Comparing granular dynamics vs. fluid dynamics via large DOF-count parallel simulation on the GPU

Milad Rakhsha, Conlain Kelly, Nicholas Olsen, Lijing Yang, Radu Serban, Dan Negrut
{rakhsha,ckelly5,nicholas.olsen,lyang296,serban,negrut}@wisc.edu
Mechanical Engineering Department, University of Wisconsin-Madison
Madison, WI

ABSTRACT

In understanding the dynamics of granular systems, a *discrete* modeling approach tracks the motion of all particles. Such an approach is computationally demanding especially when the number of particles is large, e.g., when going beyond tens of millions of particles. In these cases, one can contemplate switching to *continuum* models, which are computationally less expensive. In order to assess when such a discrete to continuum switch is justified, we report herein results for a comparison between the dynamics of granular and fluid flows obtained with an open source code that scales to handle granular systems with more than 1 billion degrees of freedom (DOFs); i.e., two orders of magnitude higher than the state of the art. On the one hand, we solve the Newton-Euler equations of motion, which govern the time evolution of the granular system. On the other hand, we solve the Navier-Stokes equations that describe the time evolution of the fluid model. Both the multibody and fluid dynamics solvers leverage parallel computing on the Graphics Processing Unit (GPU). We report similarities and differences between the dynamics of the discrete, fully resolved system and the continuum granular material model via a set of numerical experiments that include both static and highly transient scenarios. The simulation platform that anchors this contribution is publicly available on GitHub and is part of an open source code called Chrono.

KEYWORDS

granular dynamics, fluid dynamics, DEM, SPH, GPU computing

1 INTRODUCTION

Fully resolved simulation of large scale granular flows is a challenging task owing to (*i*) their complex frictional contact interactions, and (*ii*) a high computational burden as the motion of individual particles needs to be tracked in a numerical integration framework that advances simulation at time steps $\Delta t \approx 10^{-5}$ or lower. To address the latter aspect, discrete granular flow simulations use parallel computing to accelerate the computations. To the best of our knowledge, the largest granular dynamics simulation of practical relevance to date contained 2.4 billion bodies which was run on 16,384 CPUs (131,072 cores) of Japan's K-computer [4, 5], the 2012 fastest supercomputer in the world and now the 18th in the ranking of the world's supercomputers [13].

The collective behavior of individual particles may, however, be regarded as a continuum. Given that granular media can demonstrate distinctively different behaviors under various local stress conditions, we focused on fluid-like evolution of granular material that happens when it is rapidly sheared. In such cases, (*i*) shear stress in granular material shows strain-rate dependency, which is a characteristic of fluids and, (*ii*) there exists a threshold value (yield criterion), under which the grains do not flow, a characteristic of solid materials. These features suggest a viscoplastic behavior of a granular material which is similar in nature to that of non-Newtonian fluids such as Bingham liquids [1, 9].

The contributions reported in this poster are twofold. First, we demonstrate that continuum modeling represents an affordable and efficient alternative for the *discrete* simulation of granular flows. Both the discrete and continuum solutions leverage GPU computing. Second, we show that if need be, we can scale the *discrete* simulation of granular flows to levels beyond the state of the art when using commodity GPU architectures.

2 GOVERNING EQUATIONS AND NUMERICAL METHODS

2.1 Discrete Model via Newton-Euler

We rely on a penalty-based Discrete Element Method (DEM) drawing on [2], wherein particles in mutual contact experience a small overlap/local deformation that yields a normal contact force. Advancing the simulation by one time step takes place in four stages: (*i*) broad-phase collision detection (CD); (*ii*) narrow-phase CD; (*iii*) force computation; and (*iv*) time integration. In the GPU implementation, the forces can be computed as soon as contacts are detected, and the broad-phase carried out as soon as positions have been updated. Thus steps (*ii*) and (*iii*) can be combined, as can (*i*) and (*iv*). Against this backdrop, for a system of discrete elements, the motion of an element i is governed by

$$m\frac{d\mathbf{v}_i}{dt} = m\mathbf{g} + \sum_{i\neq j}\left(\mathbf{F}_n^{ij} + \mathbf{F}_t^{ij}\right) \;, \qquad (1)$$

Above, for each pair of contacting particles (i, j) the following describe the normal and tangential forces, respectively [2]:

$$\mathbf{F}_n^{ij} = f\left(\frac{\delta^{ij}}{2R}\right)\left(k_n\delta^{ij}\mathbf{n}^{ij} - \gamma_n\bar{m}\mathbf{v}_n^{ij}\right) \;, \qquad (2a)$$

$$\mathbf{F}_t^{ij} = f\left(\frac{\delta^{ij}}{2R}\right)\left(-k_t\mathbf{u}_t^{ij} - \gamma_t\bar{m}\mathbf{v}_t^{ij}\right) \text{ where } |\mathbf{F}_t^{ij}| \le \mu|\mathbf{F}_n^{ij}| \;, \qquad (2b)$$

where $\bar{m}$ is the effective mass; R is the contact radius of curvature; δ^{ij} is the normal penetration; $\mathbf{v}_n^{ij}$ and $\mathbf{v}_t^{ij}$ are the relative normal and tangential velocities of the contact points; $\mathbf{u}_t^{ij}$ is the tangential displacement or "creep" [2]; and $f(x)$ is a term depending on the force model, where $f(x) = 1$ for a Hookean model and $f(x) = \sqrt{x}$ for a Hertzian model. Details are provided in [3, 10].

2.2 Continuum Model via Navier-Stokes

The mass and momentum balance; i.e., the continuity and Navier-Stokes, equations are formulated for the fluid phase as [7]

$$\frac{d\rho}{dt} = -\rho\nabla\cdot\mathbf{u} \tag{3}$$

$$\frac{d\mathbf{u}}{dt} = \frac{1}{\rho}\nabla\cdot\boldsymbol{\sigma} + \mathbf{f}^b = -\frac{1}{\rho}\nabla p + \frac{1}{\rho}\nabla\cdot\boldsymbol{\tau} + \mathbf{f}^b \ , \tag{4}$$

where

$$\boldsymbol{\sigma} = -p\mathbf{I} + \boldsymbol{\tau} \tag{5}$$

is the stress tensor; p and τ are the volumetric and deviatoric decomposition of the stress tensor. Note that p can be expressed by negative one-third of the trace of the stress tensor (average of normal stresses), i.e. $p = -\frac{1}{3}(\sigma_{xx} + \sigma_{yy} + \sigma_{zz})$ and it is a mechanical property. Amongst the possible choices of material models, the *Newtonian* constitutive model expresses that $\tau_{ij} = \mu(\frac{\partial u_i}{\partial x_j} + \frac{\partial u_j}{\partial x_i})$. Along with the *incompressible flow* assumption ($\nabla\cdot\mathbf{u} = 0$), one can simplify the momentum equations as follows:

$$\frac{d\mathbf{u}}{dt} = -\frac{1}{\rho}\nabla p + \nu\nabla^2\mathbf{u} + \mathbf{f}^b \ , \tag{6}$$

where $\nu = \mu/\rho$ and ρ are the fluid kinematic viscosity and density, respectively; $\mathbf{f}^b$ is the density of volumetric forces such as gravity, and $\mathbf{v}$ and p are the flow velocity and pressure, respectively.

The numerical discretization of the above equations was performed via Smoothed Particle Hydrodynamics (SPH). Details of the space discretization and time integration used herein are available in [11, 12].

3 RESULTS

3.1 Scaling Analysis of Discrete Model

Preliminary results for a typical benchmark simulation (spheres dropped in a large box) [6, 8, 14] indicate linear performance scaling up to 250 million bodies on an Nvidia Tesla P100 GPU with 16GB of device memory (see Fig. 1). After 250 million bodies, the code overruns the device's memory and begins paging in and out of host memory. These results can likely be improved by memory prefetching or access "hints" to the GPU. The scaling analysis reported in Fig. 1 has been performed for two generations of Nvidia hardware: Pascal, and the more recent Turing architecture. As expected, there is a noticeable speed-up when using the more recent architecture. However, the high-end Pascal GPU has about 50% more memory, which allows it to reach large element counts. A test on an Nvidia Volta V100 GPU has scaled up to 510 million frictionless spheres. As each sphere has 3 degrees of freedom, this provides a system with 1.5 billion DOFs that runs on a single GPU. The state of the art for GPU computing allows scaling up to 10 million elements before a multi-GPU, MPI-enabled solution is considered [6].

3.2 Discrete vs. Continuum: Model Comparison

In the accompanying poster we report on a set of numerical experiments that compare the physics of granular and fluid media. Herein, we highlight the dam break with rigid cylinder. In this problem, a rigid cylindrical obstacle is placed in front of a moving wave of granular material. The quantity of interest is the overall force experienced by the cylinder over time as the granular waves

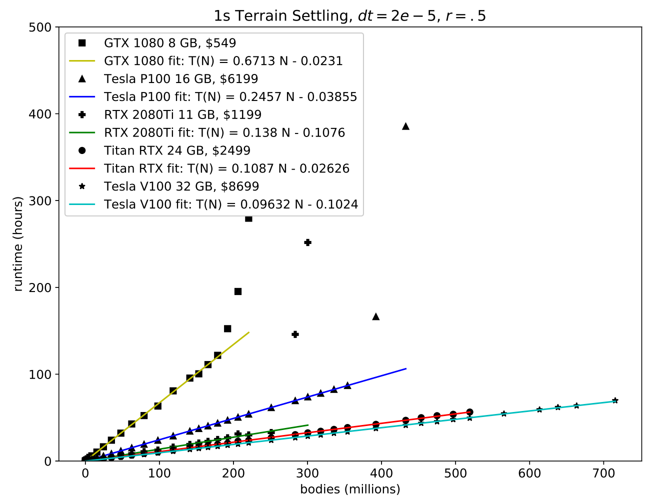

Figure 1: Scaling analysis of a 2-second terrain settling simulation

clears it. Figure 2 reports the force experienced by the cylinder for both the granular and fluid material for different fluid's viscosities, and different grain sizes. In the poster presentation, snapshots of the fluid and granular flow simulations are shown for a qualitative comparison of the time evolution of the two media.

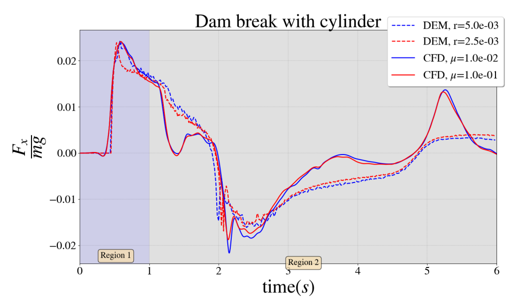

Figure 2: Comparison of the normalized horizontal force experienced by the cylinder in the dam break simulation for different viscosities of the fluid and for different particle diameters of the granular media.

4 CONCLUSION

We present two fundamentally different methodologies for simulating granular flows. The discrete model relies on DEM; the continuum model uses SPH. Both approaches are scalable in a single-instruction-multiple-data (SIMD) sense. Both solvers leverage single-GPU computing to accelerate simulation. For the discrete representation we report system sizes that go two orders of magnitude beyond the state of the art. In terms of the underlying physics, a set of numerical experiments demonstrate that a non-Newtonian fluid model is a good proxy for fully resolved granular flows.

ACKNOWLEDGMENT

Support for this project has been provided in part by US Army Research Office grant W911NF-18-1-0476.

REFERENCES

[1] Charles S Campbell. 1990. Rapid granular flows. *Annual Review of Fluid Mechanics* 22, 1 (1990), 57–90.

[2] P. Cundall and O. Strack. 1979. A discrete element model for granular assemblies. *Geotechnique* 29 (1979), 47–65.

[3] J. A. Fleischmann, R. Serban, D. Negrut, and P. Jayakumar. 2015. On the Importance of Displacement History in Soft-Body Contact Models. *ASME JCND* (2015). doi: 10.1115/1.4031197.

[4] M. Furuichi, D. Nishiura, M. Asai, and T. Hori. 2017. Poster: The first real-scale DEM simulation of a sandbox experiment using 2.4 billion particles. (2017). https://uwmadison.box.com/s/wmejroyqjv03mg4x7r9cq5vc0aj6e7c4 Supercomputing Conference.

[5] Mikito Furuichi, Daisuke Nishiura, Osamu Kuwano, Arthur Bauville, Takane Hori, and Hide Sakaguchi. 2018. Arcuate stress state in accretionary prisms from real-scale numerical sandbox experiments. *Nature Scientific Reports - www.nature.com/scientificreports/* 8 (12 2018).

[6] JQ Gan, ZY Zhou, and AB Yu. 2016. A GPU-based DEM approach for modeling of particulate systems. *Powder Technology* 301 (2016), 1172–1182.

[7] Morton E Gurtin, Eliot Fried, and Lallit Anand. 2010. *The mechanics and thermodynamics of continua.* Cambridge University Press.

[8] QF Hou, ZY Zhou, and AB Yu. 2012. Micromechanical modeling and analysis of different flow regimes in gas fluidization. *Chemical Engineering Science* 84 (2012), 449–468.

[9] Pierre Jop, Yoël Forterre, and Olivier Pouliquen. 2006. A constitutive law for dense granular flows. *Nature* 441, 7094 (2006), 727–730.

[10] A. Pazouki, M. Kwarta, K. Williams, W. Likos, R. Serban, P. Jayakumar, and D. Negrut. 2017. Compliant versus rigid contact: A comparison in the context of granular dynamics. *Phys. Rev. E* 96 (2017). Issue 4. https://doi.org/DOI:10.1103/PhysRevE.96.042905

[11] M. Rakhsha, A. Pazouki, R. Serban, and D. Negrut. 2019. Using a half-implicit integration scheme for the SPH-based solution of fluid-solid interaction problems. *Computer Methods in Applied Mechanics and Engineering* 345 (2019), 100–122.

[12] Milad Rakhsha, Lijing Yang, and Dan Negrut. [n. d.]. A comparison of three SPH methods for the solution of the Fluid-Solid Interaction problem. *Computer Methods in Applied Mechanics and Engineering - under review* ([n. d.]).

[13] E. Strohmaier, J. Dongarra, H. Simon, and M. Meuer. [n. d.]. TOP500 Supercomputer Site. http://www.top500.org.

[14] ZY Zhou, D Pinson, RP Zou, and AB Yu. 2011. Discrete particle simulation of gas fluidization of ellipsoidal particles. *Chemical Engineering Science* 66, 23 (2011), 6128–6145.