Patrick Höhn, Zoya Masih

# IO characterization for Screw Conveyor Simulations Using LIGGGHTS

Georg-August University - Goettingen
Gesellschaft für wissenschaftliche Datenverarbeitung mbH Göttingen(GWDG)

# Table of Contents

# Application motivation

- Many energy resources located far below the surface
- Costs for geothermal projects largely governed by wellbore construction
- Large cost driver non-productive time and unexpected problems
- Many challenges caused by lateral drill string vibrations and non-optimal cuttings transport
- Due to challenging in-situ measurements, simulation based approach

# Experimental study

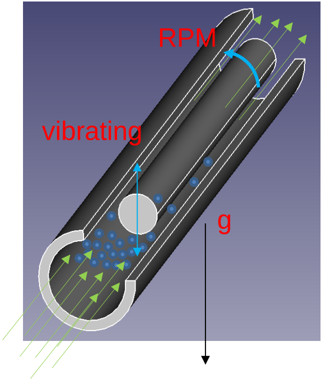


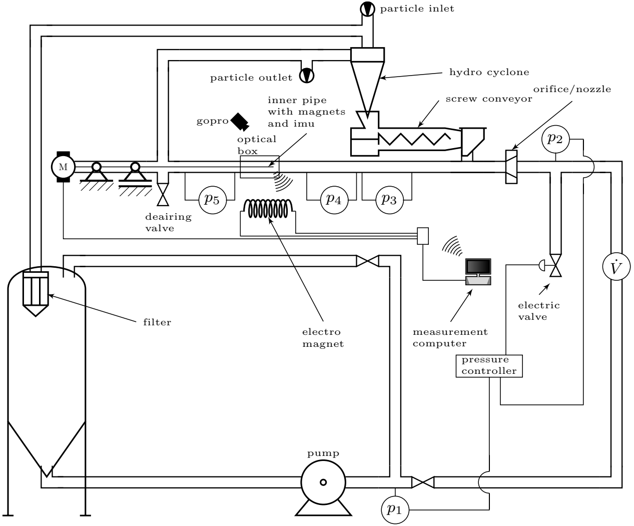

# Dimensioning of components

- Hydrocyclone
- Screw conveyor
- Pipes
- Pump
- Particle injector
- Valves
- Sensors
- Electromagnet

# Background of Screw Conveyor Dimensioning

- Sizing often based on simplified technical standards (e.g. DIN 15262)
- Simplified setup requires defined lower particle loads
- Difficult to parametrize the simulation regarding, e.g. adhesion
- More accurate estimations based on Discrete Element Method (DEM)

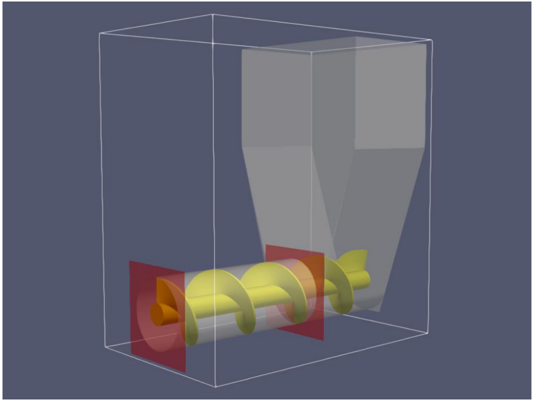

# Particle Simulation Software LIGGGHTS®

LAMMPS Improved for General Granular and Granular Heat Transfer Simulations

- Open Source Simulation Software
- Implementation of Discrete Element Method
- Based on the molecular dynamics code LAMMPS
- LAMMPS is a classical molecular dynamics simulator
- PUBLIC version developed by DCS Computing GmbH and JKU Linz since 2009
- Commercial successor of PUBLIC version Aspherix®since 2019
- Actively developed PFM-fork at JKU Linz since 2015
- Can be coupled to OpenFOAM®using CFDEM®coupling

# Motivation for IO-Characterization

- Large computational cost
- Performance engineering of the utilized software
- Understanding the IO-patterns
- Vision for porting to GPUs

# Table of Contents

# Evaluated Tools

- Evaluated for profiling and tracing: Score-P, Extrae, and Darshan
- MPI functions only traced with Score-P and Extrae
- Consistent results of two tools

# Execution on 3 processes

# Overhead in 3 Processes

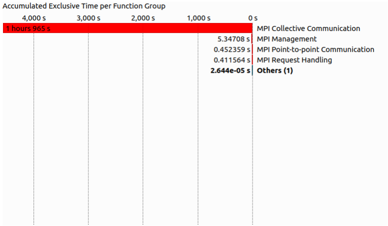


Figure: 2 Processes

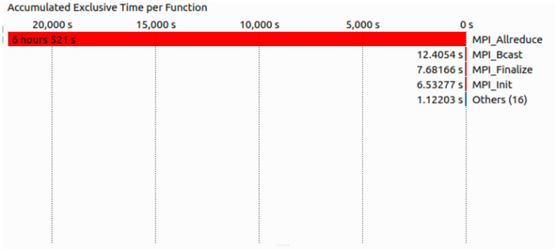


Figure: 3 Processes

# Communication on 2 and 3 Processes

- Significant communication load on Rank 0, while Rank 1 is mainly computational.

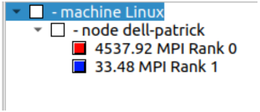


- Communication load heavily shifted to Rank 1. Ranks 0 and 2 are still involved in communication, but Rank 1 shows a dramatic increase compared to the 2-process setup.

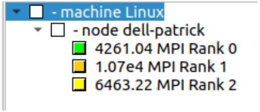

# Number of Invocations

| | |
|---|---|
| Sum | 8,719,014 |
| Function Group | MPI Collective Communication |
| Value | 8,718,862 |
| Function Group | MPI Management |
| Value | 114 |
| Function Group | MPI Point-to-point Communication |
| Value | 32 |
| Function Group | MPI Request Handling |
| Value | 6 |

Figure: Number of invocations - 2 Processes

| | |
|---|---|
| Sum | 13,078,499 |
| Function Group | MPI Collective Communication |
| Value | 13,078,271 |
| Function Group | MPI Management |
| Value | 174 |
| Function Group | MPI Point-to-point Communication |
| Value | 46 |
| Function Group | MPI Request Handling |
| Value | 8 |

Figure: Number of invocations - 3 Processes

# Domain Decomposition: Processors 1 1 2

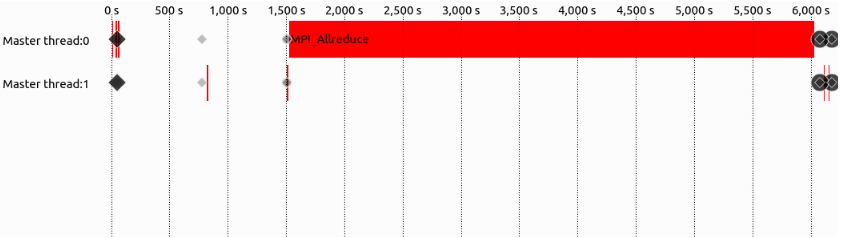


**Key Observations:**

- Thread 0: Dominated by MPI operations, indicating a heavy communication load.
- Thread 1: it's more focused on computation.

# Computation on 2 Processes

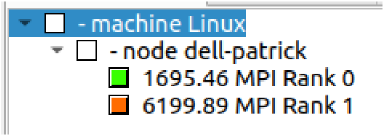


**Key Observations:**

- Rank 0: Spent 1695 seconds on MPI communication, which is a small fraction compared to its total computation time.
- Rank 1 is primarily focused on computation, with less emphasis on MPI communication than Rank 0.

# Communication Pattern

High Volume of 28kiB Messages suggests a specific communication pattern in the application

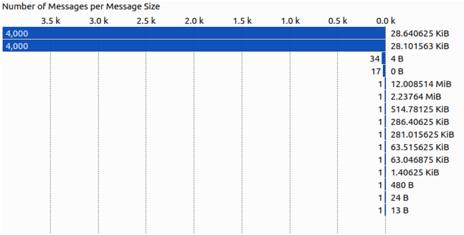


Figure: Number of messages per message size

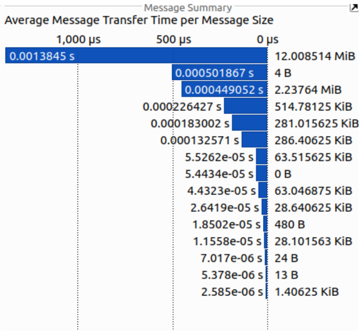


Figure: Avg message transfer time

# Workload Distribution: Visits (occ)

- The nature of the tasks differs across ranks
- with some subdomains likely being more computationally intense
- This suggests a load imbalance in computational cost, despite the even division of work by the number of tasks

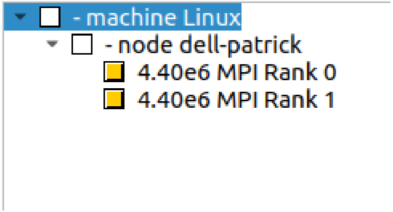

# Summary and Future Work

- These are the first steps and initial observations
- Capturing more I/O data
- Experiment with other Profiling/Tracing tools
- Scale up the simulation’s size to consider the impacts of parallelization

# References


- Philip Carns. Robert Latham. Robert Ross. Kamil Iskra. Samuel Lang and Katherine Riley. “24/7 Characterization of Petascale I/O Workloads”. In: 2009 IEEE International Conference on Cluster Computing and Workshops 13.2 (2009), pp. 1–10.
- Goniva, C., Kloss, C., Deen, N.G., Kuipers, J.A.M. and Pirker, S. (2012): “Influence of Rolling Friction Modelling on Single Spout Fluidized Bed Simulations”, Particuology, DOI 10.1016/j.partic.2012.05.002
- "Models, algorithms and validation for opensource DEM and CFD-DEM", Christoph Kloss, Christoph Goniva, Alice Hager, Stefan Amberger, Stefan Pirker - Progress in Computational Fluid Dynamics, An Int. J. 2012 - Vol. 12, No.2/3 pp. 140 - 152

# References

- Höhn, P., 2024. Investigation of the influence of cuttings transport on drill string dynamics. Clausthal-Zellerfeld.
- Barcelona Supercomputing Center – Performance Tools, Extrae, 2023, [Online]. Available: http://tools.bsc.es/extrae.
- Feld, C., et al. (2025). Scalable performance measurement infrastructure for parallel codes (Score-P) (9.0). Zenodo. https://doi.org/10.5281/zenodo.15115212