ACM SIGGRAPH / Eurographics Symposium on Computer Animation (2024)
M. Skouras and H. Wang (Editors)

Adaptive Sampling for Simulating Granular Materials

Samraat Gupta[1] and John Keyser[1]

[1]Texas A&M University

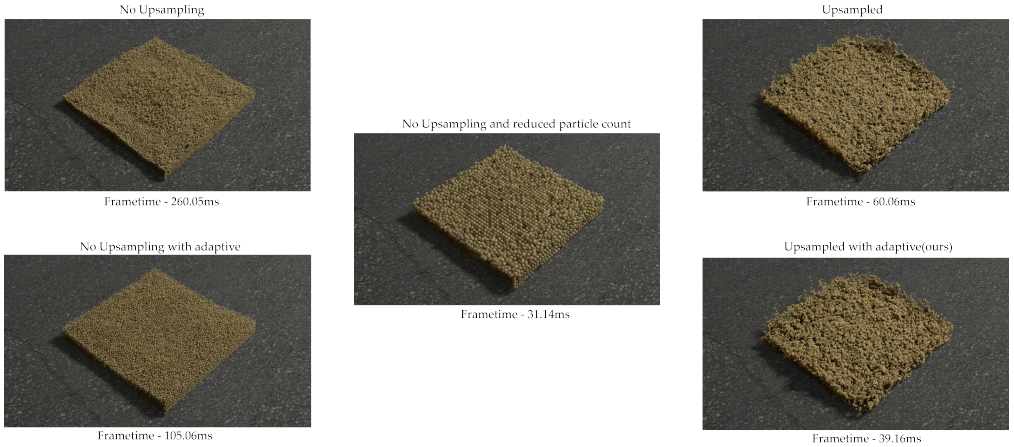

Figure 1: *Comparison of different approaches. Our method achieves a faster frame rate by reducing the number of particles simulated, while still maintaining the quality of an upsampled approach.*

Abstract

We present a method for generating simulations of granular materials more quickly within a position based dynamics framework. We do this by combining an adaptive particle sampling scheme with an upsampling approach. This allows for faster simulations in interactive applications, while maintaining visual resolution. Particles are merged or split based on their distance from the boundary, allowing for high details in areas of importance such as the surface and edges. Merging particles into a single particle reduces the number of particles for which collisions have to be simulated, thus reducing the overall simulation time. The adaptive sampling technique is then combined with an upsampling scheme that gives the coarser particle simulation the appearance of much finer resolution.

CCS Concepts

• ***Computing methodologies*** → ***Physical simulation;***

1. Introduction

Granular materials such as sand, gravel, snow, grains, etc. are important to creating realistic virtual environments. Simulating such materials at interactive rates is challenging due to the interplay of forces between the particles used to represent them. Although granular materials display many properties similar to those of fluids, they differ from fluids in that they lose kinetic energy quickly. We propose a technique to improve performance of a PBD-based (or other particle-based) granular simulation. Adaptive sampling [APKG07] [HHK08] enables simulation with adaptive numbers of particles, trading resolution in less visible areas for increased performance. Thus, we introduce a modified adaptive merging and splitting scheme for granular materials that increases performance, especially in scenarios where particles pile up, without sacrificing details in important areas.

2. Method

Our proposed framework is based on Sommer et al. [SSS22] which combines low-resolution or coarse-scale simulation with a high-resolution or fine-scale simulation in order to get high particle

DOI: 10.2312/sca.20241168

Table 1: *Frame times for the demos with comparisons.*

		Average Number		Mean Time Per Frame (ms)			
Demo		LR Particles	HR Particles	LR Updates	HR Updates	Neighbor Search	Total
Piling	Fixed	6,460	129,200	13.29	10.21	24.18	47.68
	Adaptive	3,902	129,200	6.46	8.59	21.66	36.71
Compact Box	Fixed	5,760	115,200	14.59	9.08	25.52	49.20
	Adaptive	2,881	115,200	7.068	7.19	22.49	36.75
Excavator	Fixed	5,760	115,200	16.05	11.39	25.89	53.33
	Adaptive	3,544	115,200	7.46	9.32	24.53	41.30

Figure 2: *Three demo scenarios including particle settling into a stable pile, particles settling in a compact box and particles colliding with a moving excavator bucket. Timings for each are in Table 1.*

counts without having to generate collisions for all particles. The low-resolution simulation is based on the PBD algorithm since it is a particle based approach focused on the needs of interactive applications. Around each low-resolution particle, smaller high-resolution particles are sampled to give more detail without having to generate additional collisions. The initial sampling is very important in order to avoid artifacts such as aliasing [SSS22]. Thus, high-resolution particles are sampled using randomized volume sampling [SS21]. These high-resolution particles are set in motion by the velocity field of neighboring low-resolution particles as described by by Ihmsen et al. [IWT12].

In order to reduce the number of particles that need to be updated each frame, the particles are adaptively merged or split depending on an importance criterion. Particles on the boundary are deemed to be important and are split whereas non-boundary particles are merged together to form larger particles.

A particle is determined to be a boundary particle if it is either a) an isolated particle, b) near the surface of a volume of the particles or c) has spread out over a surface.

We define merging and splitting by following energy conservation laws to keep the total energy constant and the system stable. The equations for merging are:

$$\mathbf{m}_f = \mathbf{m}_1 + \mathbf{m}_2, \mathbf{v}_f = \frac{\mathbf{m}_1\mathbf{v}_1 + \mathbf{m}_2\mathbf{v}_2}{\mathbf{m}_1 + \mathbf{m}_2}, \mathbf{r}_f = \sqrt[3]{\mathbf{r}_1^3 + \mathbf{r}_2^3} \quad (1)$$

and for splitting are:

$$\mathbf{m}_1 = \mathbf{m}_2 = \frac{\mathbf{m}_s}{2}, \mathbf{r}_1 = \mathbf{r}_2 = \sqrt[3]{\frac{\mathbf{r}_s^3}{2}} = \frac{1}{\sqrt[3]{2}}\mathbf{r}_s \quad (2)$$

The processes of merging and splitting introduce new particles in locations that can overlap with existing particles. In a standard PBD framework, this can lead to an artificial increase in kinetic energy, since these overlaps must be resolved by a constraint solve, which would induce a resulting velocity. In order to reduce the effects of these induced overlaps, we run one iteration of constraint solving following a merge/split. For the new particles, we subtract the components of the calculated velocity in the direction that the overlap is resolved in from the total velocity. This allows the particles to (mostly) resolve to a non-overlapping configuration without inducing velocity, and thus not adding kinetic energy artificially.

The method described was implemented in C++ and runs entirely on the CPU. The frame times are given on a system with a 12-core 3.7GHz AMD Ryzen 5900x processor and 32GB of RAM. See Table 1 for a summary of timing improvements. Figure 1 compares our method to alternatives and Figure 2 illustrates scenarios.

References

[APKG07] ADAMS B., PAULY M., KEISER R., GUIBAS L. J.: Adaptively sampled particle fluids. *ACM Trans. Graph. 26*, 3 (jul 2007), 48–es. doi:10.1145/1276377.1276437. 1

[HHK08] HONG W., HOUSE D. H., KEYSER J.: Adaptive particles for incompressible fluid simulation. *Vis. Comput. 24*, 7 (jul 2008), 535–543. doi:10.1007/s00371-008-0234-z. 1

[IWT12] IHMSEN M., WAHL A., TESCHNER M.: High-Resolution Simulation of Granular Material with SPH. In *Workshop on Virtual Reality Interaction and Physical Simulation* (2012), Bender J., Kuijper A., Fellner D. W., Guerin E., (Eds.), The Eurographics Association. doi:10.2312/PE/vriphys/vriphys12/053-060. 2

[SS21] SOMMER A., SCHWANECKE U.: Leaven - lightweight surface and volume mesh sampling application for particle-based simulations. In *Computer Science Research Notes* (2021), Západočeská univerzita v Plzni. doi:10.24132/csrn.2021.3101.17. 2

[SSS22] SOMMER A., SCHWANECKE U., SCHOEMER E.: Interactive high-resolution simulation of granular material. *Journal of WSCG 30*, 1-2 (2022), 9–15. doi:10.24132/jwscg.2022.2. 1, 2