

# Large-scale Parallel Discrete Element Simulations of Granular Flow

Jens H. Walther<sup>\*†</sup> and Ivo F. Sbalzarini<sup>†</sup>

<sup>\*</sup>Dept. of Mechanical Engineering, Technical University of Denmark, DK-2800 Lyngby, Denmark

<sup>†</sup>Institute of Computational Science, ETH Zurich, CH-8092 Zurich, Switzerland

E-mail: <sup>\*†</sup>jhw@mek.dtu.dk, <sup>†</sup>ivos@ethz.ch

**Abstract.** We present large-scale parallel direct numerical simulations of granular flow using a novel, portable software program for Discrete Element Method (DEM) simulations. Since particle methods provide a unifying framework for both discrete and continuous systems, we base the program on the PPM library [1], which has already been demonstrated to provide transparent parallelization and state-of-the-art parallel efficiency using particle methods for continuous systems. Adapting PPM to discrete systems, we report results from three-dimensional simulations of a sand avalanche down an inclined plane [2, 3]. All simulations use the particle interaction model by Silbert et al. [4], including Hertzian contact forces and elastic and viscoelastic deformations. We demonstrate the parallel performance and scalability of the new simulation program using up to 122 million particles on 192 processors, employing adaptive domain decomposition and load balancing techniques.

## 1 Introduction

Discrete Element Method (DEM) simulations constitute an important tool for the study of granular matter flows. Their fully resolved “atomistic” character allows to determine material properties [5] and effective macroscopic dynamics under minimal assumptions. This renders DEM simulations invaluable in the search for continuum theoretic descriptions of granular matter [6, 7, 3, 8]. DEM simulations have, for example, successfully been used to study the packing in a cylindrical container using up to 144 000 particles [9]. While such simulations are particularly well suited for the multiscale considerations [10] of linking the single-particle description to the continuum description, the size of the computationally tractable systems is frequently limited to a few hundreds of thousands of grains [11, 12]. This limitation currently hinders the use of fully resolved DEM simulations in studying the continuum limit of granular materials. One possible solution consists in using coarse-grained methods, such as the recently presented spot model [12], in order to reach more realistic system sizes.

In the present work we take a different approach based on high-performance computing. We present an efficient parallel DEM implementation on the basis of the general-purpose, scalable, and portable Parallel Particle Mesh (PPM) library [1]. While the PPM library was so far mostly used for simulations using continuous particle methods, it also supports fully discrete particle simulations [13]. We have thus developed and implemented a specialized DEM client for the PPM library, taking advantage of method-specific optimizations. Within the PPM framework, the actual simulation client is written in a sequential way, using the parallelization primitives of the library. The library then automatically performs the parallelization, including domain decomposition, load balancing, communication scheduling, and distributed file I/O.

Previous approaches to parallel DEM simulations include an adaptation of the molecular dynamics code DL\_POLY [14] to DEM simulations [15] and a specialized parallel code for finite and discrete element simulations based on domain-decomposition and dynamic load (re-)balancing [16]. Due to the limited scalability of these codes, the sizes of the simulated systems were, however, below 500 000 particles.

Using the present implementation, we perform DEM simulations of frictional, viscoelastic spheres using up to 122 million particles on 192 processors. These simulations consider a sand avalanche down an inclined plane [17] and allow us to study front propagation speeds and moving layer characteristics.

Large-scale DEM simulations have a significant potential in many areas of research and technology, including the environmental sciences to, e.g., simulate land slides or avalanches [18, 11], the study of complex self-organization phenomena as, e.g., in singing sand dunes [8], and research toward answering open questions in the continuum modeling of granular matter [6].

## 2 Grain interaction model

In order to realistically model sand and prevent spurious crystallization artifacts, we consider spherical grains with randomly perturbed radii. Each particle is represented by the location of its center  $\mathbf{r}_i$ , its radius  $R_i$ , its mass  $m_i$ , and its polar moment of inertia  $I_i$ . The collision of a particle  $i$  with particle  $j$  is modeled according to Silbert et al. [4] (with the correction published later [19]), generalized to varying radii. The contact forces have a normal component due to elastic deformation of the spheres and a tangential component due to viscous friction. We use the Hertzian contact pressure model. The normal elastic deformation of a contact is given by

$$\delta_{ij} = (R_i + R_j) - r_{ij}, \quad (1)$$

with  $\mathbf{r}_{ij} = \mathbf{r}_i - \mathbf{r}_j$  the distance vector between the two particle centers and  $r_{ij} = \|\mathbf{r}_{ij}\|_2$  its length. The normal and tangential components of the slip velocity at the point of contact are

$$\mathbf{v}_{n_{ij}} = (\mathbf{v}_{ij} \cdot \mathbf{n}_{ij}) \mathbf{n}_{ij}, \quad (2)$$

$$\mathbf{v}_{t_{ij}} = \mathbf{v}_{ij} - \mathbf{v}_{n_{ij}} - (\boldsymbol{\omega}_i R_i + \boldsymbol{\omega}_j R_j) \times \mathbf{n}_{ij}, \quad (3)$$

where  $\mathbf{n}_{ij} = \mathbf{r}_{ij}/r_{ij}$  is the unit normal vector onto the contact plane,  $\boldsymbol{\omega}_i$  the angular velocity of a particle, and  $\mathbf{v}_{ij} = \mathbf{v}_i - \mathbf{v}_j$  the relative velocity between the two particle centers. The evolution of the elastic tangential displacement  $\mathbf{u}_{t_{ij}}$  is governed by

$$\frac{d\mathbf{u}_{t_{ij}}}{dt} = \mathbf{v}_{t_{ij}} \quad (4)$$

and is integrated over the duration of a contact with initial condition  $\mathbf{u}_{t_{ij}}(0) = 0$  at first contact. The normal and tangential forces acting on the particles then become:

$$\mathbf{F}_{n_{ij}} = \sqrt{\frac{\delta_{ij}}{R_i + R_j}} (k_n \delta_{ij} \mathbf{n}_{ij} - \gamma_n m_{\text{eff}} \mathbf{v}_{n_{ij}}), \quad (5)$$

$$\mathbf{F}_{t_{ij}} = \sqrt{\frac{\delta_{ij}}{R_i + R_j}} (-k_t \mathbf{u}_{t_{ij}} - \gamma_t m_{\text{eff}} \mathbf{v}_{t_{ij}}), \quad (6)$$

where  $k_{n,t}$  are the elastic constants in normal and tangential direction, respectively, and  $\gamma_{n,t}$  the corresponding viscoelastic constants. The effective collision mass is  $m_{\text{eff}} = m_i m_j / (m_i + m_j)$ . At each interaction time step, the elastic tangential displacement  $\mathbf{u}_{t_{ij}}$  is truncated in order to enforce Coulomb's law  $\|\mathbf{F}_{t_{ij}}\|_2 < \|\mu \mathbf{F}_{n_{ij}}\|_2$ . Truncating the elastic displacement if this condition is violated amounts to the two spheres slipping against each other without inducing further deformations. The truncation is performed by scaling the tangential force as

$$\mathbf{F}_{t_{ij}} \leftarrow \mathbf{F}_{t_{ij}} \frac{\|\mu \mathbf{F}_{n_{ij}}\|_2}{\|\mathbf{F}_{t_{ij}}\|_2} \quad (7)$$

and adjusting accordingly the stationary elastic displacement:

$$\mathbf{u}_{t_{ij}} = -\frac{1}{k_t} \left( \mathbf{F}_{t_{ij}} \sqrt{\frac{R_i + R_j}{\delta_{ij}}} + \gamma_t m_{\text{eff}} \mathbf{v}_{t_{ij}} \right). \quad (8)$$

The total resultant force on particle  $i$  is then computed by summing the contributions of all particles  $j$  with which it currently interacts, thus:

$$\mathbf{F}_i^{\text{tot}} = m_i \mathbf{g} + \sum_j (\mathbf{F}_{n_{ij}} + \mathbf{F}_{t_{ij}}), \quad (9)$$

where  $\mathbf{g}$  is the acceleration due to gravity. The total torque acting on particle  $i$  is given by

$$\mathbf{T}_i^{\text{tot}} = -R_i \sum_j \mathbf{n}_{ij} \times \mathbf{F}_{t_{ij}}. \quad (10)$$

We integrate the equations of motion using the second-order accurate leap frog scheme

$$\mathbf{v}_i^{n+1} = \mathbf{v}_i^n + \frac{\delta t}{m_i} \mathbf{F}_i^{\text{tot}}, \quad \mathbf{r}_i^{n+1} = \mathbf{r}_i^n + \delta t \mathbf{v}_i^{n+1} \quad (11)$$

$$\boldsymbol{\omega}_i^{n+1} = \boldsymbol{\omega}_i^n + \frac{\delta t}{I_i} \mathbf{T}_i^{\text{tot}}, \quad (12)$$

where  $(\mathbf{r}_i^n, \mathbf{v}_i^n, \boldsymbol{\omega}_i^n)$  denotes the state of the  $i$ -th particle at time step  $n$ , and  $\delta t$  the time step size. The elastic tangential displacement of each particle-particle contact (Eq. 4) is integrated using the explicit Euler scheme with the same time step size.

### 3 Implementation using the PPM library

The present DEM implementation is based on the PPM library [1], which provides an efficient and generic parallelization framework for numerical simulations using particle methods. The PPM library supports a wide range of static and adaptive domain decomposition schemes. These decompositions divide the computational domain into *sub-domains* with sufficient granularity to provide adequate load balancing. The concurrent presence of different decompositions allows to perform each step of the computational algorithm in its optimal environment with respect to load balance and the computation-to-communication ratio. For the actual computations, the individual sub-domains are treated as independent problems and extended with *ghost layers* to allow for communication between them. These ghost layers contain *ghost particles* that are copies of true particles that either reside on a neighboring processor, or account for periodic boundary conditions.

A PPM *topology* is defined by the decomposition of space into sub-domains with the corresponding boundary conditions, and the assignment of these sub-domains onto processors. Sub-domains can be assigned to the processors in various ways. The PPM-internal method assigns contiguous blocks of sub-domains to processors until the accumulated cost for a processor is greater than the theoretical average cost under uniform load distribution. The average is weighted with the relative processor speeds to account for heterogeneous machine architectures. In addition, four different Metis-based [20] and a user-defined assignment are available. Multiple topologies may coexist and library routines are provided to *map* particle data between them as described below.

In order to achieve good load balance, the *SAR heuristic* [21] is used in the PPM library to decide when problem re-decomposition is advised, i.e., when the cost of topology re-definition is amortized by the gain in load balance. Moreover, all topology definition routines can account for the true computational cost of each particle, for example defined by the actual number of its interactions, and the effective speeds of all processors.

PPM topologies implicitly define a data-to-processor assignment. Mapping routines provide the functionality of sending particles to the proper processor, i.e., to the one that “owns” the corresponding sub-domain(s) of the computational space. Three different mapping types are provided:

- a *global mapping*, involving an all-to-all communication among processors,
- a *local mapping* for neighborhood communication, and
- *ghost mappings* to update the ghost layers.

The global mapping is used to perform the initial data-to-processor assignment or to switch from one topology to another (e.g. after problem re-decomposition for load balancing reasons), whereas the local mapping is mainly used to account for particle motion during a simulation. Communication is scheduled by solving the minimum edge coloring problem using the efficient approximation algorithm by Vizing [22]. Ghost mappings are provided to receive ghost particles and to send ghost contributions back to the corresponding real element.

For efficiency, all mapping types are organized as stacks. A mapping operation consists of four steps: (1) defining the mapping, (2) pushing data onto the send stack, (3) performing the actual send and receive operations, and (4) popping the data from the receive stack. This architecture allows data stored in different arrays to be sent together to minimize network latency, and mapping definitions to be re-used by repeatedly calling the push/send/pop sequence for the same, persisting mapping definition.

The evaluation of Particle-Particle (PP) interactions is a key component of DEM as the overall dynamics of the system are governed by local particle interactions. The PPM library implements symmetric and non-symmetric PP computations using a novel type of Cell lists [1] and Verlet lists [23].

The present DEM implementation uses asymmetric PP interactions on the basis of Verlet lists [23] that are constructed in linear time using an intermediate Cell list. The size of the Verlet sphere is chosen slightly larger than the particle diameter to capture all possible contacts. In order to prevent an update of the Verlet list at every time step, we enlarge the Verlet sphere by a safety margin, called the *skin*. Thus, the interaction lists only need to be reconstructed once a particle has traveled more than the skin thickness. In our implementation we use the conservative lower bound

$$N_{Verlet} = \frac{c - d_{max}}{2\max(v_i)\delta t} \quad (13)$$

for the number of time steps between two list updates. Hereby,  $c$  denotes the size of the Verlet sphere including the skin,  $d_{max}$  the diameter of the largest particle in the domain and  $\max(v_i)$  the magnitude of the largest occurring particle velocity. The serial implementation of the Verlet list is at most  $3^4/4\pi \approx 6$  times faster than the corresponding Cell list. However, the parallel implementation of the Verlet list offers the additional improvement that the particles need only be mapped onto the processors after the Verlet list is updated. Between Verlet list updates the only required communication is the *state* of the *ghost* particles – the particular ghost mapping remains constant and is stored and used repeatedly until the next Verlet update.

## 4 Results for a sand avalanche

In order to simulate an avalanche down an inclined plane, we use the grain interaction model as described above, with the following parameter values:  $k_n = 7849 \text{ N/m}$ ,  $k_t = 2243 \text{ N/m}$ ,  $\gamma_n = 3401 \text{ N/m}$ ,  $\gamma_t = 3401.0 \text{ s}^{-1}$ , and  $\mu = 1$ . The reference system contains 635 780 grains with uniformly distributed radii between 1.00 and 1.12 mm placed in a computational domain of size  $1800 \times 8.6 \times 250 \text{ mm}$ . The grains are initially placed on a regular lattice with a void fraction of 0.47 and confined in horizontal and vertical direction by solid walls modeled by two layers of immobile grains. Periodic boundary condition are imposed in the span-wise  $y$ -direction. The radius of the Verlet sphere is  $c = 15 \text{ mm}$ , and the time step size is  $10^{-6} \text{ s}$ . These parameters require an update of the Verlet lists every 150 time steps.

All simulations are performed on the Cray XT3 computer at the Swiss National Supercomputing Centre (CSCS). This distributed memory machine has 1664 AMD Opteron


processors at 2.6 GHz clock speed (5.2 GFlop/s peak performance). Each processor has 1 GB of main memory, and they are interconnected in a full 3D torus network topology through a 6-way Cray SeaStar network. The machine is running version 1.4 of the UNICOS/lc operating system, and all codes are compiled using the Portland Group Fortran 90 compiler, version 6.1-4 (64bit).

The performance of the simulations is tested for both a fixed-size and a scaled-size problem. In the fixed-size problem, the number of particles is kept constant, i.e., the work load per processor decreases with increasing number of processors. In the scaled problem, the particle number grows proportionally to the number of processors, resulting in a constant work load per processor. In each test, we measure the elapsed wall-clock time  $t_{ij}$  for each time step  $j$  on each processor  $i = 1, \dots, N_{proc}$ . To account for synchronous communication steps, we report the maximum of these times over all processors to compute speedup  $S$  and efficiency  $e$ :

$$S(N_{proc}) = \frac{t(1)}{\text{mean}_j \max_i t_{ij}(N_{proc})} \cdot \frac{N(N_{proc})}{N(1)}, \quad (14)$$

$$e(N_{proc}) = \frac{S(N_{proc})}{N_{proc}}, \quad (15)$$


where  $t(1)$  is the time on a single processor (linearly extrapolated if not measured),  $t_{ij}(N_{proc})$  is the time on  $N_{proc}$  processors,  $N(1)$  is the problem size on a single processor, and  $N(N_{proc})$  is the problem size on  $N_{proc}$  processors.



**Fig. 1.** Parallel speedup and efficiency of the PPM DEM client for the scaled-size problem. The initial number of particles on two processors is 635 780, scaling up to 122 million particles on 192 processors. The PPM library decomposes the domain into 4213 sub-domains that are automatically distributed among the processors. All timings are performed on the Cray XT3.

The performance of the code is shown in Figs. 1 and 2 for the scaled and fixed size problems, respectively. We observe a good speedup up to 192 processors and a parallel efficiency of 40% for both the scaled and fixed-size problem on 192 processors. The scaled problem uses up to 122 million fully resolved particles and takes less than 2.6 second of wall-clock time per time step. One time step of the fixed problem with 635 780 particles takes 1 second of wall-clock time on a single processor and 0.011 seconds on 192 processors. For the ten-fold larger fixed-size problem, one time step takes 0.13 seconds on 192 processors, thus demonstrating the good linear scaling of the simulation code with particle number.

Figure 3 shows snapshots from one of the avalanche simulations. The system is first equilibrated in the horizontal position for 400 000 time steps before the box is tilted and the right wall removed. The evolution of the avalanche is then simulated for 1.2 million time steps. The relatively low damping of the grains results in a spray-like front (Fig. 3b-c). The front is bend upward and forms a jet as it impacts onto the lower wall (Fig. 3e-f).



**Fig. 2.** Parallel speedup and efficiency of the PPM DEM client for fixed-size problems with 635 780 (+), 3 178 900 (×), and 6 357 800 (\*) particles on up to 192 processors. The PPM library decomposes the domain into 4213 sub-domains that are automatically distributed among the processors. All timings are performed on the Cray XT3.

## 5 Conclusions

We have presented an efficient parallel DEM implementation based on the PPM library. Using the particle interaction model by Silbert et al. [4, 19], the implementation has maintained a parallel efficiency of 40% on up to 192 processors of a Cray XT3 machine, and we have presented results from simulations using up to 122 million fully resolved, frictional, viscoelastic particles. This constitutes the first fully discrete application of the PPM library and, to our knowledge, the largest DEM simulations performed so far. A simulation with 630 000 particles takes less than one second per time step on a single processor. Since the present simulation code scales well with both the number of particles and the number of processors, this translates to less than 1 second per time step using 40.7 million particles on 128 processors, or 0.017 seconds per time step using 635 780 particles on 128 processors.


We have applied the new code to the simulation of sand avalanches down an inclined plane. This system is frequently used as a model to study the macroscopic behavior of granular flows [2, 3]. Future work will include calibrating the model parameters by comparing the front propagation speeds to laboratory experiments of sand avalanches, and using the simulation code to help understand the global behavior of granular flows, including phenomena such as fingering [24], shear zones [25, 26], and generation of sound [8]. The large number of particles that can be treated by the present implementation will allow the extraction of macroscopic material properties [5] and the validation of the corresponding governing equations on the continuum level [7].

## 6 Acknowledgements

The authors gratefully acknowledge the experimental validation data provided by Stéphane Douady and Sylvain Courrech du Pont of École Normale Supérieure, Paris, France. We also thank the Swiss National Supercomputing Centre (CSCS) for providing the compute resources required by this work. JHW acknowledges additional support by the Julie Damm’s Studiefond.

## References

1. Sbalzarini, I.F., Walther, J.H., Bergdorf, M., Hieber, S.E., Kotsalis, E.M., Koumoutsakos, P.: PPM – a highly efficient parallel particle-mesh library for the simulation of continuum systems. *J. Comput. Phys.* **215**(2) (2006) 566–588



**Fig. 3.** Visualization of individual grains from direct DEM simulation of a sand avalanche down an inclined plane. Initially, 120 000 grains with a void fraction of 0.468 and uniformly distributed radii between 1.00 and 1.12 mm are equilibrated in a box of dimensions  $740 \times 200 \times 8$  mm. At time 0 the box is inclined at an angle of  $31.2^\circ$  and the right wall is removed. We show the flow of the avalanche down the plane 0.1, 0.2,  $\dots$ , 0.6 seconds after removing the wall (a-f).

2. Daerr, A., Douady, S.: Two types of avalanche behaviour in granular media. *Nature* **399** (1999) 241–243
3. Douady, S., Andreotti, B., Daerr, A., Cladé, P.: From a grain to avalanches: on the physics of granular surface flows. *C. R. Physique* **3** (2002) 177–186
4. Silbert, L.E., Ertas, D., Grest, G.S., Halsey, T.C., Levine, D., Plimpton, S.J.: Granular flow down an inclined plane: Bagnold scaling and rheology. *Phys. Rev. E* **64** (2001) 051302
5. Hunt, M.L.: Discrete element simulations for granular material flows: effective thermal conductivity and self-diffusivity. *Int. J. Heat Mass Transfer* **40**(13) (1997) 3059–3068
6. de Gennes, P.G.: Granular matter: a tentative view. *Rev. Mod. Phys.* **71**(2) (1999) 374–382
7. Douady, S., Andreotti, B., Daerr, A.: On granular surface flow equations. *Eur. Phys. J. B* **11** (1999) 131–142
8. Douady, S., Manning, A., Hersen, P., Elbelhiti, H., Protière, S., Daerr, A., Kabbachi, B.: Song of the dunes as a self-synchronized instrument. *Phys. Rev. Lett.* **97** (2006) 018002
9. Landry, J.W., Grest, G.S., Silbert, L.E., Plimpton, S.J.: Confined granular packings: Structure, stress, and forces. *Phys. Rev. E* **67** (2003) 041303
10. Heyes, D.M., Baxter, J., Tüzün, U., Qin, R.S.: Discrete-element method simulations: from micro to macro scales. *Phil. Trans. R. Soc. Lond. A* **362** (2004) 1853–1865
11. Richards, K., Bithell, M., Dove, M., Hodge, R.: Discrete-element modelling: methods and applications in the environmental sciences. *Phil. Trans. R. Soc. Lond. A* **362** (2004) 1797–1816
12. Rycroft, C.H., Bazant, M.Z., Grest, G.S., Landry, J.W.: Dynamics of random packings in granular flow. *Phys. Rev. E* **73** (2006) 051306
13. Altenhoff, A., Walther, J.H., Koumoutsakos, P.: A stochastic boundary forcing for dissipative particle dynamics. *J. Comput. Phys.* (2007) in press
14. Todorov, I.T., Smith, W.: DL\_POLY\_3: the CCP5 national UK code for molecular-dynamics simulations. *Phil. Trans. R. Soc. Lond. A* **362** (2004) 1835–1852
15. Dutt, M., Hancock, B., Bentham, C., Elliott, J.: An implementation of granular dynamics for simulating frictional elastic particles based on the DL\_POLY code. *Comput. Phys. Comm.* **166** (2005) 26–44
16. Owen, D.R.J., Feng, Y.T.: Parallelised finite/discrete element simulation of multi-fracturing solids and discrete systems. *Engineering Computations* **18**(3/4) (2001) 557–576
17. Balmforth, N.J., Kerswell, R.R.: Granular collapse in two dimensions. *J. Fluid Mech.* **538** (2005) 399–428
18. Cleary, P.W., Prakash, M.: Discrete-element modelling and smoothed particle hydrodynamics: potential in the environmental sciences. *Phil. Trans. R. Soc. Lond. A* **362** (2004) 2003–2030

19. Silbert, L.E., Grest, G.S., Plimpton, S.J.: Boundary effects and self-organization in dense granular flows. *Phys. Fluids* **14**(8) (2002) 2637–2646
20. Karypis, G., Kumar, V.: A fast and high quality multilevel scheme for partitioning irregular graphs. *SIAM J. Sci. Comput.* **20**(1) (1998) 359–392
21. Moon, B., Saltz, J.: Adaptive runtime support for direct simulation Monte Carlo methods on distributed memory architectures. In: *Proceedings of the IEEE Scalable High-Performance Computing Conference, IEEE* (1994) 176–183
22. Vizing, V.G.: On an estimate of the chromatic class of a p-graph. *Diskret. Anal.* **3** (1964) 25–30 in Russian.
23. Verlet, L.: Computer experiments on classical fluids. I. Thermodynamical properties of Lennard-Jones molecules. *Phys. Rev.* **159**(1) (1967) 98–103
24. Pouliquen, O., Delour, J., Savage, S.B.: Fingering in granular flows. *Nature* **386**(816–817) (1997)
25. Fenistein, D.: Wide shear zones in granular bulk flow. *Nature* **425** (2003) 256
26. Cheng, X., Lechman, J.B., Fernandez-Barbero, A., Grest, G.S., Jaeger, H.M., Karczmar, G.S., Möbius, M.E., Nagel, S.R.: Three-dimensional shear in granular flow. *Phys. Rev. Lett.* **96** (2006) 038001