Vectorization of Multi-Center, Highly-Parallel

Wolfgang Eckhardt
eckhardw@in.tum.de

Alexander Heinecke
heinecke@in.tum.de

TECHNISCHE UNIVERSITÄT MÜNCHEN

Rigid-body Molecular Dynamics Simulation

Molecular Dynamics Simulation

Our MD simulation implements the following model [1]:

- Treat molecules as spherical rigid bodies composed of multiple interaction sites.
- Determine force between molecules i and j at distance r_{ij} according to the Lennard-Jones-12-6 potential:

$$F(r_{ij}) = 24\epsilon \frac{1}{r_{ij}} \left[\left(\frac{\sigma}{r_{ij}} \right)^6 - 2 \left(\frac{\sigma}{r_{ij}} \right)^{12} \right].$$

- Calculate total force F_i and torque τ_i on molecule i as $F_i = \sum_{j,j \neq i} \sum_{n \in sites_i} \sum_{m \in sites_j} F_{nm}(r_{nm})$, $\tau_i = \sum_{n \in sites_i} d_n \times F_n$.
- Integrate translational and rotational motion with the Rotational Leapfrog Algorithm.

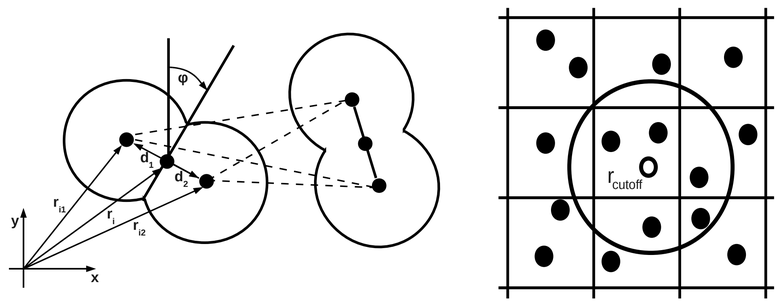

Left: Model of rigid-body interaction.
Right: Schematic of the Linked-Cell structure.

To reduce the algorithmic complexity of the force calculation, the potential can be truncated at a cutoff-distance r_c, reducing the algorithmic complexity from $O(n^2)$ to $O(n)$.

Particles interact if the distance between their centers of mass is less than r_c.

Linked-Cell Algorithms

The search for particles i and j with $r_{ij} < r_c$ can be efficiently implemented with the Linked-Cell algorithm:

- Subdivide domain in cells of edge-length $l \geq r_c$.
- Sort particles into cells.
- In order to find particle pairs, only neighbouring cells have to be searched.

Challenge:

Vectorization cannot be done automatically by the compiler, due to the if-condition. Other approaches rely on gather/scatter hardware [1], or they apply vectorization over the 3d coordinates [2], reducing possible speedups.

Efficient Implementation for AVX and Intel MIC

SSE/AVX: Based on our implementation for atoms [3,4], we derive an implementation for multi-centered molecules.

We store particle data in a SoA-structure per cell and vectorize over the number of molecules. The SoA-structure contains positions and interaction parameters for each interaction site. As the cutoff-condition is evaluated based on the center-of-mass coordinates of the molecule, we redundantly store the molecule coordinates with each center. The center-based distance computation results in quadratic increase in computations:

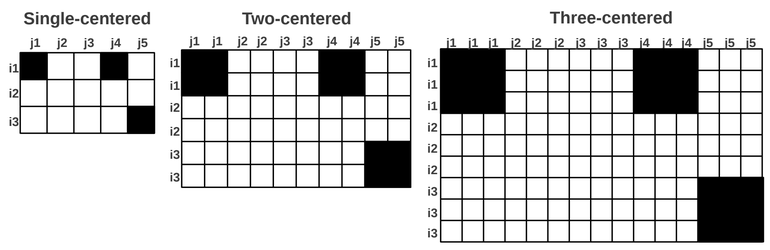

Intel MIC: For larger vectors, our AVX implementation suffers from low vector load because of the cutoff condition:

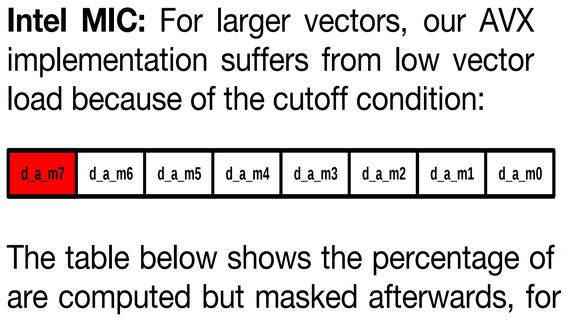

The table below shows the percentage of interactions, that are computed but masked afterwards, for SSE and AVX in double precision for liquid at cut-off radius $r_c = 5\sigma$.

	1 center	2 centers	3 centers	4 centers
SSE	19.1 %	0 %	9.8 %	0 %
AVX	38.3 %	24.6 %	23.5 %	0 %

Gather and *scatter* allow for full vector exploitation, so we store element indices instead of masks during the computation of distances.

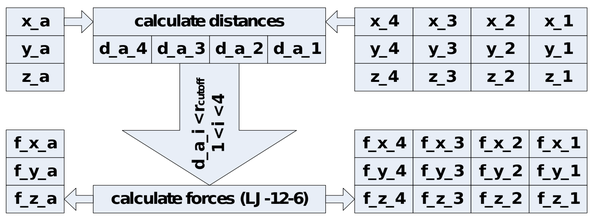

Vectorization of the Lennard-Jones-12-6 force calculation kernel.

Avoid the quadratic increase by lookup-table for distances:

- For a molecule, compute distances to all other centers and store interaction mask.
- If the current molecule interacts, iterate over its centers and compute force with all interacting other centers.

$\Rightarrow$ linear increase of computations

If the molecule does not interact with any other molecule $\Rightarrow$ skip force computation completely. Else, compute forces for whole vector and mask, where $r_{ij} > r_c$.

Pseudo-code of the kernel:

```
for ( i = 0; i < cell1._num_mol; ++i ) {
    compute_molecule = 0;
    for ( j = 0; j < cell2._num_ljcenters; ++j ) {
        dx =cell1._mol_x[i] − cell2._mol_center_x[j];
        dy = cell1._mol_y[i] − cell2._mol_center_y[j];
        dz = cell1._mol_z[i] − cell2._mol_center_z[j];
        r2 = dx ∗ dx + dy ∗ dy + dz ∗ dz;

        forceMask =m_r2 < _rc2 ? 0xFFFFFFFF : 0;
        compute_molecule |= forceMask;
        _dist_lookup[j] = forceMask;
    }

    if ( !compute_molecule ) continue;

    for ( local_i = 0; local_i < cell1._num_ljc[i]; local_i++ ) {
        for ( j = 0; j < cell2._num_ljcenters; ++j ) {
            compute_potential(cell1, i_center_idx, cell2, j, _dist_lookup);
        }
        i_center_idx++;
    }
}
```

Computing distances between a molecule and all centers and evalutation of the cutoff condition delivers the bit mask we use to compress center indices.
The cost of VCOMPRESS is measurable.

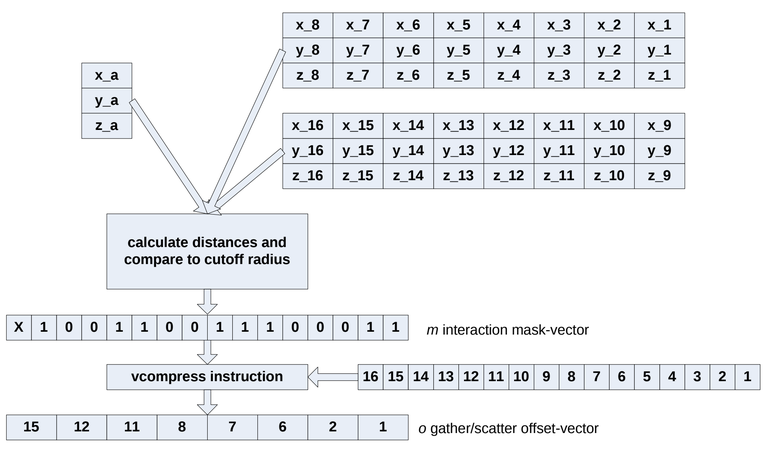

The index vector is stored, instead of the bit-mask as in the case of AVX. Only if a molecule interacts with any center, the force computation is executed. The interaction partners' data is gathered, and broadcasted after the computation.

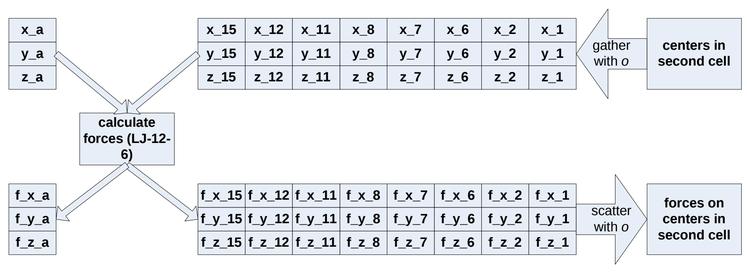

Ongoing work:

- Improved shared-memory parallelization
- Load-balancing for heterogeneous systems to exploit both MIC and host CPU

References

[1] D. C. Rapaport: The Art of Molecular Dynamics Simulation. Cambridge University Press, 2004.
[2] Peng et. al.: Exploiting hierarchical parallelisms for molecular dynamics simulation on multicore clusters. In *The Journal of Supercomputing*, 57, p. 20 - 33, 2011.
[3] W. Eckhardt, A. Heinecke: An efficient Vectorization of Linked-Cell Particle Simulations. In *ACM International Conference on Computing Frontiers*, p. 241 - 243, May 2012.
[4] W. Eckhardt, A. Heinecke et. al.: 591 TFLOPS Multi-Trillion Particles Simulation on SuperMUC. In *International Supercomputing Conference (ISC) Proceedings 2013*, p. 1 - 12, June 2013.

Rigid-Body Molecular Dynamics Simulations

Wolfgang Hölzl
hoelzlw@in.tum.de

Hans-Joachim Bungartz
bungartz@in.tum.de

SPONSORED BY THE
Federal Ministry of Education and Research

Results for AVX

Performance on SuperMUC@LRZ, Munich (2 x 8-core Intel SNB@2.7 GHz per node):
Runtime for $10.7 \cdot 10^6$ molecules initially arranged on a cubic lattice at liquid density.

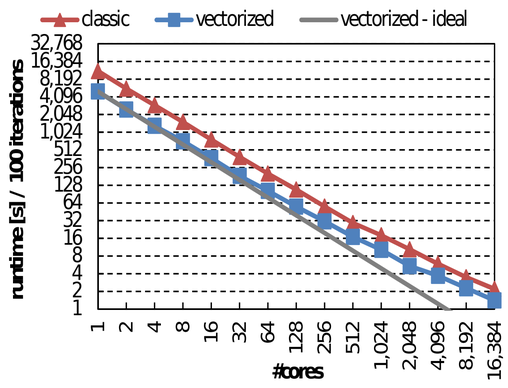

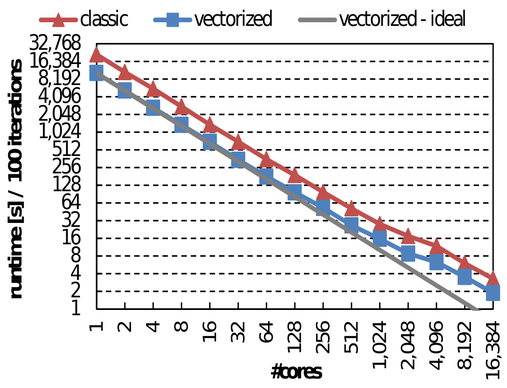

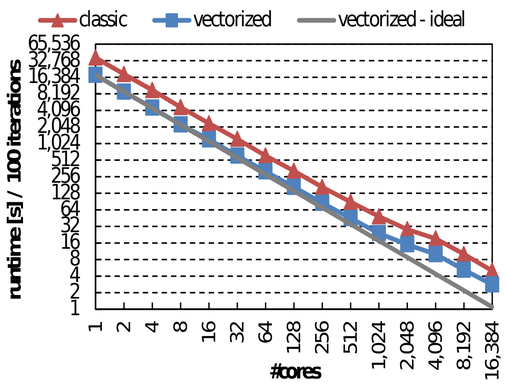

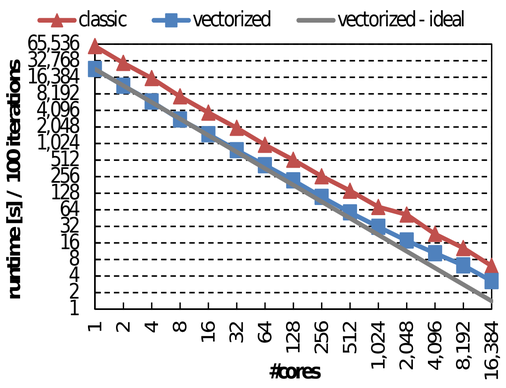

- Performance gain through AVX $\approx$ 3x
- Scalability the better, the more complex the molecule:
  - Argon: 1 center
  - Ethan: 2 center
  - CO2: 3 center
  - Aceton: 4 center

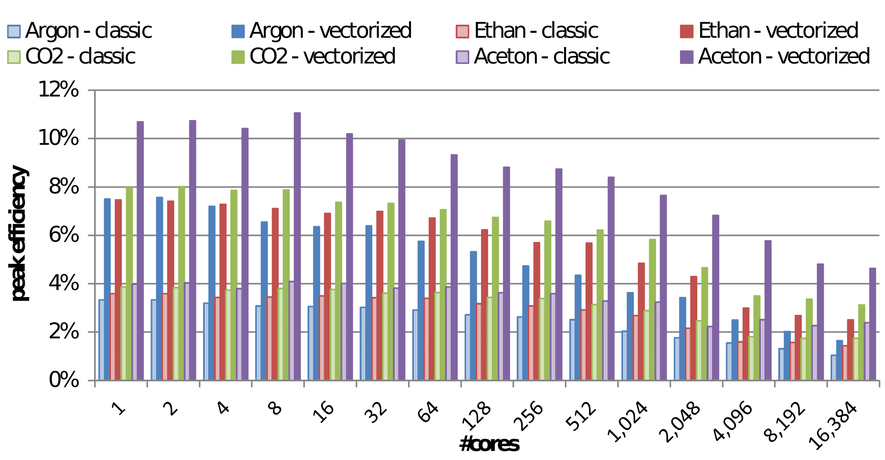

Comparison of Performance on Intel Ivy Bridge vs. Haswell:

- i7-3770K: Ivy Bridge Quad-Core 3.5 GHz
- i7-4770K: Haswell Quad-Core 3.5 GHz

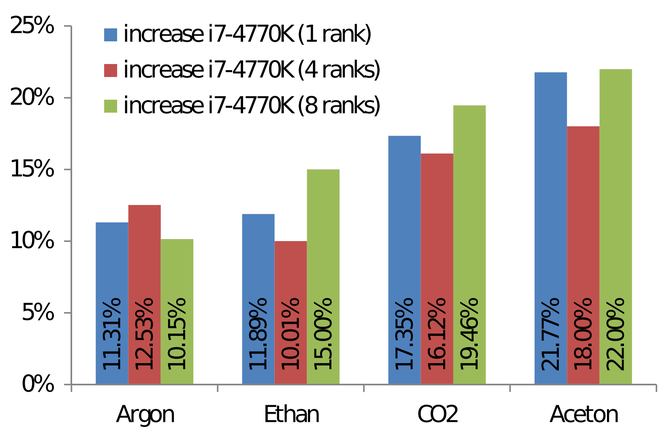

Performance of AVX-128 version for mon-atomic fluids (32 Bytes per atom only):
Performance on SuperMUC@LRZ, Munich (2 x 8-core Intel SNB@2.7 GHz per node) and on Hermit@HLRZ, Stuttgart (2 x 16-core AMD Interlagos@2,3 GHz)

- Strong scaling: $9.5 \cdot 10^8$ particles.
- Weak scaling: $1.6 \cdot 10^7$ particles per node.

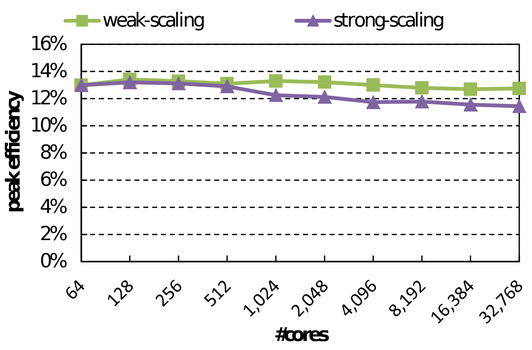

Efficiency Hermit

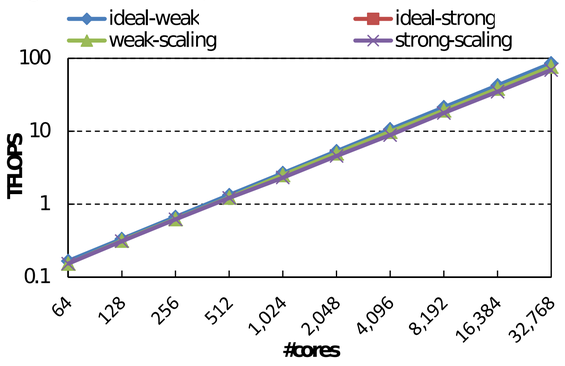

Scaling Hermit

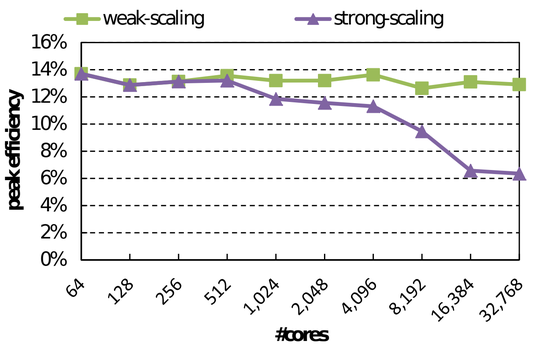

Efficiency SuperMUC

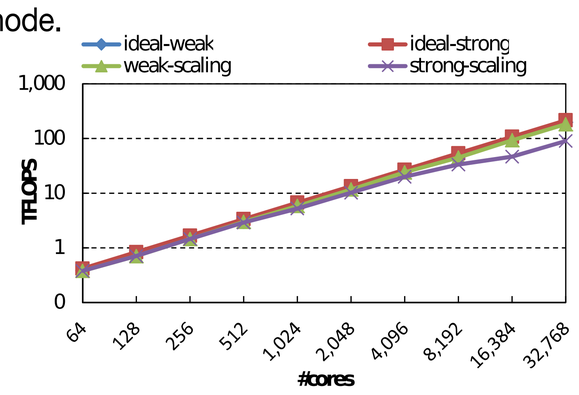

Scaling SuperMUC

Results for Intel MIC

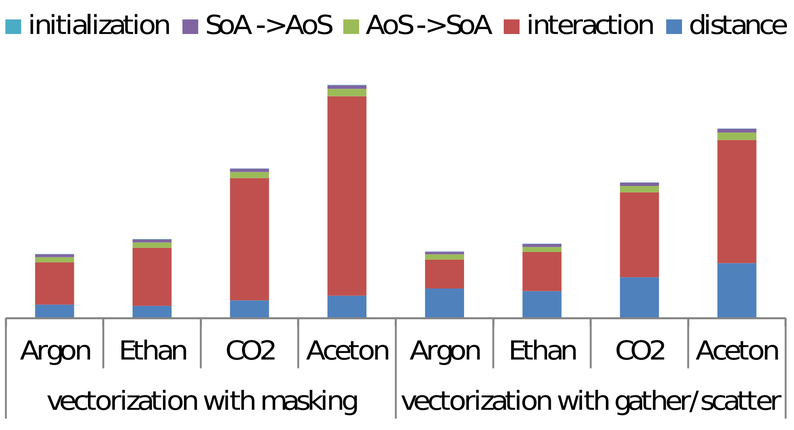

Performance evaluation on Intel MIC:

- Transferring our vectorization approach directly to Xeon Phi (KNC@1.056 GHz) results in roughly half SuperMUC node speed.
- Using gather/scatter boosted to SuperMUC node speed, but still pure MPI with 120 ranks.
- Using gather/scatter, overhead is shifted to the distance computation (creation of index vector), but the force computation is performed more efficiently (less elements masked).

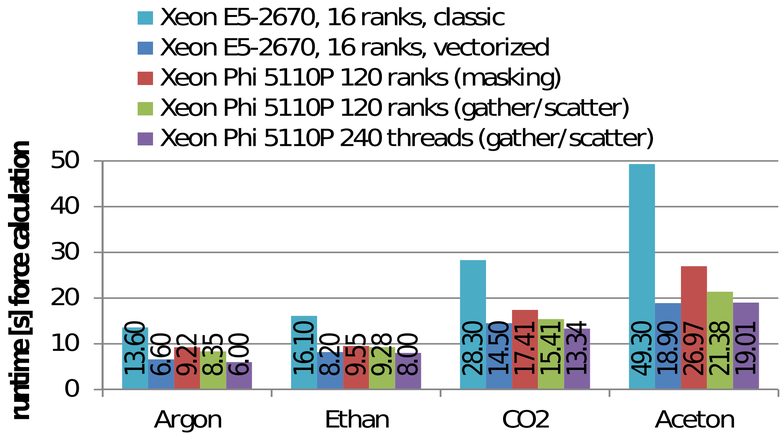

Acknowledgements

The authors gratefully acknowledge the Gauss Centre for Supercomputing e.V. (www.gauss-centre.eu) for funding this project by providing computing time on the GCS Supercomputer SuperMUC at Leibniz Supercomputing Centre (LRZ, www.lrz.de). This work has been partly funded by the German federal ministry of education and research (BMBF) in the project SkaSim within the research programme IKT 2020 under Grant Number 01IH13005F. Furthermore, this material is based upon work supported by the National Science Foundation under Grant Number 1137097 and by the University of Tennessee through the Beacon Project. We especially thank Vince Betro and Glenn Brook for their support when using the Beacon cluster. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation or the University of Tennessee.

Technische Universität München
Chair of Scientific Computing