A Scalable, Linear-Time Dynamic Cutoff Algorithm for Molecular Dynamics

Paul Springer, Ahmed E. Ismail, **Paolo Bientinesi**

Aachen Institute for Advanced Study in Computational Engineering Science

ISC 2015, Frankfurt, 14.07.15

Deutsche Forschungsgemeinschaft

DFG

Molecular Dynamics

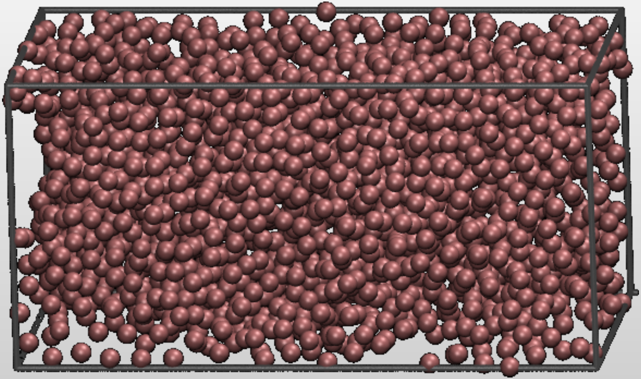

$$V_C(r) = \frac{1}{4\pi\epsilon}\frac{Q}{r}$$

$$V_{LJ}(r) = 4\epsilon\left(\left(\frac{\sigma}{r}\right)^{12} - \left(\frac{\sigma}{r}\right)^{6}\right)$$

Interfacial systems

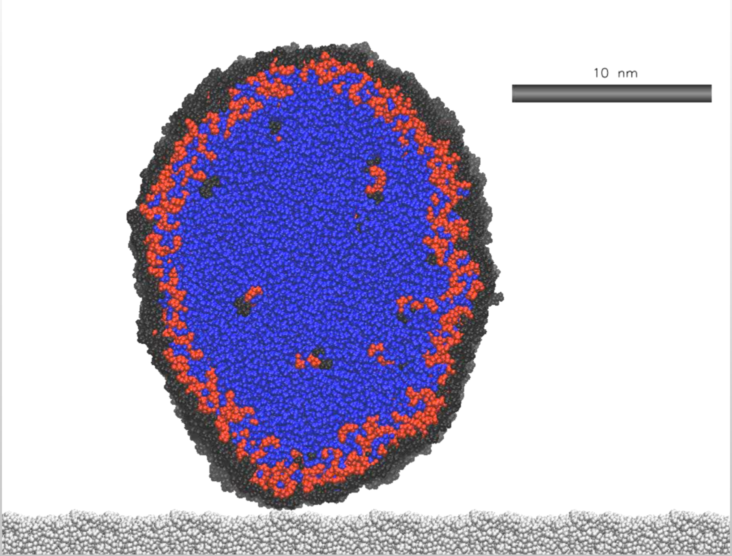

Short-range method: "cutoff"

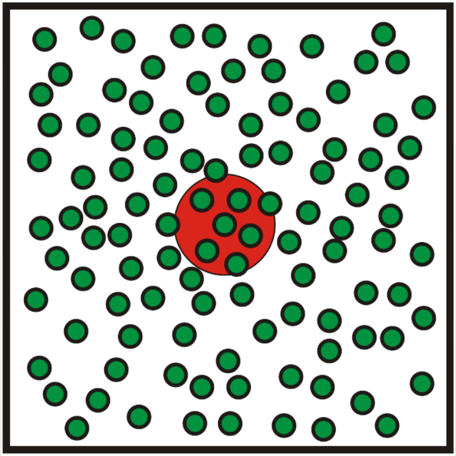

Interface

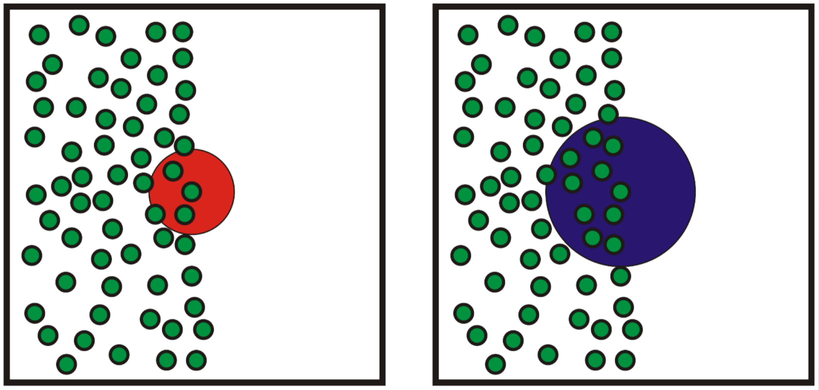

Interface

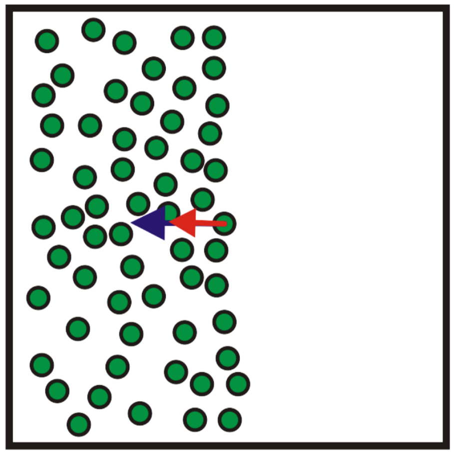

Interface

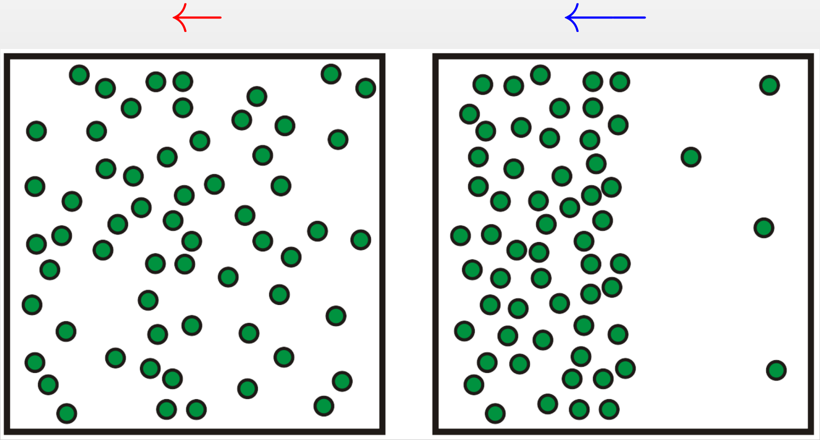

Long-range solvers

Ewald methods, based on FFT. $O(n \log(n))$

Long-range solvers

Ewald methods, based on FFT. $O(n \log(n))$

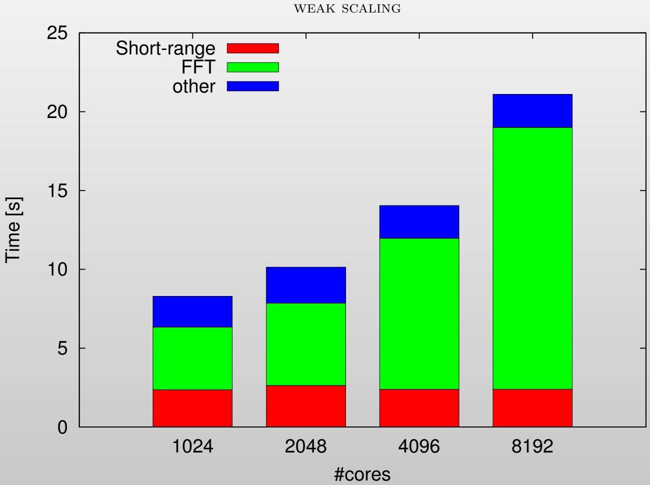

PPPM long-range solver, 1200 particles/core, IBM BG/Q.

Recap

- Small cutoff $\rightarrow$ accuracy loss at the interface
- Large cutoff $\rightarrow$ flops wasted in the bulk phase

Recap

- Small cutoff $\rightarrow$ accuracy loss at the interface
- Large cutoff $\rightarrow$ flops wasted in the bulk phase
- **Idea**
  Cutoff chosen dynamically, according to the distance from the interface

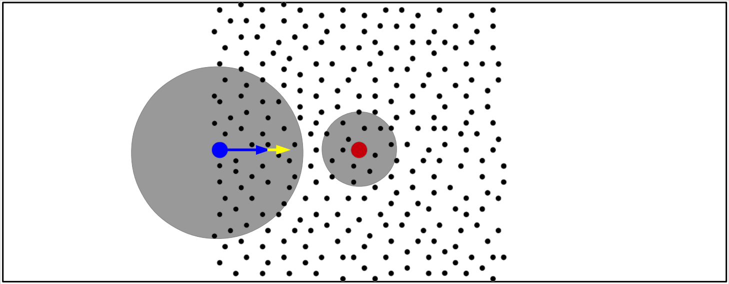

Challenges & Objectives

Cutoff chosen particle by particle

- How to detect the interface?
- How to compute the distance particle-interface?

- Complexity $O(n)$
- Only local communication

Dynamic Cutoff Method

1. Interface detection*
2. Distance calculation + cutoff assignemnt*

. (Almost) traditional short-range method

   3. Neighbor-list*
   4. Forces
   5. Positions

*: not in every iteration

Dynamic Cutoff Method

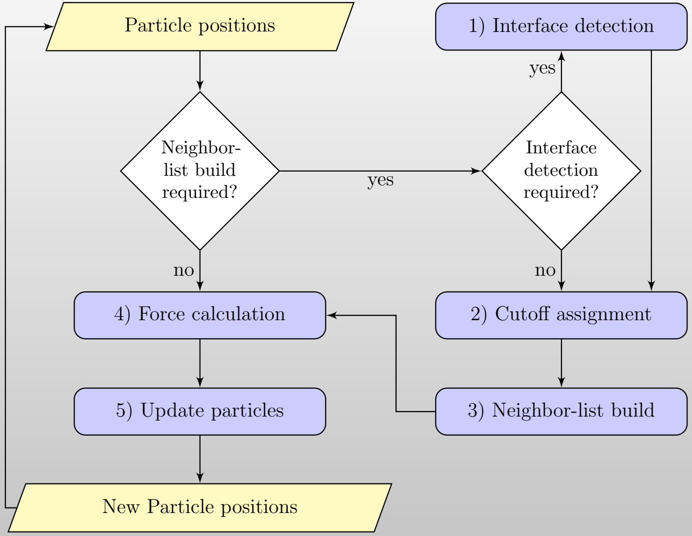

1) Interface detection

- Bin particles to 3D bins

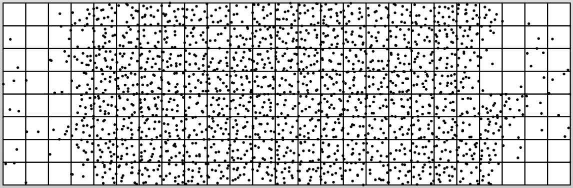

1) Interface detection

- Bin particles to 3D bins
- Treat particle densities as gray values

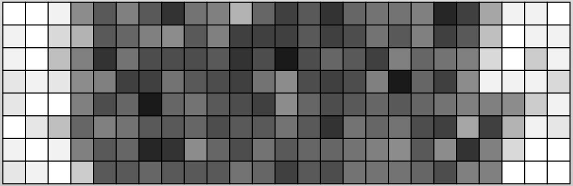

1) Interface detection

- Bin particles to 3D bins
- Treat particle densities as gray values
- Apply image segmentation

  Minimization of Mumford-Shah functional; periodic boundary conditions; finite differences + filtering, local communication only

2) Distance from interface

- Distance is only required close to the interface

- Interface-box distance (Fast Sweeping Method) $+$ box-particle distance (trilinear interpolation)

  local communication; one scalar reduction

2) Distance from interface

- Distance is only required close to the interface
- Interface-box distance (Fast Sweeping Method) + box-particle distance (trilinear interpolation)

  local communication; one scalar reduction

2) Distance from interface

- Distance is only required close to the interface
- Interface-box distance (Fast Sweeping Method) + box-particle distance (trilinear interpolation)

  local communication; one scalar reduction

2) Distance from interface

- Distance is only required close to the interface
- Interface-box distance (Fast Sweeping Method) +
  box-particle distance (trilinear interpolation)

  local communication; one scalar reduction

2) Distance from interface

- Distance is only required close to the interface
- Interface-box distance (Fast Sweeping Method) + box-particle distance (trilinear interpolation)

  local communication; one scalar reduction

2) Distance from interface

- Distance is only required close to the interface
- Interface-box distance (Fast Sweeping Method) +
  box-particle distance (trilinear interpolation)

  local communication; one scalar reduction

Cutoff-based Fast Sweeping Method

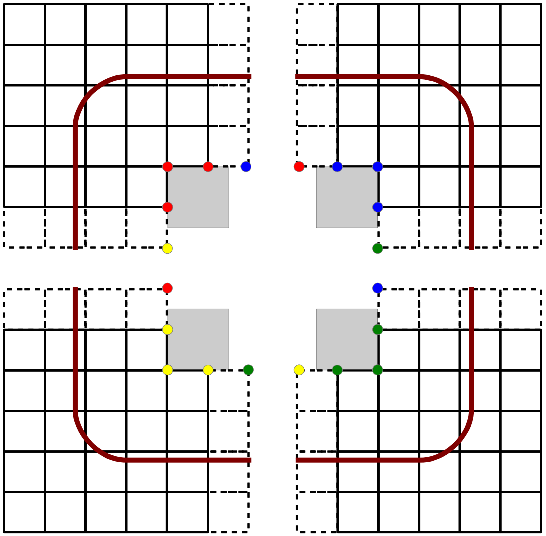

Initial state

Cutoff-based Fast Sweeping Method

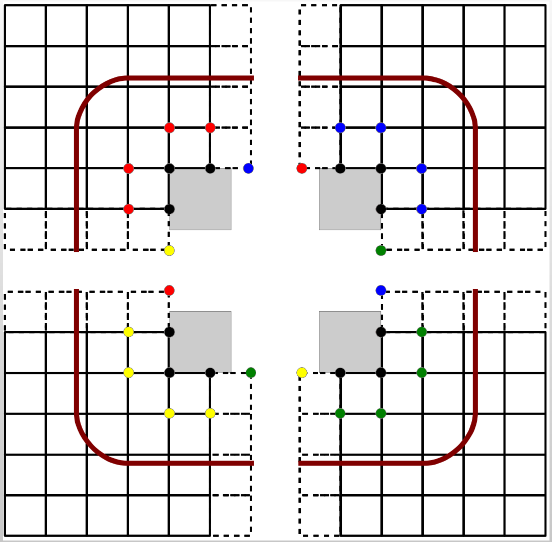

Local CFSM step

Cutoff-based Fast Sweeping Method

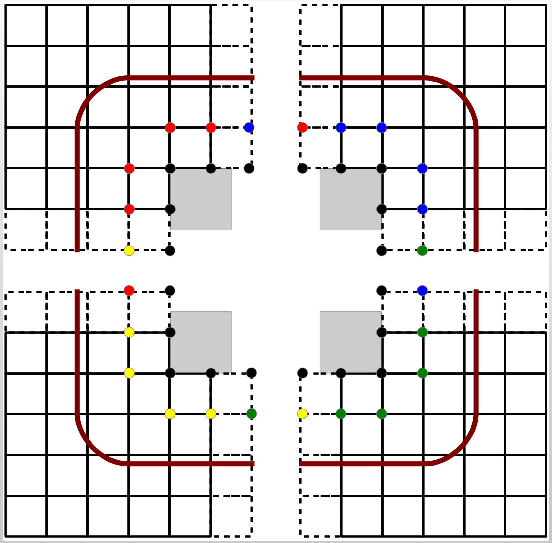

Ghost Exchange

Cutoff-based Fast Sweeping Method

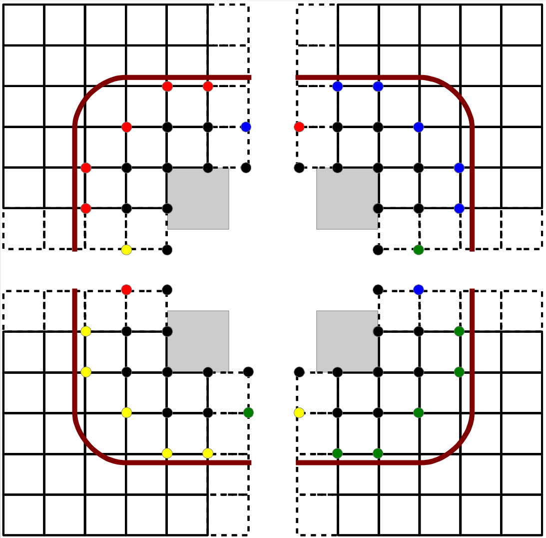

Local CFSM step

Cutoff-based Fast Sweeping Method

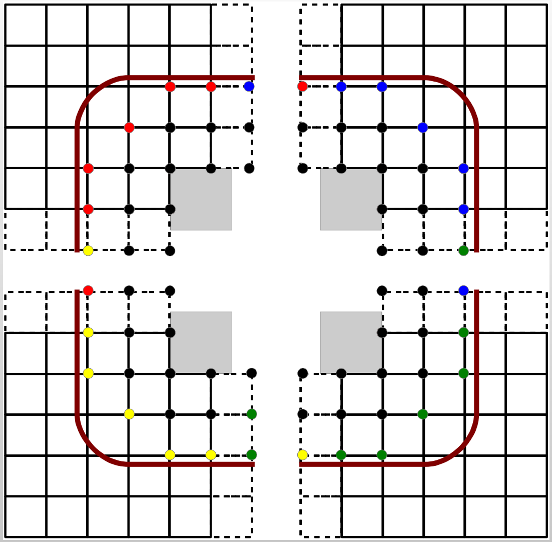

Ghost Exchange

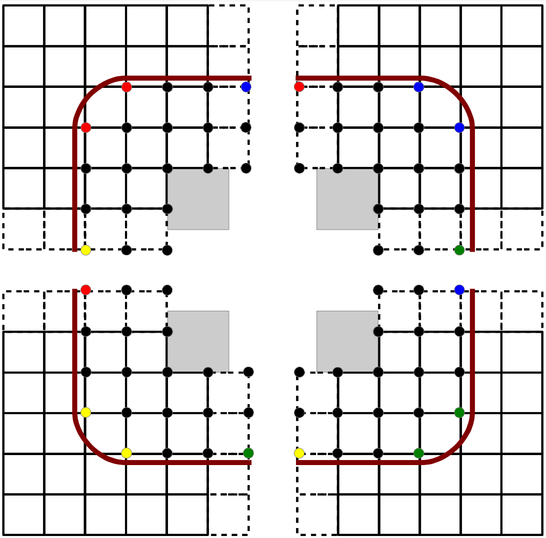

Local CFSM step

Dynamic Cutoff Method

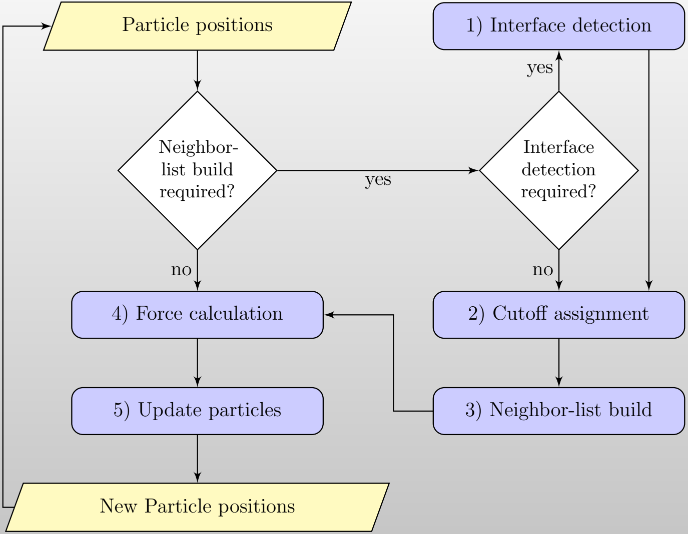

3) Neighbor-list

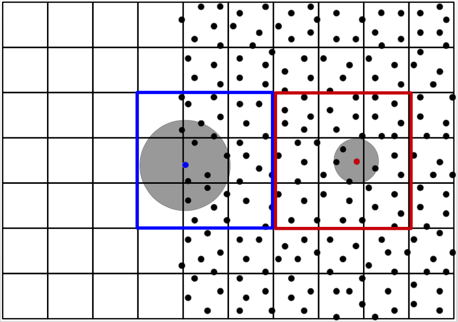

- Spatial binning. Search limited to neighboring boxes $\Rightarrow$ $\mathcal{O}(N)$
- DCM $\Rightarrow$ Newton's 3rd law not applicable anymore

3) Neighbor-list

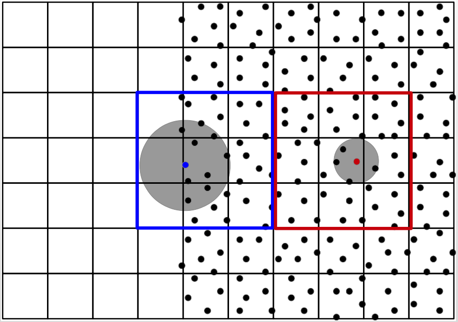

- Spatial binning. Search limited to neighboring boxes $\Rightarrow \mathcal{O}(N)$
- DCM $\Rightarrow$ Newton's 3rd law not applicable anymore
- Bins of size $r_c^{max} \times r_c^{max} \times r_c^{max}$ $\Rightarrow$ performance loss in bulk phase

3) Neighbor-list

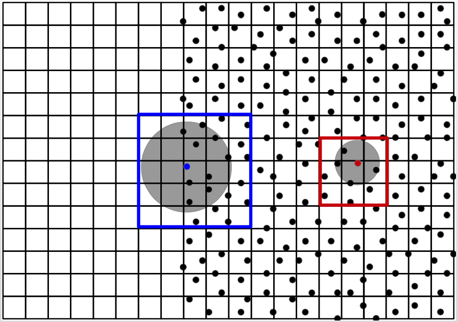

- Spatial binning. Search limited to neighboring boxes $\Rightarrow \mathcal{O}(N)$
- DCM $\Rightarrow$ Newton's 3rd law not applicable anymore
- Bins of size $r_c^{max} \times r_c^{max} \times r_c^{max} \Rightarrow$ performance loss in bulk phase
- Bins of size $r_c^{min} \times r_c^{min} \times r_c^{min} \Rightarrow$ speedups of $\mathbf{4} - \mathbf{6}\times$

Experiments

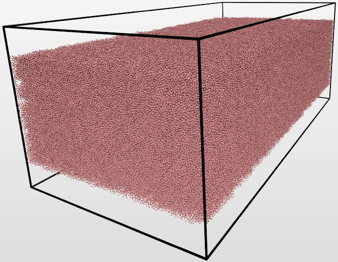

a) planar interface

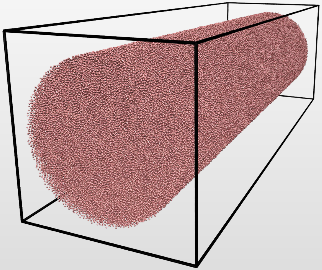

b) non-planar interface

System (a) $\rightarrow$ JUQUEEN, IBM, BG/Q

System (b) $\rightarrow$ SuperMUC, Intel, Sandy Bridge

LAMMPS Molecular Dynamics Simulator `lammps.sandia.gov`

Strong Scaling

Size: $1200 \times 32768 \approx 4 \cdot 10^7$ particles

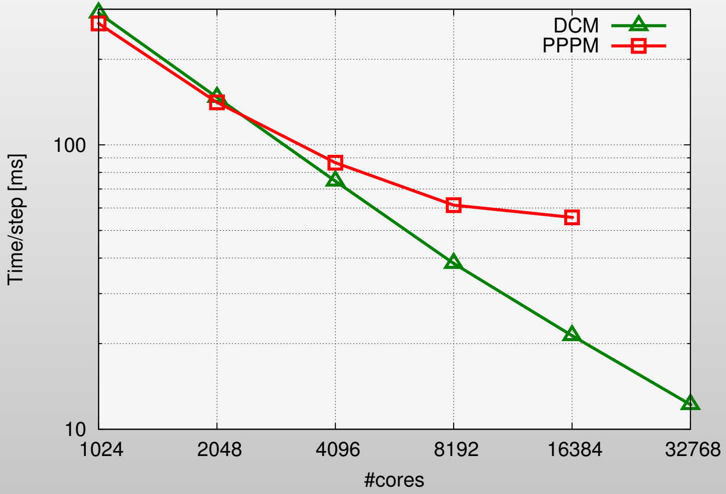

SuperMUC

Strong Scaling

Size: $1200 \times 32768 \approx 4 \cdot 10^7$ particles

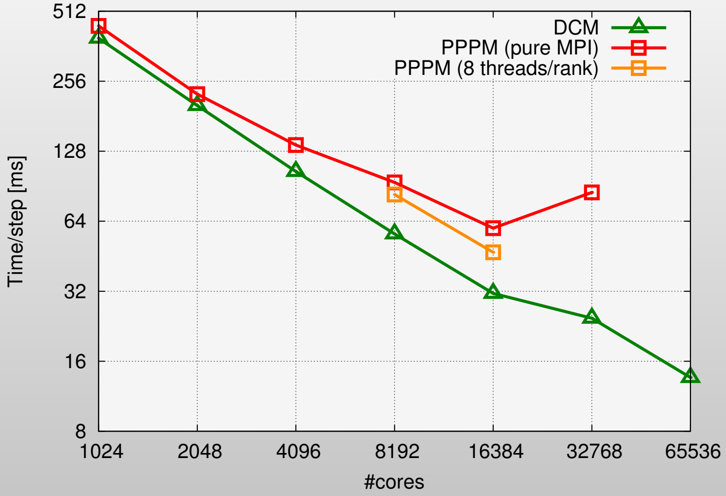

JUQUEEN

Weak Scaling

1200 particles per core

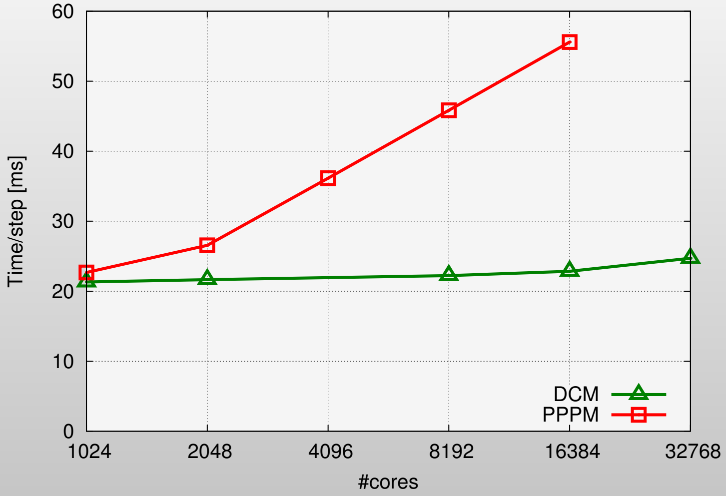

SuperMUC

Weak Scaling

1200 particles per core

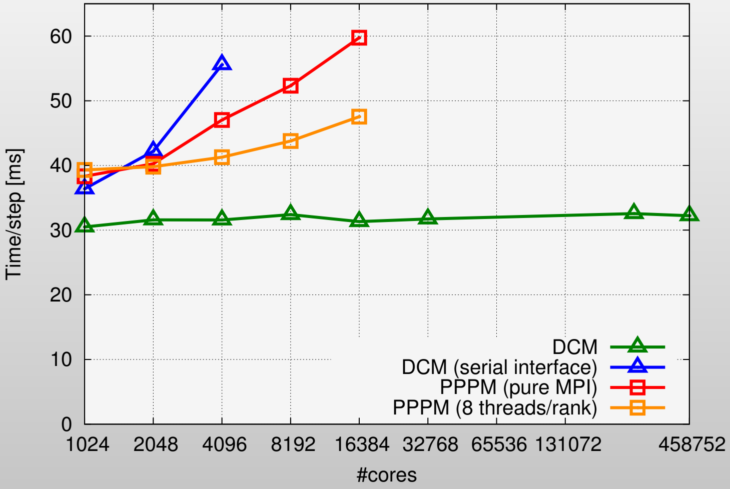

JUQUEEN

Weak Scaling

1200 particles per core

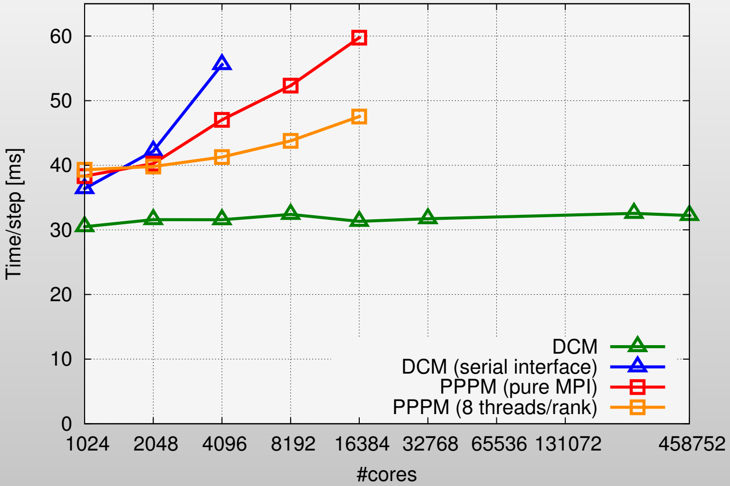

DCM	3.0/**7.0**	3.0/**8.0**	3.5/**8.0**	3.0/**9.0**	3.5/**9.0**
Speedup	2.31	2.66	2.38	2.42	2.23

Accuracy

Absolute error with respect to long-range Ewald summation

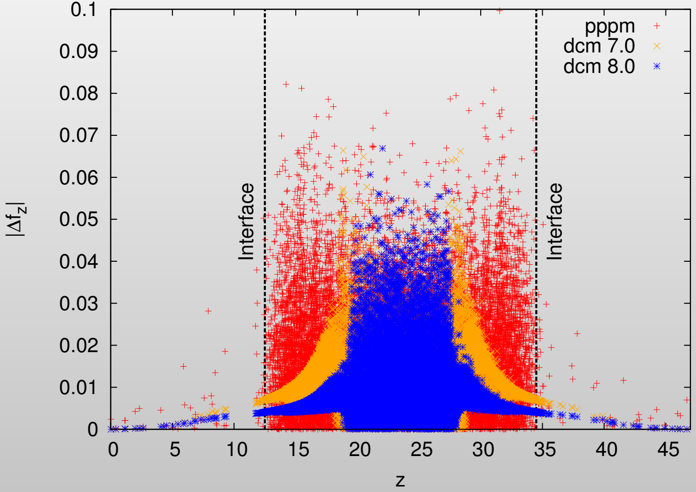

Accuracy vs. performance

How to select min/max cutoffs

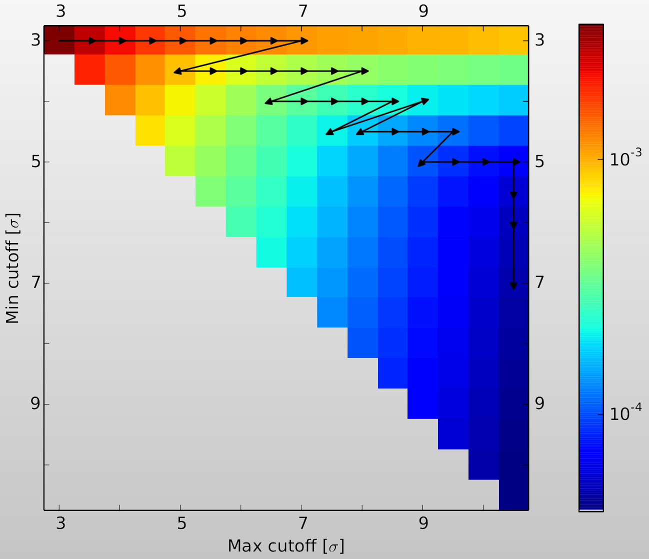

Root-mean-square error of the force

Conclusions

- Idea

  Cutoff chosen particle by particle

  $\Rightarrow$ interface detection, distance calculation

- Results
  - Nearly perfect strong and weak scalability
  - 5.5×10^8 particles on $458,752$ cores
  - As accurate as PPPM, much more scalable than PPPM

- Future work
  - Accelerators?
  - Use DCM as short-range solver within Ewald-based solvers

Conclusions

- Idea

  Cutoff chosen particle by particle

  $\Rightarrow$ interface detection, distance calculation

- Results
  - Nearly perfect strong and weak scalability
  - 5.5×10^8 particles on $458,752$ cores
  - As accurate as PPPM, much more scalable than PPPM

- Future work
  - Accelerators?
  - Use DCM as short-range solver within Ewald-based solvers

Thank you for your attention