

Tiny Tasks, Huge Impact

March 10, 2023 | I. Kabadshow, M. Zych, A. Beckmann, L. Morgenstern | Aristote Workshop, Paris

Member of the Helmholtz Association

My Background at JSC

Fast Multipole Library

- Linear scaling Coulomb solver for MD
- Dynamic protonation and long-range electrostatics for GROMACS

Tasking Library

- Fine-grained tasking
- Joint lab (Aidas) CEA/FZJ: incorporate tasking into Smilei PIC code





















Coarse-Grained P Phase-based-Parallelism: Exam	arallelization				
Input P2N	M2M M2L	L2L L2P	P2P	Output	
 Different amount of available loop-level parallelism within each phase Some phases contain sub-dependencies Synchronizations might be problematic 					
Member of the Helmholtz Association	March 10, 2023	Slide 10		JÜLICH Forschungszentrum	



FMM Full Dependency Graph

How to deal with it

- We do not want to set it up explicitly upfront
- Going through the graph implicitely
- Worksharing should honor locality
- Workstealing should honor locality
- No artifical synchronization (just synching fine-grained tasks)
- Only ready-to-execute tasks should be stored
- Utilize critical path
- Utilize backfill tasks

Member of the Helmholtz Association	March 10, 2023	Slide 12	JÜLICH Forschungszentrum

Tasking Framework Overview Mostly independent of the algorithm Static Data-Flow Dispatcher? Load-Balancer MultiQueue Scheduler Executor **Threading Wrapper** Thread Core ICH. March 10, 2023 Slide 13 Member of the Helmholtz Association

Tasking With Type-Driven Priority Scheduling

March 10, 2023 | |

Member of the Helmholtz Association	March 10, 2023	Slide 14	JÜLICH Forschungszentrum



FMSolvr specific queue for multiple producer and multiple consumer



MultiQueue

Dependents on the algorithm

A How the user sees it

```
using multi_queue = MultiQueue<M2LTask;
using multi_queue = MultiQueue<P2MTask, M2MTask, M2LTask;
using multi_queue = MultiQueue<P2MTask, M2MTask, M2LTask, L2LTask, L2PTask, P2PTask;
```







Scaling

March 10, 2023 | |

Member of the Helmholtz Association March 10, 2023 Slide 20	JÜLICH Forschungszentrum
---	------------------------------------

Strong-Scaling

103680 Particles with p=10, d=4, on 2×26-core Intel Xeon Platinum 8170 (FMM parameter unfavorable for tasking)



Strong-Scaling

103680 Particles with p=10, d=4, on 2×26-core Intel Xeon Platinum 8170 (FMM parameter unfavorable for tasking)





Strong-Scaling

1000 Particles with p=5, d=3, on 2×26-core Intel Xeon Platinum 8170 (FMM parameter unfavorable for tasking)



Next Steps

- Fully decouple from FMSolvr
- Simplify template meta-ptogramming with latest C++20 standard
- Extend to more parts of Smilei (Aidas project)
- Modularize the tasking itself (queues, locks, work stealing)



March 10, 2023

Slide 25



Thank you for your attention.

Questions?

Contact: lvo Kabadshow i.kabadshow@fz-juelich.de



Tiny Tasks, Huge Impact

March 10, 2023 | I. Kabadshow, M. Zych, A. Beckmann, L. Morgenstern | Aristote Workshop, Paris

