-J Force	EAM Force
<u>% Runtime Per Region</u>	Serial: ~ 2%	Serial: ~ 1%
	Parallel: ~ 98%	Parallel: ~ 99%
<u>% Per Subroutine Runtime</u>	Lenard-Jones: ~ 93%	EAM: ~ 94%
	Others: ~ 7%	Others: ~ 6%

☺ **Top Hotspots**  
This section lists the most active functions in your application. Optimize these functions to improve performance.

Function	Module	CPU Time <sup>Ⓢ</sup>
<a href="#">ljForce_V\$omp\$parallel_for@172</a>	CoMD-openmp	67.118s
<a href="#">ljForce_V\$omp\$parallel_for@157</a>	CoMD-openmp	0.820s
<a href="#">putAtomInBox</a>	CoMD-openmp	0.648s
<a href="#">sortAtomsInCell</a>	CoMD-openmp	0.490s
<a href="#">qsort_r</a>	libc.so.6	0.430s
[Others]	N/A*	1.984s

\*N/A is applied to non-summable metrics.

☺ **Top Hotspots**  
This section lists the most active functions in your application. Optimize these functions to improve performance.

Function	Module	CPU Time <sup>Ⓢ</sup>
<a href="#">eamForce_V\$omp\$parallel_for@249</a>	CoMD-openmp	59.518s
<a href="#">eamForce_V\$omp\$parallel_for@327</a>	CoMD-openmp	56.420s
<a href="#">eamForce_V\$omp\$parallel_for@238</a>	CoMD-openmp	1.649s
<a href="#">sortAtomsInCell</a>	CoMD-openmp	0.850s
<a href="#">putAtomInBox</a>	CoMD-openmp	0.749s
[Others]	N/A*	3.323s

\*N/A is applied to non-summable metrics.

# L-J Force - Simulation Output:

## CoMD Problem Definition:

Force Kernel:

- Lennard - Jones

Number of Atoms:

- 32000

Number of TimeSteps:

- 100

Parallelization:

- MPI
- OpenMP

## Build Information:

```
5 ### CoMD can be built in either double or single precision and with or
6 ### without MPI.  Select desired precision and MPI here.
7
8 # double precision (ON/OFF)
9 DOUBLE_PRECISION = ON
10 # MPI for parallel (ON/OFF)
11 DO_MPI = ON
12
13 ### Set your desired C compiler and any necessary flags.  Note that CoMD
14 ### uses some c99 features.  You can also set flags for optimization and
15 ### specify paths to include files that the compiler can't find on its
16 ### own.  If you need any -L or -l switches to get C standard libraries
17 ### (such as -lm for the math library) put them in C_LIB.
18 CC = mpiicc
19 CFLAGS = -std=c99 -qopenmp -xmic-avx512 -qopt-report=5
20 OPTFLAGS = -g -O3
21 INCLUDES =
22 C_LIB = -lm
```

## V&V Information:

Simulation Validation:

Initial energy : -1.166063303477

Final energy : -1.166049767266

eFinal/eInitial : 0.999988

Final atom count : 32000, no atoms lost

## Output:

Timing Statistics Across 1 Ranks:						
Timer	Rank:	Min(s)	Rank:	Max(s)	Avg(s)	Stdev(s)
total	0:	2.2790	0:	2.2790	2.2790	0.0000
loop	0:	2.0142	0:	2.0142	2.0142	0.0000
timestep	0:	2.0083	0:	2.0083	2.0083	0.0000
position	0:	0.0093	0:	0.0093	0.0093	0.0000
velocity	0:	0.0066	0:	0.0066	0.0066	0.0000
redistribute	0:	1.6301	0:	1.6301	1.6301	0.0000
atomHalo	0:	1.1872	0:	1.1872	1.1872	0.0000
force	0:	0.3833	0:	0.3833	0.3833	0.0000
commHalo	0:	0.1995	0:	0.1995	0.1995	0.0000
commReduce	0:	0.0005	0:	0.0005	0.0005	0.0000

Timings for Rank 0					
Timer	# Calls	Avg/Call (s)	Total (s)	% Loop	
total	1	2.2790	2.2790	113.15	
loop	1	2.0142	2.0142	100.00	
timestep	10	0.2008	2.0083	99.71	
position	100	0.0001	0.0093	0.46	
velocity	200	0.0000	0.0066	0.33	
redistribute	101	0.0161	1.6301	80.93	
atomHalo	101	0.0118	1.1872	58.94	
force	101	0.0038	0.3833	19.03	
commHalo	303	0.0007	0.1995	9.90	
commReduce	39	0.0000	0.0005	0.02	

# Source Code: CoMD

LJ-Force Kernel

Initialization Step

# L-J Force - Baseline:

## Compiler:

- Intel 17/up1

## Optimization Reports:

```
163 LOOP BEGIN at ljForce.c(157,1)
164   remark #15344: loop was not vectorized: vector dependence prevents vectorization
165   remark #15346: vector dependence: assumed OUTPUT dependence between s->atoms->f[ii] (22:4) and s->atoms->U[ii] (162:7)
166   remark #15346: vector dependence: assumed OUTPUT dependence between s->atoms->U[ii] (162:7) and s->atoms->f[ii] (22:4)
167   remark #25439: unrolled with remainder by 2
168 LOOP END
169
170 LOOP BEGIN at ljForce.c(157,1)
171 <Remainder>
172 LOOP END
173
```

## Optimization Report Summary:

Initialization Step of Baseline Code:

1. No vectorization
2. Function call in loop:
  - Line: 160
  - Function: "zeroReal3"
3. Remainder loop
4. Assumed dependence

## Source Code:

```
145 int ljForce(SimFlat* s)
146 {
147     LjPotential* pot = (LjPotential *) s->pot;
148     real_t sigma = pot->sigma;
149     real_t epsilon = pot->epsilon;
150     real_t rCut = pot->cutoff;
151     real_t rCut2 = rCut*rCut;
152
153     // zero forces and energy
154     real_t ePot = 0.0;
155     s->ePotential = 0.0;
156     int fSize = s->boxes->nTotalBoxes*MAXATOMS;
157 #pragma omp parallel for
158     for (int ii=0; ii<fSize; ++ii)
159     {
160         zeroReal3(s->atoms->f[ii]);
161         s->atoms->U[ii] = 0.;
162     }
163 }
```

# L-J Force:

## Summary of Fix List:

1. No vectorization
2. Remainder loop
3. Assumed dependence

Fixing Item 1 - "No vectorization":

Add:

```
"#pragma omp simd"
```

Or Add:

```
"#pragma omp parallel for simd"
```

Why Not?:

```
"#pragma simd"
```

# L-J Force:

## Source Code Before:

```
145 int ljForce(SimFlat* s)
146 {
147     LjPotential* pot = (LjPotential *) s->pot;
148     real_t sigma = pot->sigma;
149     real_t epsilon = pot->epsilon;
150     real_t rCut = pot->cutoff;
151     real_t rCut2 = rCut*rCut;
152
153     // zero forces and energy
154     real_t ePot = 0.0;
155     s->ePotential = 0.0;
156     int fSize = s->boxes->nTotalBoxes*MAXATOMS;
157 #pragma omp parallel for
158     for (int ii=0; ii<fSize; ++ii)
159     {
160         zeroReal3(s->atoms->f[ii]);
161         s->atoms->U[ii] = 0.;
162     }
163 }
```

## Source Code After:

```
145 int ljForce(SimFlat* s)
146 {
147     LjPotential* pot = (LjPotential *) s->pot;
148     real_t sigma = pot->sigma;
149     real_t epsilon = pot->epsilon;
150     real_t rCut = pot->cutoff;
151     real_t rCut2 = rCut*rCut;
152
153     // zero forces and energy
154     real_t ePot = 0.0;
155     s->ePotential = 0.0;
156     int fSize = s->boxes->nTotalBoxes*MAXATOMS;
157 #pragma omp simd
158     for (int ii=0; ii<fSize; ++ii)
159     {
160         zeroReal3(s->atoms->f[ii]);
161         s->atoms->U[ii] = 0.;
162     }
163 }
```

# L-J Force:

## Optimization Reports:

```
172 LOOP BEGIN at ljForce.c(158,4)
173 remark #15389: vectorization support: reference s->atoms->U[ii] has unaligned access [ ljForce.c(161,7) ]
174 remark #15381: vectorization support: unaligned access used inside loop body
175 remark #15416: vectorization support: non-unit strided store was generated for the variable <s->atoms->f[ii], stride is 3 [ mytype.h(22,4) ]
176 remark #15416: vectorization support: non-unit strided store was generated for the variable <*(a+8)>, stride is 3 [ mytype.h(23,4) ]
177 remark #15416: vectorization support: non-unit strided store was generated for the variable <*(a+16)>, stride is 3 [ mytype.h(24,4) ]
178 remark #15305: vectorization support: vector length 16
179 remark #15309: vectorization support: normalized vectorization overhead 0.079
180 remark #15301: OpenMP SIMD LOOP WAS VECTORIZED
181 remark #15442: entire loop may be executed in remainder
182 remark #15451: unmasked unaligned unit stride stores: 1
183 remark #15453: unmasked strided stores: 3
184 remark #15475: --- begin vector cost summary ---
185 remark #15476: scalar cost: 16
186 remark #15477: vector cost: 10.250
187 remark #15478: estimated potential speedup: 1.320
188 remark #15488: --- end vector cost summary ---
189 LOOP END
```

### Optimization Report Summary:

#### "Relative Quality" of Vectorization:

- scalar cost ~16
- vector cost ~10.250
- Est. Speed-Up ~1.320 (low)

#### Necessary fixes:

1. Non-SIMD-Enabled function - "zeroReal3"
2. Unaligned access:
  - "s->atom->U" structure
  - "s->atom->f" structure
3. Non-unit (3) stride store: "s->atom->f"



# L-J Force:

Fixing Item 1 - "Function call in loop":

- Convert "zeroReal3" function:
  - SIMD-Enabled function
    - Part 1: Declaration
    - Part 2: Definition

OR:

- SIMD-Enabled omp function
  - Part 1: Declaration
  - Part 2: Definition:
    - Add "notinbranch" clause
    - "omp declare simd" directive

Source Code Before:

```
20 static void zeroReal3(real3 a)
21 {
22     a[0] = 0.0;
23     a[1] = 0.0;
24     a[2] = 0.0;
25 }
```

Source Code After:

```
19
20 __attribute__((vector)) static void zeroReal3(real3 a)
21 {
22     a[0] = 0.0;
23     a[1] = 0.0;
24     a[2] = 0.0;
25 }
```

# L-J Force:

Fixing Item 2 - "Unaligned access":

- Unaligned access:
  - "s->atom->U" structure
  - "s->atom->f" structure
- How to Align:
  - Part 1: "\_mm\_malloc"
  - Part 2: "\_mm\_free"
  - Part 3: "inform compiler"
    - Before use

Aligned Allocation & Free: "memUtils.h":

```
16 static void* comdMalloc(size_t iSize)
17 {
18 #ifdef ASSUME_ALIGN
19     return _mm_malloc(iSize, _ALIGN_INT);
20 #else
21     return malloc(iSize);
22 #endif
23 }
```

```
35 static void comdFree(void *ptr)
36 {
37 #ifdef ASSUME_ALIGN
38     _mm_free(ptr);
39 #else
40     free(ptr);
41 #endif
42 }
```

```
158 #ifdef __INTEL_COMPILER
159     #define ASSUME_ALIGN
160     #define _ALIGN_INT 64
161 #endif
162
163 #ifdef ASSUME_ALIGN
164     __assume_aligned(s->atoms->f, _ALIGN_INT);
165     __assume_aligned(s->atoms->U, _ALIGN_INT);
166 #endif
```

# L-j Force:

Fixing Item 3 - "Non-unit stride store - 3":

- Convert So(AoA) to:
  - SoA\_0, SoA\_1, SoA\_2

## Source Code Before:

```
153 // zero forces and energy
154 real_t ePot = 0.0;
155 s->ePotential = 0.0;
156 int fSize = s->boxes->nTotalBoxes*MAXATOMS;
157
158 #ifdef __INTEL_COMPILER
159     #define ASSUME_ALIGN
160     #define _ALIGN_INT 64
161 #endif
162
163 #ifdef ASSUME_ALIGN
164     __assume_aligned(s->atoms->f, _ALIGN_INT);
165     __assume_aligned(s->atoms->U, _ALIGN_INT);
166 #endif
167
168 #pragma omp simd
169 for (int ii=0; ii<fSize; ++ii)
170 {
171     zeroReal3(s->atoms->f[ii]);
172     s->atoms->U[ii] = 0.;
173 }
174
```

## Source Code After:

```
153 // zero forces and energy
154 real_t ePot = 0.0;
155 s->ePotential = 0.0;
156 int fSize = s->boxes->nTotalBoxes*MAXATOMS;
157
158 #ifdef __INTEL_COMPILER
159     #define ASSUME_ALIGN
160     #define _ALIGN_INT 64
161 #endif
162
163 #ifdef ASSUME_ALIGN
164     __assume_aligned(s->atoms->f_0, _ALIGN_INT);
165     __assume_aligned(s->atoms->f_1, _ALIGN_INT);
166     __assume_aligned(s->atoms->f_2, _ALIGN_INT);
167     __assume_aligned(s->atoms->U, _ALIGN_INT);
168 #endif
169
170 #pragma omp simd
171 for (int ii=0; ii<fSize; ++ii)
172 {
173     s->atoms->f_0[ii] = 0.;
174     s->atoms->f_1[ii] = 0.;
175     s->atoms->f_2[ii] = 0.;
176     s->atoms->U[ii] = 0.;
177 }
178
```

# L-J Force:

Final Optimization Report:

1. All Aligned Access
2. No Peel/Remainder Loops
3. Relative Vectorization Quality (up ~90%)

```
150 Report from: Loop nest, Vector & Auto-parallelization optimizations [loop, vec, par]
151
152
153 LOOP BEGIN at ljForce.c(171,4)
154 remark #15388: vectorization support: reference s->atoms->f_0[ii] has aligned access [ ljForce.c(173,7) ]
155 remark #15388: vectorization support: reference s->atoms->f_1[ii] has aligned access [ ljForce.c(174,7) ]
156 remark #15388: vectorization support: reference s->atoms->f_2[ii] has aligned access [ ljForce.c(175,7) ]
157 remark #15388: vectorization support: reference s->atoms->U[ii] has aligned access [ ljForce.c(176,7) ]
158 remark #15305: vectorization support: vector length 8
159 remark #15399: vectorization support: unroll factor set to 2
160 remark #15301: OpenMP SIMD LOOP WAS VECTORIZED
161 remark #15449: unmasked aligned unit stride stores: 4
162 remark #15475: --- begin vector cost summary ---
163 remark #15476: scalar cost: 14
164 remark #15477: vector cost: 1.500
165 remark #15478: estimated potential speedup: 8.750
166 remark #15488: --- end vector cost summary ---
167 LOOP END
168
```

# Source Code: CoMD

L-J Force Kernel

Force Calculation Step

# L-J Force:

## Showing 4 Nested Loops:

```
171 // loop over local boxes
172 #pragma omp parallel for reduction(+:ePot)
173 for (int iBox=0; iBox<s->boxes->nLocalBoxes; iBox++)
174 {
175     int nIBox = s->boxes->nAtoms[iBox];
176
177     // loop over neighbors of iBox
178     for (int jTmp=0; jTmp<nNbrBoxes; jTmp++)
179     {
180         int jBox = s->boxes->nbrBoxes[iBox][jTmp];
181
182         assert(jBox>=0);
183
184         int nJBox = s->boxes->nAtoms[jBox];
185
186         // loop over atoms in iBox
187         for (int iOff=MAXATOMS*iBox; iOff<(iBox*MAXATOMS+nIBox); iOff++)
188         {
189
190             // loop over atoms in jBox
191             for (int jOff=jBox*MAXATOMS; jOff<(jBox*MAXATOMS+nJBox); jOff++)
192             {
```

## Source Code Summary:

- Force Calculation Step:
  - 4 Nested Loops
  - "omp reduction" on:
    - "ePot"
  - Imperfect nested loops
    - Loop @Line 195
    - Loop @Line 180-182
  - 2 perfectly nested loops:
    - Loop @Line 187
    - Loop @Line 191

# L-J Force:

## Innermost Loop :

```
190 // loop over atoms in jBox
191 for (int jOff=jBox*MAXATOMS; jOff<(jBox*MAXATOMS+nJBox); jOff++)
192 {
193     real3 dr;
194     real_t r2 = 0.0;
195     for (int m=0; m<3; m++)
196     {
197         dr[m] = s->atoms->r[iOff][m]-s->atoms->r[jOff][m];
198         r2+=dr[m]*dr[m];
199     }
200
201     if ( r2 <= rCut2 && r2 > 0.0)
202     {
203
204         // Important note:
205         // from this point on r actually refers to 1.0/r
206         r2 = 1.0/r2;
207         real_t r6 = s6 * (r2*r2*r2);
208         real_t eLocal = r6 * (r6 - 1.0) - eShift;
209         s->atoms->U[iOff] += 0.5*eLocal;
210         ePot += 0.5*eLocal;
211
212         // different formulation to avoid sqrt computation
213         real_t fr = - 4.0*epsilon*r6*r2*(12.0*r6 - 6.0);
214         for (int m=0; m<3; m++)
215         {
216             s->atoms->f[iOff][m] -= dr[m]*fr;
217         }
218     }
219 } // loop over atoms in jBox
```

## Source Code Summary:

- Force Calculation Step:
  - Inner most loop constains:
    - Sub-Loops @Line 195
    - Sub-Loops @Line 214
  - Omp reduction on variable
    - "ePot"

# L-J Force:

## Optimization Reports:

### Optimization Report Summary:

1. No vectorization:
  - 4 Nested-Loops
2. Imperfect loop nest
3. Assumed flow dependence:
  - Loop @187
  - Loop @191

```
LOOP BEGIN at ljForce.c(173,4)
  remark #15523: loop was not vectorized: loop control variable iBox was found, but loop iteration count cannot be computed before executing the loop

LOOP BEGIN at ljForce.c(178,7)
  remark #15520: loop was not vectorized: loop with multiple exits cannot be vectorized unless it meets search loop idiom criteria [ ljForce.c(182,10) ]

LOOP BEGIN at ljForce.c(187,10)
  remark #25096: Loop Interchange not done due to: Imperfect Loop Nest (Either at Source or due to other Compiler Transformations)
  remark #25452: Original Order found to be proper, but by a close margin
  remark #15344: loop was not vectorized: vector dependence prevents vectorization
  remark #15346: vector dependence: assumed FLOW dependence between dr[m] (197:19) and dr[m] (216:22)
  remark #15346: vector dependence: assumed ANTI dependence between dr[m] (216:22) and dr[m] (197:19)

LOOP BEGIN at ljForce.c(191,13)
  remark #15344: loop was not vectorized: vector dependence prevents vectorization
  remark #15346: vector dependence: assumed FLOW dependence between dr[m] (197:19) and dr[m] (216:22)
  remark #15346: vector dependence: assumed ANTI dependence between dr[m] (216:22) and dr[m] (197:19)

LOOP BEGIN at ljForce.c(195,16)
  remark #25436: completely unrolled by 3 (pre-vector)
LOOP END

LOOP BEGIN at ljForce.c(214,19)
  remark #25436: completely unrolled by 3 (pre-vector)
LOOP END
LOOP END
LOOP END
LOOP END
LOOP END
```



# L-J Force:

## 1. Fixing Item 1 - "No vectorization":

- Add "simd" directive to innermost loop
  - Sub-Loops: 196, 215 get unrolled
- Note:
  - There are 4 nested loops

## Optimization Reports:

```
LOOP BEGIN at ljForce.c(173,4)
remark #15523: loop was not vectorized: loop control variable iBox was found, but loop iteration count cannot be computed before executing the loop

LOOP BEGIN at ljForce.c(178,7)
remark #15520: loop was not vectorized: loop with multiple exits cannot be vectorized unless it meets search loop idiom criteria [ ljForce.c(182,10) ]

LOOP BEGIN at ljForce.c(187,10)
remark #25096: Loop Interchange not done due to: Imperfect Loop Nest (Either at Source or due to other Compiler Transformations)
remark #25452: Original Order found to be proper, but by a close margin
remark #15344: loop was not vectorized: vector dependence prevents vectorization
remark #15346: vector dependence: assumed FLOW dependence between dr[m] (198:19) and dr[m] (217:22)
remark #15346: vector dependence: assumed ANTI dependence between dr[m] (217:22) and dr[m] (198:19)

LOOP BEGIN at ljForce.c(192,13)
remark #15316: simd loop was not vectorized: scalar assignment in simd loop is prohibited, consider private, lastprivate or reduction clauses [ ljForce.c(210,19) ]
remark #15552: loop was not vectorized with "simd"

LOOP BEGIN at ljForce.c(196,16)
remark #25436: completely unrolled by 3 (pre-vector)
LOOP END

LOOP BEGIN at ljForce.c(215,19)
remark #25436: completely unrolled by 3 (pre-vector)
LOOP END
LOOP END
LOOP END
LOOP END
LOOP END
```

# L-J Force:

## Summary of Todo:

### 1. Align Data & Hint Compiler:

- "s->atoms->r"
- "s->atoms->f" ...done!

### 2. Convert So(AoA) to multiple SoA:

- "s->atoms->r"
- "s->atoms->f" ...done!

### 3. Convert reduction in inner loop on:

- "s->atoms->r" to "simd reduction"
- "s->atoms->f" to "simd reduction"
- Hint:
  - Introduce summation variable
    - "sum\_R"
    - "sum\_F"

# L-J Force

## Source Code

### Before:

```
190 // loop over atoms in jBox
191 for (int jOff=jBox*MAXATOMS; jOff<(jBox*MAXATOMS+nJBox); jOff++)
192 {
193     real3 dr;
194     real_t r2 = 0.0;
195     for (int m=0; m<3; m++)
196     {
197         dr[m] = s->atoms->r[iOff][m]-s->atoms->r[jOff][m];
198         r2+=dr[m]*dr[m];
199     }
200
201     if ( r2 <= rCut2 && r2 > 0.0)
202     {
203
204         // Important note:
205         // from this point on r actually refers to 1.0/r
206         r2 = 1.0/r2;
207         real_t r6 = s6 * (r2*r2*r2);
208         real_t eLocal = r6 * (r6 - 1.0) - eShift;
209         s->atoms->U[iOff] += 0.5*eLocal;
210         ePot += 0.5*eLocal;
211
212         // different formulation to avoid sqrt computation
213         real_t fr = - 4.0*epsilon*r6*r2*(12.0*r6 - 6.0);
214         for (int m=0; m<3; m++)
215         {
216             s->atoms->f[iOff][m] -= dr[m]*fr;
217         }
218     }
219 } // loop over atoms in jBox
```

```
192 #pragma omp parallel for default(shared) reduction(+:ePot) //private(ePot)
193 for (int iBox=0; iBox<s->boxes->nLocalBoxes; iBox++)
194 {
195     int nIBox = s->boxes->nAtoms[iBox];
196
197     // loop over neighbors of iBox
198     for (int jTmp=0; jTmp<nNbrBoxes; jTmp++)
199     {
200         int jBox = s->boxes->nbrBoxes[iBox][jTmp];
201
202         assert(jBox>=0);
203
204         int nJBox = s->boxes->nAtoms[jBox];
205         real_t sum_U;
206         real_t sum_F_0;
207         real_t sum_F_1;
208         real_t sum_F_2;
209
210         // loop over atoms in iBox
211         for (int iOff=MAXATOMS*iBox; iOff<(iBox*MAXATOMS+nIBox); iOff++)
212         {
213             #pragma simd reduction(+:sum_U,ePot) reduction(-:sum_F_0,sum_F_1,sum_F_2)
214             // loop over atoms in jBox
215             for (int jOff=jBox*MAXATOMS; jOff<(jBox*MAXATOMS+nJBox); jOff++)
216             {
217                 real3 dr;
218                 real_t r2 = 0.0;
219
220                 dr[0] = s->atoms->r_0[iOff] - s->atoms->r_0[jOff];
221                 r2+=dr[0]*dr[0];
222
223                 dr[1] = s->atoms->r_1[iOff] - s->atoms->r_1[jOff];
224                 r2+=dr[1]*dr[1];
225
226                 dr[2] = s->atoms->r_2[iOff] - s->atoms->r_2[jOff];
227                 r2+=dr[2]*dr[2];
228
229                 if ( r2 <= rCut2 && r2 > 0.0)
230                 {
231
232                     // Important note:
233                     // from this point on r actually refers to 1.0/r
234                     real_t r3 = 1.0/r2;
235                     real_t r6 = s6 * (r3*r3*r3);
236                     real_t eLocal = r6 * (r6 - 1.0) - eShift;
237                     //s->atoms->U[iOff] += 0.5*eLocal;
238                     sum_U += 0.5*eLocal;
239                     ePot += 0.5*eLocal;
240
241                     // different formulation to avoid sqrt computation
242                     real_t fr = - 4.0*epsilon*r6*r3*(12.0*r6 - 6.0);
243
244                     sum_F_0 -= dr[0]*fr;
245                     sum_F_1 -= dr[1]*fr;
246                     sum_F_2 -= dr[2]*fr;
247                 }
248             } // loop over atoms in jBox
249             s->atoms->U[iOff] += sum_U;
250             s->atoms->f_0[iOff] += sum_F_0;
251             s->atoms->f_1[iOff] += sum_F_1;
252             s->atoms->f_2[iOff] += sum_F_2;
253             sum_F_0 = 0.;
254             sum_F_1 = 0.;
255             sum_F_2 = 0.;
256             sum_U = 0.;
257         } // loop over atoms in iBox
```

## Source Code

### After:

```
228
229
230
231
232
233
234
235
236
237
238
239
240
241
242
243
244
245
246
247
248
249
250
251
252
253
254
255
256
257
```

# L-J Force - Optimized

## Performance Report":

- ~40% Speedup in Force Calculation
- ~10% Speedup overall

Old Output:

Timings for Rank 0				
Timer	# Calls	Avg/Call (s)	Total (s)	% Loop
total	1	2.2790	2.2790	113.15
loop	1	2.0142	2.0142	100.00
timestep	10	0.2008	2.0083	99.71
position	100	0.0001	0.0093	0.46
velocity	200	0.0000	0.0066	0.33
redistribute	101	0.0161	1.6301	80.93
atomHalo	101	0.0118	1.1872	58.94
force	101	0.0038	0.3833	19.03
commHalo	303	0.0007	0.2121	9.90
commReduce	39	0.0000	0.0005	0.02

Timing Statistics Across 1 Ranks:				
Timer	Rank: Min(s)	Rank: Max(s)	Avg(s)	Stdev(s)
total	0: 2.2790	0: 2.2790	2.2790	0.0000
loop	0: 2.0142	0: 2.0142	2.0142	0.0000
timestep	0: 2.0083	0: 2.0083	2.0083	0.0000
position	0: 0.0093	0: 0.0093	0.0093	0.0000
velocity	0: 0.0066	0: 0.0066	0.0066	0.0000
redistribute	0: 1.6301	0: 1.6301	1.6301	0.0000
atomHalo	0: 1.1872	0: 1.1872	1.1872	0.0000
force	0: 0.3833	0: 0.3833	0.3833	0.0000
commHalo	0: 0.1995	0: 0.1995	0.1995	0.0000
commReduce	0: 0.0005	0: 0.0005	0.0005	0.0000

New Output:

Timings for Rank 0				
Timer	# Calls	Avg/Call (s)	Total (s)	% Loop
total	1	2.1776	2.1776	116.31
loop	1	1.8723	1.8723	100.00
timestep	10	0.1866	1.8664	99.68
position	100	0.0001	0.0107	0.57
velocity	200	0.0000	0.0069	0.37
redistribute	101	0.0160	1.6136	86.18
atomHalo	101	0.0115	1.1624	62.08
force	101	0.0025	0.2545	13.59
commHalo	303	0.0006	0.1828	10.30
commReduce	39	0.0000	0.0005	0.03

Timing Statistics Across 1 Ranks:				
Timer	Rank: Min(s)	Rank: Max(s)	Avg(s)	Stdev(s)
total	0: 2.1776	0: 2.1776	2.1776	0.0000
loop	0: 1.8723	0: 1.8723	1.8723	0.0000
timestep	0: 1.8664	0: 1.8664	1.8664	0.0000
position	0: 0.0107	0: 0.0107	0.0107	0.0000
velocity	0: 0.0069	0: 0.0069	0.0069	0.0000
redistribute	0: 1.6136	0: 1.6136	1.6136	0.0000
atomHalo	0: 1.1624	0: 1.1624	1.1624	0.0000
force	0: 0.2545	0: 0.2545	0.2545	0.0000
commHalo	0: 0.1928	0: 0.1928	0.1928	0.0000
commReduce	0: 0.0005	0: 0.0005	0.0005	0.0000

# L-J Force - Revision:

## Source Code Before:

```
157 #pragma omp parallel for
158   for (int ii=0; ii<fSize; ++ii)
159   {
160       zeroReal3(s->atoms->f[ii]);
161       s->atoms->U[ii] = 0.;
162   }
163
```

## Source Code After Revision I:

```
157 #pragma omp simd
158   for (int ii=0; ii<fSize; ++ii)
159   {
160       zeroReal3(s->atoms->f[ii]);
161       s->atoms->U[ii] = 0.;
162   }
163
```

## Source Code After Revision II:

```
169
170 #pragma omp parallel for simd
171   for (int ii=0; ii<fSize; ++ii)
172   {
173       s->atoms->f_0[ii] = 0.;
174       s->atoms->f_1[ii] = 0.;
175       s->atoms->f_2[ii] = 0.;
176       s->atoms->U[ii] = 0.;
177   }
```

# L-J Force - Optimized

## Revised Performance Report":

- Revised Initialization Step
- ~51% Speedup in Force Calculation
- ~10% Speedup overall

Old Output:

Timings for Rank 0				
Timer	# Calls	Avg/Call (s)	Total (s)	% Loop
total	1	2.2790	2.2790	113.15
loop	1	2.0142	2.0142	100.00
timestep	10	0.2008	2.0083	99.71
position	100	0.0001	0.0093	0.46
velocity	200	0.0000	0.0066	0.33
redistribute	101	0.0161	1.6301	80.93
atomHalo	101	0.0118	1.1872	58.94
force	101	0.0038	0.3833	19.03
commHalo	303	0.0007	0.2121	9.90
commReduce	39	0.0000	0.0005	0.02

Timing Statistics Across 1 Ranks:				
Timer	Rank: Min(s)	Rank: Max(s)	Avg(s)	Stdev(s)
total	0: 2.2790	0: 2.2790	2.2790	0.0000
loop	0: 2.0142	0: 2.0142	2.0142	0.0000
timestep	0: 2.0083	0: 2.0083	2.0083	0.0000
position	0: 0.0093	0: 0.0093	0.0093	0.0000
velocity	0: 0.0066	0: 0.0066	0.0066	0.0000
redistribute	0: 1.6301	0: 1.6301	1.6301	0.0000
atomHalo	0: 1.1872	0: 1.1872	1.1872	0.0000
force	0: 0.3833	0: 0.3833	0.3833	0.0000
commHalo	0: 0.1995	0: 0.1995	0.1995	0.0000
commReduce	0: 0.0005	0: 0.0005	0.0005	0.0000

New Output:

Timings for Rank 0				
Timer	# Calls	Avg/Call (s)	Total (s)	% Loop
total	1	2.0499	2.0499	112.82
loop	1	1.8170	1.8170	100.00
timestep	10	0.1811	1.8108	99.66
position	100	0.0001	0.0102	0.56
velocity	200	0.0000	0.0063	0.35
redistribute	101	0.0160	1.6190	89.10
atomHalo	101	0.0117	1.1825	65.08
force	101	0.0019	0.1947	10.72
commHalo	303	0.0007	0.2121	10.90
commReduce	39	0.0000	0.0005	0.03

Timing Statistics Across 1 Ranks:				
Timer	Rank: Min(s)	Rank: Max(s)	Avg(s)	Stdev(s)
total	0: 2.0499	0: 2.0499	2.0499	0.0000
loop	0: 1.8170	0: 1.8170	1.8170	0.0000
timestep	0: 1.8108	0: 1.8108	1.8108	0.0000
position	0: 0.0102	0: 0.0102	0.0102	0.0000
velocity	0: 0.0063	0: 0.0063	0.0063	0.0000
redistribute	0: 1.6190	0: 1.6190	1.6190	0.0000
atomHalo	0: 1.1825	0: 1.1825	1.1825	0.0000
force	0: 0.1947	0: 0.1947	0.1947	0.0000
commHalo	0: 0.1981	0: 0.1981	0.1981	0.0000
commReduce	0: 0.0005	0: 0.0005	0.0005	0.0000

# Source Code: CoMD

EAM-Force Kernel

Force Calculation Step

# EAM Force - Baseline Simulation Output:

## Problem Definitions:

- 1Million Particles
- 100 Time-steps
- Mpi + OpenMP

## V&V Information:

Simulation Validation:	
Initial energy	: -3.460523233094
Final energy	: -3.460530068412
eFinal/eInitial	: 1.000002
Final atom count	: 1024000, no atoms lost

## Output:

Timing Statistics Across 32 Ranks:						
Timer	Rank:	Min(s)	Rank:	Max(s)	Avg(s)	Stdev(s)
total	17:	21.8428	18:	21.8431	21.8430	0.0001
loop	0:	21.5517	27:	21.5518	21.5518	0.0000
timestep	0:	21.5504	11:	21.5514	21.5514	0.0002
position	29:	0.0503	26:	0.0570	0.0542	0.0016
velocity	26:	0.0952	29:	0.1044	0.1005	0.0023
redistribute	22:	1.6953	15:	1.8536	1.7928	0.0409
atomHalo	22:	1.1331	15:	1.2960	1.2376	0.0414
force	15:	19.6944	22:	19.9278	19.7831	0.0656
eamHalo	29:	0.1215	26:	0.3058	0.2114	0.0598
commHalo	22:	0.4628	26:	0.6899	0.5920	0.0409
commReduce	22:	0.0226	8:	0.0997	0.0646	0.0255

Timings for Rank 0					
Timer	# Calls	Avg/Call (s)	Total (s)	% Loop	
total	1	21.8431	21.8431	101.35	
loop	1	21.5517	21.5517	100.00	
timestep	10	2.1550	21.5504	99.99	
position	100	0.0005	0.0524	0.24	
velocity	200	0.0005	0.1013	0.47	
redistribute	101	0.0179	1.8107	8.40	
atomHalo	101	0.0124	1.2524	5.81	
force	101	0.1954	19.7399	91.59	
eamHalo	101	0.0014	0.1414	0.67	
commHalo	606	0.0009	0.5417	2.51	
commReduce	39	0.0024	0.0929	0.43	



# EAM Force:

## Summary of Todo List I:

### 1. Align Data & Hint Compiler:

- pot->phi
- pot->rho
- s->atoms->U
- pot->dfEmbed
- pot->rhobar
- "s->atoms->r"
- "s->atoms->f" ...done!

### 2. Convert So(AoA) to multiple SoA:

- "s->atoms->r"
- "s->atoms->f" ...done!

### 3. Convert reduction in inner loop on:

- "s->atoms->r" to "simd reduction"
- "s->atoms->f" to "simd reduction"
- Hint:
  - Introduce summation variable
    - "sum\_R"
    - "sum\_F"

# EAM Force:

## Summary of Todo II:

1. Convert "interpolate" subroutine to:

- SIMD-Enabled function
  - Part 1: Declaration
  - Part 2: Definition

Or:

- "omp" SIMD-Enabled function
  - Part 1: Declaration
  - Part 2: Definition
    - "omp declare simd"
    - "align clause"
    - "inbranch clause"

# EAM Force:

## Summary of Todo III:

Hint compiler on alignment "eam.c":

1. "table->values" variable:

- @Function
  - InterpolationObject\*  
initInterpolationObject


1. "buf" variable:

- @Function
  - void eamReadSetfl
  - void eamReadFuncfl

# EAM Force:

## Source Code Before:

```
233  real_t rCut2 = pot->cutoff*pot->cutoff;
234  real_t etot = 0.;
235
236  // zero forces / energy / rho / rhoprime
237  int fsize = s->boxes->nTotalBoxes*MAXATOMS;
238  #pragma omp parallel for
239  for (int ii=0; ii<fsize; ii++)
240  {
241      zeroReal3(s->atoms->f[ii]);
242      s->atoms->U[ii] = 0.;
243      pot->dfEmbed[ii] = 0.;
244      pot->rhobar[ii] = 0.;
245  }
246
247  int nNbrBoxes = 27;
248  // loop over local boxes
```



## Source Code After:

```
238  int fsize = s->boxes->nTotalBoxes*MAXATOMS;
239
240  #ifdef __INTEL_COMPILER
241      #define ASSUME_ALIGN
242      #define _ALIGN_INT 64
243  #endif
244
245  #ifdef ASSUME_ALIGN
246      __assume_aligned(s->atoms->f_0, _ALIGN_INT);
247      __assume_aligned(s->atoms->f_1, _ALIGN_INT);
248      __assume_aligned(s->atoms->f_2, _ALIGN_INT);
249      __assume_aligned(s->atoms->U, _ALIGN_INT);
250      __assume_aligned(pot->dfEmbed, _ALIGN_INT);
251      __assume_aligned(pot->rhobar, _ALIGN_INT);
252  #endif
253
254  #pragma omp simd
255  for (int ii=0; ii<fsize; ii++)
256  {
257      s->atoms->f_0[ii] = 0.;
258      s->atoms->f_1[ii] = 0.;
259      s->atoms->f_2[ii] = 0.;
260      s->atoms->U[ii] = 0.;
261      pot->dfEmbed[ii] = 0.;
262      pot->rhobar[ii] = 0.;
263  }
264
```

## Modifications - Initialization step:

- Aligned memory allocation
  - Part 1: “\_\_mm\_malloc”
  - Part 2: “hint compiler”
- 1X(SoA(oA)) => 3X(SoA)


# EAM Force

Source Code Before:

Modifications - Force Calculation Pass I:

- Simd reduction
  - s->atom->U
  - s->atom->F\_(all)
  - pot->rhobar

```
248 // loop over local boxes
249 #pragma omp parallel for reduction(+:etot)
250 for (int iBox=0; iBox<s->boxes->nLocalBoxes; iBox++)
251 {
252     int nIBox = s->boxes->nAtoms[iBox];
253
254     // loop over neighbor boxes of iBox (some may be halo boxes)
255     for (int jTmp=0; jTmp<nNrBoxes; jTmp++)
256     {
257         int jBox = s->boxes->nbrBoxes[iBox][jTmp];
258         int nJBox = s->boxes->nAtoms[jBox];
259
260         // loop over atoms in iBox
261         for (int iOff=MAXATOMS*iBox; iOff<(iBox*MAXATOMS+nIBox); iOff++)
262         {
263             // loop over atoms in jBox
264             for (int jOff=MAXATOMS*jBox; jOff<(jBox*MAXATOMS+nJBox); jOff++)
265             {
266
267                 real3 dr;
268                 real_t r2 = 0.0;
269                 for (int k=0; k<3; k++)
270                 {
271                     dr[k]=s->atoms->r[iOff][k]-s->atoms->r[jOff][k];
272                     r2+=dr[k]*dr[k];
273                 }
274
275                 if(r2 <= rCut2 && r2 > 0.0)
276                 {
277                     real_t r = sqrt(r2);
278
279                     real_t phiTmp, dPhi, rhoTmp, dRho;
280                     interpolate(pot->phi, r, &phiTmp, &dPhi);
281                     interpolate(pot->rho, r, &rhoTmp, &dRho);
282
283
284                     for (int k=0; k<3; k++)
285                     {
286                         s->atoms->f[iOff][k] -= dPhi*dr[k]/r;
287                     }
288
289                     // Calculate energy contribution
290                     s->atoms->U[iOff] += 0.5*phiTmp;
291                     etot += 0.5*phiTmp;
292
293                     // accumulate rhobar for each atom
294                     pot->rhobar[iOff] += rhoTmp;
295                 }
296
297             } // loop over atoms in jBox
298         } // loop over atoms in iBox
299     } // loop over neighbor boxes
300 } // loop over local boxes
301
```



```
264 #ifndef ASSUME_ALIGN
265     __assume_aligned(s->atoms->r_0, ALIGN_INT);
266     __assume_aligned(s->atoms->r_1, ALIGN_INT);
267     __assume_aligned(s->atoms->r_2, ALIGN_INT);
268     __assume_aligned(pot->phi, ALIGN_INT);
269     __assume_aligned(pot->rho, ALIGN_INT);
270     __assume_aligned(pot->f, ALIGN_INT);
271 #endif
272 #endif
273 InterpolationObject* table_Phi = pot->phi;
274 InterpolationObject* table_Rho = pot->rho;
275
276 int nNrBoxes = 27;
277 // loop over local boxes
278 #pragma omp parallel for reduction(+:etot)
279 #pragma omp parallel for default(shared) reduction(+)
280 for (int iBox=0; iBox<s->boxes->nLocalBoxes; iBox++)
281 {
282     int nIBox = s->boxes->nAtoms[iBox];
283
284     // loop over neighbor boxes of iBox (some may be
285     for (int jTmp=0; jTmp<nNrBoxes; jTmp++)
286     {
287         int jBox = s->boxes->nbrBoxes[iBox][jTmp];
288         int nJBox = s->boxes->nAtoms[jBox];
289
290         real_t sum_U;
291         real_t sum_Rho;
292         real_t sum_F_0;
293         real_t sum_F_1;
294         real_t sum_F_2;
295         // loop over atoms in iBox
296         for (int iOff=MAXATOMS*iBox; iOff<(iBox*MAXATOMS+nIBox); iOff++)
297         {
298             // loop over atoms in jBox
299             #pragma simd reduction(+:sum_U,sum_Rho,sum_F_0,sum_F_1,sum_F_2)
300             for (int jOff=MAXATOMS*jBox; jOff<(jBox*MAXATOMS+nJBox); jOff++)
301             {
302
303                 real3 dr;
304                 real_t r2 = 0.0;
305
306                 dr[0] = s->atoms->r_0[iOff] - s->atoms->r_0[jOff];
307                 r2+=dr[0]*dr[0];
308                 dr[1] = s->atoms->r_1[iOff] - s->atoms->r_1[jOff];
309                 r2+=dr[1]*dr[1];
310                 dr[2] = s->atoms->r_2[iOff] - s->atoms->r_2[jOff];
311                 r2+=dr[2]*dr[2];
312
313                 if(r2 <= rCut2 && r2 > 0.0)
314                 {
315
```

```
313         if(r2 <= rCut2 && r2 > 0.0)
314         {
315
316             real_t r = sqrt(r2);
317
318             real_t phiTmp, dPhi, rhoTmp, dRho;
319             interpolate(table_Phi, r, &phiTmp, &dPhi);
320             interpolate(table_Rho, r, &rhoTmp, &dRho);
321
322             real_t dPhibyR = dPhi/r;
323             sum_F_0 += dPhibyR*dr[0];
324             sum_F_1 += dPhibyR*dr[1];
325             sum_F_2 += dPhibyR*dr[2];
326
327             // Calculate energy contribution
328             sum_U += 0.5*phiTmp;
329             etot += 0.5*phiTmp;
330             // accumulate rhobar for each atom
331             sum_Rho += rhoTmp;
332         }
333     } // loop over atoms in jBox
334
335     s->atoms->f_0[iOff] -= sum_F_0;
336     s->atoms->f_1[iOff] -= sum_F_1;
337     s->atoms->f_2[iOff] -= sum_F_2;
338     s->atoms->U[iOff] += sum_U;
339     pot->rhobar[iOff] += sum_Rho;
340     sum_F_0 = 0.;
341     sum_F_1 = 0.;
342     sum_F_2 = 0.;
343     sum_Rho = 0.;
344     sum_U = 0.;
345 } // loop over atoms in iBox
346 } // loop over neighbor boxes
347 } // loop over local boxes
348
```


# EAM Force:

## Modifications - Force Calculation Pass II:

- Added "pragma simd" directive
  - s->atom->U

### Source Code Before:

```
302 // Compute Embedding Energy
303 // loop over all local boxes
304 #pragma omp parallel for reduction(+:etot)
305 for (int iBox=0; iBox<s->boxes->nLocalBoxes; iBox++)
306 {
307     int nIBox = s->boxes->nAtoms[iBox];
308
309     // loop over atoms in iBox
310     for (int iOff=MAXATOMS*iBox; iOff<(MAXATOMS*iBox+nIBox); iOff++)
311     {
312         real_t fEmbed, dfEmbed;
313         interpolate(pot->f, pot->rhobar[iOff], &fEmbed, &dfEmbed);
314         pot->dfEmbed[iOff] = dfEmbed; // save derivative for halo exchange
315         s->atoms->U[iOff] += fEmbed;
316         etot += fEmbed;
317     }
318 }
```



### Source Code After:

```
350 InterpolationObject* table_F = pot->f;
351
352 // Compute Embedding Energy
353 // loop over all local boxes
354 #pragma omp parallel for reduction(+:etot)
355 for (int iBox=0; iBox<s->boxes->nLocalBoxes; iBox++)
356 {
357     int nIBox = s->boxes->nAtoms[iBox];
358
359     // loop over atoms in iBox
360     #pragma simd
361     for (int iOff=MAXATOMS*iBox; iOff<(MAXATOMS*iBox+nIBox); iOff++)
362     {
363         real_t fEmbed, dfEmbed;
364         interpolate(table_F, pot->rhobar[iOff], &fEmbed, &dfEmbed);
365         pot->dfEmbed[iOff] = dfEmbed; // save derivative for halo exchange
366         s->atoms->U[iOff] += fEmbed;
367         etot += fEmbed;
368     }
369 }
```


# EAM Force:

## Modifications - Force Calculation Pass III:

- Simd reduction: "s->atoms->f(All)"

### Source Code Before:

```
325 // third pass
326 // loop over local boxes
327 #pragma omp parallel for
328 for (int iBox=0; iBox<s->boxes->nLocalBoxes; iBox++)
329 {
330     int nIBox = s->boxes->nAtoms[iBox];
331
332     // loop over neighbor boxes of iBox (some may be halo boxes)
333     for (int jTmp=0; jTmp<nNbrBoxes; jTmp++)
334     {
335         int jBox = s->boxes->nbrBoxes[iBox][jTmp];
336         int nJBox = s->boxes->nAtoms[jBox];
337
338         // loop over atoms in iBox
339         for (int iOff=MAXATOMS*iBox; iOff<(MAXATOMS*iBox+nIBox); iOff++)
340         {
341             // loop over atoms in jBox
342             for (int jOff=MAXATOMS*jBox; jOff<(MAXATOMS*jBox+nJBox); jOff++)
343             {
344
345                 real_t r2 = 0.0;
346                 real3 dr;
347                 for (int k=0; k<3; k++)
348                 {
349                     dr[k]=s->atoms->r[iOff][k]-s->atoms->r[jOff][k];
350                     r2+=dr[k]*dr[k];
351                 }
352
353                 if(r2 <= rCut2 && r2 > 0.0)
354                 {
355                     real_t r = sqrt(r2);
356                     real_t rhoTmp, dRho;
357                     interpolate(pot->rho, r, &rhoTmp, &dRho);
358
359                     for (int k=0; k<3; k++)
360                     {
361                         s->atoms->f[iOff][k] -= (pot->dfEmbed[iOff]+pot->dfEmbed[jOff])*dRho*dr[k]/r;
362                     }
363
364                 }
365             }
366         } // loop over atoms in jBox
367     } // loop over atoms in iBox
368 } // loop over neighbor boxes
369 } // loop over local boxes
370
```



### Source Code After:

```
377 // third pass
378 // loop over local boxes
379 #pragma omp parallel for //private(table_Rho)
380 for (int iBox=0; iBox<s->boxes->nLocalBoxes; iBox++)
381 {
382     int nIBox = s->boxes->nAtoms[iBox];
383
384     // loop over neighbor boxes of iBox (some may be halo boxes)
385     for (int jTmp=0; jTmp<nNbrBoxes; jTmp++)
386     {
387         int jBox = s->boxes->nbrBoxes[iBox][jTmp];
388         int nJBox = s->boxes->nAtoms[jBox];
389         real_t sum_F_0;
390         real_t sum_F_1;
391         real_t sum_F_2;
392         // loop over atoms in iBox
393         for (int iOff=MAXATOMS*iBox; iOff<(MAXATOMS*iBox+nIBox); iOff++)
394         {
395             // loop over atoms in jBox
396             #pragma omp simd reduction(+:sum_F_0,sum_F_1,sum_F_2)
397             for (int jOff=MAXATOMS*jBox; jOff<(MAXATOMS*jBox+nJBox); jOff++)
398             {
399
400                 real_t r2 = 0.0;
401                 real3 dr;
402
403                 dr[0] = s->atoms->r_0[iOff] - s->atoms->r_0[jOff];
404                 r2+=dr[0]*dr[0];
405                 dr[1] = s->atoms->r_1[iOff] - s->atoms->r_1[jOff];
406                 r2+=dr[1]*dr[1];
407                 dr[2] = s->atoms->r_2[iOff] - s->atoms->r_2[jOff];
408                 r2+=dr[2]*dr[2];
409
410                 if(r2 <= rCut2 && r2 > 0.0)
411                 {
412                     real_t r = sqrt(r2);
413
414                     real_t rhoTmp, dRho;
415                     interpolate(pot->rho, r, &rhoTmp, &dRho);
416
417                     real_t dRhobyx = dRho/r;
418                     sum_F_0 += (pot->dfEmbed[iOff]+pot->dfEmbed[jOff])*dRhobyx*dr[0];
419                     sum_F_1 += (pot->dfEmbed[iOff]+pot->dfEmbed[jOff])*dRhobyx*dr[1];
420                     sum_F_2 += (pot->dfEmbed[iOff]+pot->dfEmbed[jOff])*dRhobyx*dr[2];
421                 }
422             } // loop over atoms in jBox
423         }
424         s->atoms->f_0[iOff] -= sum_F_0;
425         s->atoms->f_1[iOff] -= sum_F_1;
426         s->atoms->f_2[iOff] -= sum_F_2;
427         sum_F_0 = 0.;
428         sum_F_1 = 0.;
429         sum_F_2 = 0.;
430     } // loop over atoms in iBox
431 } // loop over neighbor boxes
432 } // loop over local boxes
433
```

# EAM Force - Optimized:

## Performance Report":

- ~27% Speedup in EAM Force kernel
- ~28% Speedup overall

Old Output:

Timing Statistics Across 32 Ranks:						
Timer	Rank:	Min(s)	Rank:	Max(s)	Avg(s)	Stdev(s)
total	17:	21.8428	18:	21.8431	21.8430	0.0001
loop	0:	21.5517	27:	21.5518	21.5518	0.0000
timestep	0:	21.5504	11:	21.5514	21.5514	0.0002
position	29:	0.0503	26:	0.0570	0.0542	0.0016
velocity	26:	0.0952	29:	0.1044	0.1005	0.0023
redistribute	22:	1.6953	15:	1.8536	1.7928	0.0409
atomHalo	22:	1.1331	15:	1.2960	1.2376	0.0414
force	15:	19.6944	22:	19.9278	19.7831	0.0656
eamHalo	29:	0.1215	26:	0.3058	0.2114	0.0598
commHalo	22:	0.4628	26:	0.6899	0.5920	0.0409
commReduce	22:	0.0226	8:	0.0997	0.0646	0.0255

Timings for Rank 0				
Timer	# Calls	Avg/Call (s)	Total (s)	% Loop
total	1	21.8431	21.8431	101.35
loop	1	21.5517	21.5517	100.00
timestep	10	2.1550	21.5504	99.99
position	100	0.0005	0.0524	0.24
velocity	200	0.0005	0.1013	0.47
redistribute	101	0.0179	1.8107	8.40
atomHalo	101	0.0124	1.2333	5.81
force	101	0.1954	19.7399	91.59
eamHalo	101	0.0014	0.1414	0.67
commHalo	606	0.0009	0.5417	2.51
commReduce	39	0.0024	0.0929	0.43

New Output:

Timing Statistics Across 32 Ranks:						
Timer	Rank:	Min(s)	Rank:	Max(s)	Avg(s)	Stdev(s)
total	18:	16.0573	0:	16.0577	16.0575	0.0001
loop	0:	15.8279	18:	15.8279	15.8279	0.0000
timestep	0:	15.8266	29:	15.8276	15.8275	0.0002
position	21:	0.0562	11:	0.0598	0.0581	0.0007
velocity	11:	0.0898	29:	0.0963	0.0936	0.0013
redistribute	22:	1.6353	20:	1.7249	1.6868	0.0230
atomHalo	22:	1.0884	20:	1.1809	1.1431	0.0231
force	20:	14.1193	22:	14.2059	14.1581	0.0224
eamHalo	4:	0.1230	11:	0.2064	0.1627	0.0224
commHalo	22:	0.4234	11:	0.5417	0.4845	0.0243
commReduce	3:	0.0060	14:	0.0199	0.0121	0.0038

Timings for Rank 0				
Timer	# Calls	Avg/Call (s)	Total (s)	% Loop
total	1	16.0577	16.0577	101.45
loop	1	15.8279	15.8279	100.00
timestep	10	1.5827	15.8266	99.99
position	100	0.0006	0.0581	0.37
velocity	200	0.0005	0.0944	0.60
redistribute	101	0.0168	1.6958	10.71
atomHalo	101	0.0114	1.1431	7.26
force	101	0.1401	14.1507	89.40
eamHalo	101	0.0014	0.1414	0.87
commHalo	606	0.0007	0.4415	2.79
commReduce	39	0.0003	0.0116	0.07



# L-J Force - Revision:

Source Code Before:

```
238 #pragma omp parallel for
239 for (int ii=0; ii<fsize; ii++)
240 {
241     zeroReal3(s->atoms->f[ii]);
242     s->atoms->U[ii] = 0.;
243     pot->dfEmbed[ii] = 0.;
244     pot->rhobar[ii] = 0.;
245 }
```

Source Code After Revision II:

```
254 #pragma omp parallel for simd
255 for (int ii=0; ii<fsize; ii++)
256 {
257     s->atoms->f_0[ii] = 0.;
258     s->atoms->f_1[ii] = 0.;
259     s->atoms->f_2[ii] = 0.;
260     s->atoms->U[ii] = 0.;
261     pot->dfEmbed[ii] = 0.;
262     pot->rhobar[ii] = 0.;
263 }
```

# EAM Force - Optimized

## Revised Performance Report":

- Revised Initialization Step
- ~40% Speedup in EAM Force Kernel
- ~32% Speedup overall

Old Output:

Timing Statistics Across 32 Ranks:						
Timer	Rank:	Min(s)	Rank:	Max(s)	Avg(s)	Stdev(s)
total	17:	21.8428	18:	21.8431	21.8430	0.0001
loop	0:	21.5517	27:	21.5518	21.5518	0.0000
timestep	0:	21.5504	11:	21.5514	21.5514	0.0002
position	29:	0.0503	26:	0.0570	0.0542	0.0016
velocity	26:	0.0952	29:	0.1044	0.1005	0.0023
redistribute	22:	1.6953	15:	1.8536	1.7928	0.0409
atomHalo	22:	1.1331	15:	1.2960	1.2376	0.0414
force	15:	19.6944	22:	19.9278	19.7831	0.0656
eamHalo	29:	0.1215	26:	0.3058	0.2114	0.0598
commHalo	22:	0.4628	26:	0.6899	0.5920	0.0409
commReduce	22:	0.0226	8:	0.0997	0.0646	0.0255

Timings for Rank 0				
Timer	# Calls	Avg/Call (s)	Total (s)	% Loop
total	1	21.8431	21.8431	101.35
loop	1	21.5517	21.5517	100.00
timestep	10	2.1550	21.5504	99.99
position	100	0.0005	0.0524	0.24
velocity	200	0.0005	0.1013	0.47
redistribute	101	0.0179	1.8107	8.40
atomHalo	101	0.0124	1.2323	5.81
force	101	0.1954	19.7399	91.59
eamHalo	101	0.0014	0.1414	0.67
commHalo	606	0.0009	0.5417	2.51
commReduce	39	0.0024	0.0929	0.43

New Output:

Timing Statistics Across 32 Ranks:						
Timer	Rank:	Min(s)	Rank:	Max(s)	Avg(s)	Stdev(s)
total	6:	14.8587	0:	14.8650	14.8591	0.0011
loop	0:	14.5427	12:	14.5433	14.5433	0.0001
timestep	0:	14.5299	3:	14.5423	14.5419	0.0022
position	0:	0.0450	15:	0.0522	0.0491	0.0017
velocity	15:	0.0744	27:	0.0830	0.0790	0.0023
redistribute	22:	2.4576	8:	2.6951	2.5706	0.0843
atomHalo	22:	1.7593	8:	2.0172	1.8961	0.0850
force	8:	11.6473	22:	12.1074	11.8970	0.1937
eamHalo	0:	0.1805	6:	0.6386	0.4179	0.1923
commHalo	22:	0.6640	6:	1.1347	0.9483	0.1255
commReduce	27:	0.0379	24:	0.2679	0.1376	0.1096

Timings for Rank 0				
Timer	# Calls	Avg/Call (s)	Total (s)	% Loop
total	1	14.8650	14.8650	102.22
loop	1	14.5427	14.5427	100.00
timestep	10	1.4530	14.5299	99.91
position	100	0.0004	0.0450	0.31
velocity	200	0.0004	0.0823	0.57
redistribute	101	0.0261	2.6375	18.14
atomHalo	101	0.0194	1.9388	13.48
force	101	0.1158	11.6982	80.44
eamHalo	101	0.0018	0.1818	1.24
commHalo	606	0.0013	0.7607	5.23
commReduce	39	0.0064	0.2511	1.73

Questions & Comments

? / !