A STUDY ON THE INFLUENCE OF THE PARTICLE PACKING FRACTION ON THE PERFORMANCE OF A MULTILEVEL CONTACT DETECTION ALGORITHM

V. OGARKO AND S. LUDING

Multi Scale Mechanics (MSM), CTW, UTwente,
PO Box 217, 7500 AE Enschede, Netherlands;
v.ogarko@utwente.nl, s.luding@utwente.nl

Key words: contact detection, discrete element, polydisperse, particle size distribution, packing fraction

Abstract. We investigate the influence of the packing fraction of highly polydisperse particle systems on the performance of a high-performance multilevel contact detection algorithm as applied for molecular dynamics type simulations. For best performance, this algorithm requires two or more hierarchy levels in order to cope with the strongly different particle size classes. In order to predict the optimal number of levels, an empirical parameter corresponding to the "overhead" is identified. For homogeneous systems, the density is not much affecting the performance of the algorithm, however, the optimal number of levels slightly increases with density as well as the speed-up as compared to a single-level method.

1 INTRODUCTION

Contact detection is a fundamental computational problem arising in computer simulations of systems consisting of many discrete objects such as particles or atoms. The particle based modeling methods like the Discrete Element Method (DEM) [1] play an important role for physics-based simulations in many diverse fields. The performance of the computation relies on several factors, which include the physical model, on the one hand, and the contact detection method used, on the other. The collision detection of short-range pairwise interactions between particles is usually one of the most time-consuming tasks in calculations [2].

The most commonly used method for contact detection of nearly monosized particles with short-ranged forces is the Linked-Cell method [3]. Nevertheless, the Linked-Cell method is unable to efficiently deal with particles of greatly varying sizes [4]. This can effectively be addressed by the use of methods based on *hierarchical grids* [4, 5, 6].

The multilevel contact detection algorithm presented in Ref. [6] contains the number of hierarchy levels and the cell sizes at each level as adjustable parameters. A method to find the optimal parameters for an arbitrary polydisperse size distribution of objects is also shown in [6]. As an input parameter, this method requires the "overhead" of the data structure. In addition to the detailed study of different size distributions in Ref. [6], we examine the influence of the density of the particle system on the performance of the algorithm [6] and on the "overhead" of the data structure by means of molecular dynamics simulations.

2   CONTACT DETECTION ALGORITHM

In this section we briefly describe the main ideas of the multilevel contact detection algorithm and the method to find the optimal parameters, which are necessary for understanding the results of this study. For more details see Ref. [6].

2.1   Algorithm

The algorithm is designed to determine all the pairs in a set of N spherical particles in a d-dimensional Euclidean space that are in close neighborhood and possibly do overlap. A *hierarchial grid* structure used in the algorithm is a set of L regular grids with different cell sizes. Every regular grid is associated with a hierarchy level $h \in [1, L]$, where L is the integer number of hierarchy levels. Each level h has a different cell size $s_h \in \mathbb{R}$, where the cells are d-dimensional cubes. Grids are ordered with increasing cell size so that $h = 1$ corresponds to the grid with smallest cell size, i.e. $s_h < s_{h+1}$. For a given number of levels and cell sizes, the hierarchical grid cells are defined by the following spatial mapping of points $\vec{x} \in \mathbb{R}^d$ to a cell at specified level h:

$$(\vec{x}, h) \mapsto (\lfloor x_1/s_h \rfloor, ..., \lfloor x_d/s_h \rfloor, h), \tag{1}$$

where $\lfloor x \rfloor$ denotes the floor function [1].

The algorithm is made up of two phases. In the first "mapping phase" all the particles are mapped into hierarchical grid cells based on their location and size as follows. The diameter of the smallest particle within a given level of hierarchy should be larger than the cell size of the previous level of hierarchy, and the diameter of the largest particle at some hierarchical level should be smaller or equal to the cell size of this level. The cell-sizes and the distribution of particles among levels are detailed below.

In the second "contact detection phase" for every particle in the system the potential contact partners are determined, and the geometrical intersection tests with them are made. The contact detection is split into two steps, and the search is done by looping over all particles and performing the first and second steps consecutively for each particle. The first step is the contact search at the level of insertion of a given particle, using the classical Linked-Cell method [3]. The second step is the *cross-level search*. For a

---

[1]The largest integer not greater than x

given particle, one searches for potential contacts at every level h lower than the level of insertion. For detailed description of these steps and further optimizations see Ref. [6].

2.2 Optimal parameters

It is shown in Ref. [6] that in the case of polydisperse systems, i.e., systems where all the particles' sizes are different, the performance of the multilevel contact detection algorithm strongly depends on the selected parameters (the number of levels L and cell sizes s_h). Furthermore, a hypothesis on the optimal distribution of particles by levels was formulated in [6], which states:

Hypothesis 1 *Let m_h be the average number of particles per cell at level h, that is, $m_h = N_h/N_h^c$, where N_h is the number of particles at level h, and N_h^c is the number of cells at this level. Then the optimal distribution of particles by levels satisfies the following condition:*

$$m := m_i = m_j, \; for \; all \; \; i,j \in [1, L]. \tag{2}$$

If the particles are mapped to the levels, such that Eq. (2) is approximated, the CPU time spent for contact detection, T_{CD}, scales as

$$T_{CD} \sim NL(m + K), \tag{3}$$

where K is a constant corresponding to the "overhead" of the algorithm, i.e., the time spent to access cells to be tested, and $m = m(L, \text{PSD})$ is the number of particles per cell (see Ref. [6]). To compute m for a given L and for the particle system at hand, one needs to choose cell sizes s_h so that Eq. (2) is approximately satisfied. How to do this is explained in Ref. [6]. In the next section we investigate the influence of the volume fraction of the particle system, i.e., the ratio between the volume of the particles and the volume of the system, on the value of the "overhead" K.

3 SIMULATION RESULTS

We use homogeneous and isotropic disordered systems of colliding elastic spherical particles in a cubical box with hard walls. The motion of particles is governed by Newton's second law with a linear elastic contact force during overlap. Every particle undergoes only translational motion (without rotation) and gravity is set to zero.

3.1 Experimental setup

The systems are prepared in two stages. Starting from a random uniform distribution of points in a cubical box, the radius of the particles grows linearly with time. We use non-overlapping spheres [7], with an Event-Driven code [8], whose growing rate conforms to and conserves a prescribed size distribution. Initial velocities are set randomly in order to keep the system dynamic and random, for details see Ref. [8]. When the target volume fraction is reached, the growth process is stopped.

Figure 1: Particle systems with $N = 125001$ with (a) volume fraction $\nu = 0.1$, (b) volume fraction $\nu = 0.4$, with uniform volume distribution. The size ratio ω is 50. Colour is by relative size.

With the final configuration from the first step, the simulation switches to the relaxation stage with "soft" particles, i.e., particles move according to interparticle forces [1]. The linear elastic normal contact force model is used [9], which leads to a certain contact duration. At the beginning of this stage the velocities of the particles are scaled in a way that a collision between two of the smallest particles would reach an average maximum overlap of one percent of their radius. We let the simulations run for a few collisions per particle for equilibration before making the measurement of the performance, in order to have contacts (overlaps) between particles in the system.

We use uniform volume size distribution, i.e., the probability distribution of the volumes of the particles is constant. This type of size distribution implies that the particle system contains rather many small particles. We denote the size ratio between the radius of the largest and the smallest particle in the system as ω.

3.2 Numerical performance results

We measure the total CPU time for varying number of hierarchy levels L for particle systems with different volume fractions. For this, we calculate the value of m for each $L \in [1, 50]$ using the method given in Ref. [6], and run simulations where the multilevel grid is used with L levels, and the cell sizes s_h are computed in accordance with hypothesis (1). To present the total CPU time, we use the slowdown factor SF, that is the total CPU time divided by the smallest CPU time for a given system. In Fig. 2 the results of this numerical experiment are shown for systems with uniform volume distribution, with $N = 125001$ and $\omega = 50$. The analytical prediction (3) is also plotted with $K = 0.2$,

scaled in such a way that SF $= 1$ corresponds to the minimum of the right hand side of Eq. (3). Note that even though the prediction (3) is for CPU time spent only for contact detection, T_{CD}, the total CPU time for fixed number of particles also scales as T_{CD}. This is because the CPU time spent in the force calculation and integration does not depend on the grid parameters used. From the numerical experimental results shown in Fig. 2 it can be seen that the prediction of Eq. (3) used with $K = 0.2$ can be used for determining the optimum number of levels, and the value of K is not dependent on the packing fraction. We also confirm that the dependence on density is very weak for rather large packing fractions of 0.62 (data from Ref. [6] – not shown here). Furthermore, the optimal number of levels, L^*, slightly increases with the volume fraction from $L^* = 5$ for $\nu = 0.1$ to $L^* = 8$ for $\nu = 0.62$, see Fig. 3(a). Finally, the best speed-up over the Linked-Cell method ($L = 1$), S, increases nearly linearly with the volume fraction, as it can be seen from Fig. 3(b). Note that even for relatively low volume fraction $\nu = 0.1$, the speed-up is high, i.e., $S \approx 150$, so the use of the multilevel algorithm is highly advantageous for this type of particle size distributions (many small particles) for the whole range of volume fractions.

(a)   (b)

Figure 2: Slowdown factor for different numbers of levels for systems with $N = 125001$ and $\omega = 50$ with uniform volume distribution; (a) volume fraction $\nu = 0.1$, (b) volume fraction $\nu = 0.4$. The prediction of Eq. (3) is used with $K = 0.2$ for plotting solid lines. Note that the data can be obtained only for integer values of L.

Figure 3: (a) The optimal number of levels for different densities. (b) The speed-up relative to the single-level contact search for different packing fractions. The data for $\nu = 0.62$ is taken from Ref. [6].

4 SUMMARY AND CONCLUSIONS

We have shown that the multilevel algorithm presented in Ref. [6] can be used to significantly improve the performance of simulations of particle systems with rather many small particles at very different volume fractions. The speed-up relative to the single-level Linked-Cell method increases nearly linearly with the volume fraction for the homogeneous systems tested. Moreover, the value of the "overhead" parameter, which is necessary to analytically find the optimal number of hierarchy levels is not dependent on the packing fraction for these systems. This leads to an increasing number of levels with the density for optimal performance of the contact detection.

ACKNOWLEDGMENTS

REFERENCES

[1] P. A. Cundall and O. D. L. Strack. A discrete numerical model for granular assemblies. *Géotechnique*, **29** (1), 47-65 (1979).

[2] J. R. Williams and R. O'Connor. Discrete element simulation and the contact problem. *Archives of Computational Methods in Engineering*, **6** (4), 279-304 (1999).

[3] M. P. Allen and D. J. Tildesley, *Computer Simulations of Liquids.* Clarendon Press, 1989.

[4] T. Iwai, C. W. Hong, and P. Greil. Fast particle pair detection algorithms for particle simulations. *International Journal of Modern Physics C*, **10** (5), 823-837 (1999).

[5] B. Mirtich, Efficient Algorithms for Two-Phase Collision Detection in *Practical Motion Planning in Robotics: Current Approaches and Future Directions*, Ed. K. Gupta and A. P. del Pobil, Wiley, New York, 1998.

[6] V. Ogarko and S. Luding, A fast multilevel algorithm for contact detection of arbitrarily polydisperse objects, submitted 2011.
`http://www2.msm.ctw.utwente.nl/sluding/PAPERS/HGrid.pdf`

[7] B. D. Lubachevsky and F. H. Stillinger. Geometric Properties of Random Disk Packings. *J. Stat. Phys.*, **60**, 561-583 (1990).

[8] M. Skoge, A. Donev, F. H. Stillinger and S. Torquato. Packing hyperspheres in high-dimensional Euclidean spaces. *Phys. Rev. E*, **74** (4), 041127 (2006).

[9] S. Luding. Cohesive, frictional powders: contact models for tension. *Granular Matter*, **10** (4), 235-246 (2008).