High-Performance GPU Implementation of the Material Point Method for Equal-Mass Asteroid Impact Simulations

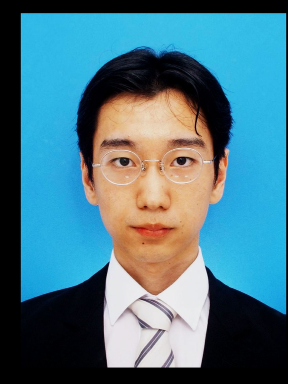

Haru Fukuta (M1) Graduate school of Science and Engineering, Chiba University, 25wm0212@student.gs.chiba-u.jp

Tomoaki Ishiyama Digital Transformation Enhancement Council, Chiba University, ishiyama@chiba-u.jp

CHIBA UNIVERSITY

1. Introduction

- The **Material Point Method (MPM)** integrates continuum mechanics into the Particle-In-Cell (PIC) method, making it ideal for simulating large deformations and fractures.
- Smoothed Particle Hydrodynamics is commonly used for equal-mass asteroid impacts (Sugiura et al. 2018 [2]), MPM utilizes a hybrid particle-grid approach, offering distinct advantages for contact mechanics.

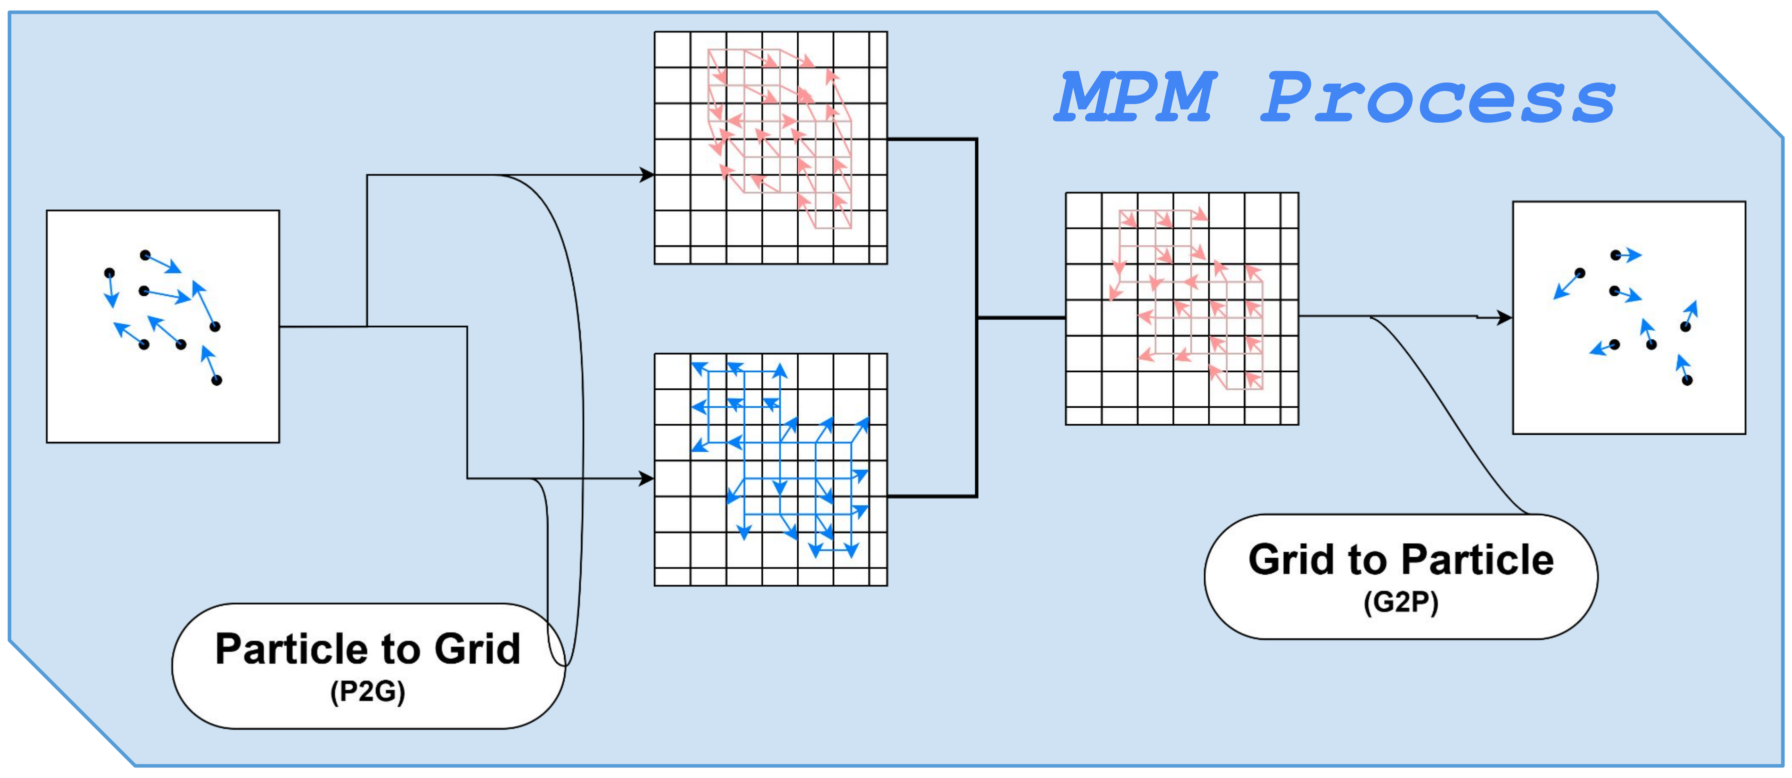

HPC Challenge: Post-impact fragments scatter over a vast spatial domain. To address memory inefficiency, we adopted the **GPU-based Sparse Grid method** (Gao et al. 2018 [3]) and optimized the implementation for **a single NVIDIA A100 GPU** at the Center for Computational Astrophysics, National Astronomical Observatory of Japan (NAOJ).

2. Method

Model Description We implemented the fracture model described by Benz & Asphaug (1995) [1] and the rock material model from Jutzi (2015) [2]. The Tillotson Equation of State (EOS) is adopted to calculate pressure.

Sparse Grid Construction for MPM on GPU

Approach: Implemented a hash-based sparse grid method for GPU-based MPM (following **Gao et al. 2018 [3]**) to allocate memory only for active grid blocks populated by particles.

Data Structure: Spatial hashing maps active cells to a linear array in GPU global memory.

Optimization: Guided by **Nsys profiling**, core physics utilize custom CUDA kernels, while **CUDA Graphs** efficiently manage dependencies and minimize synchronization overhead during dynamic grid reconstruction.

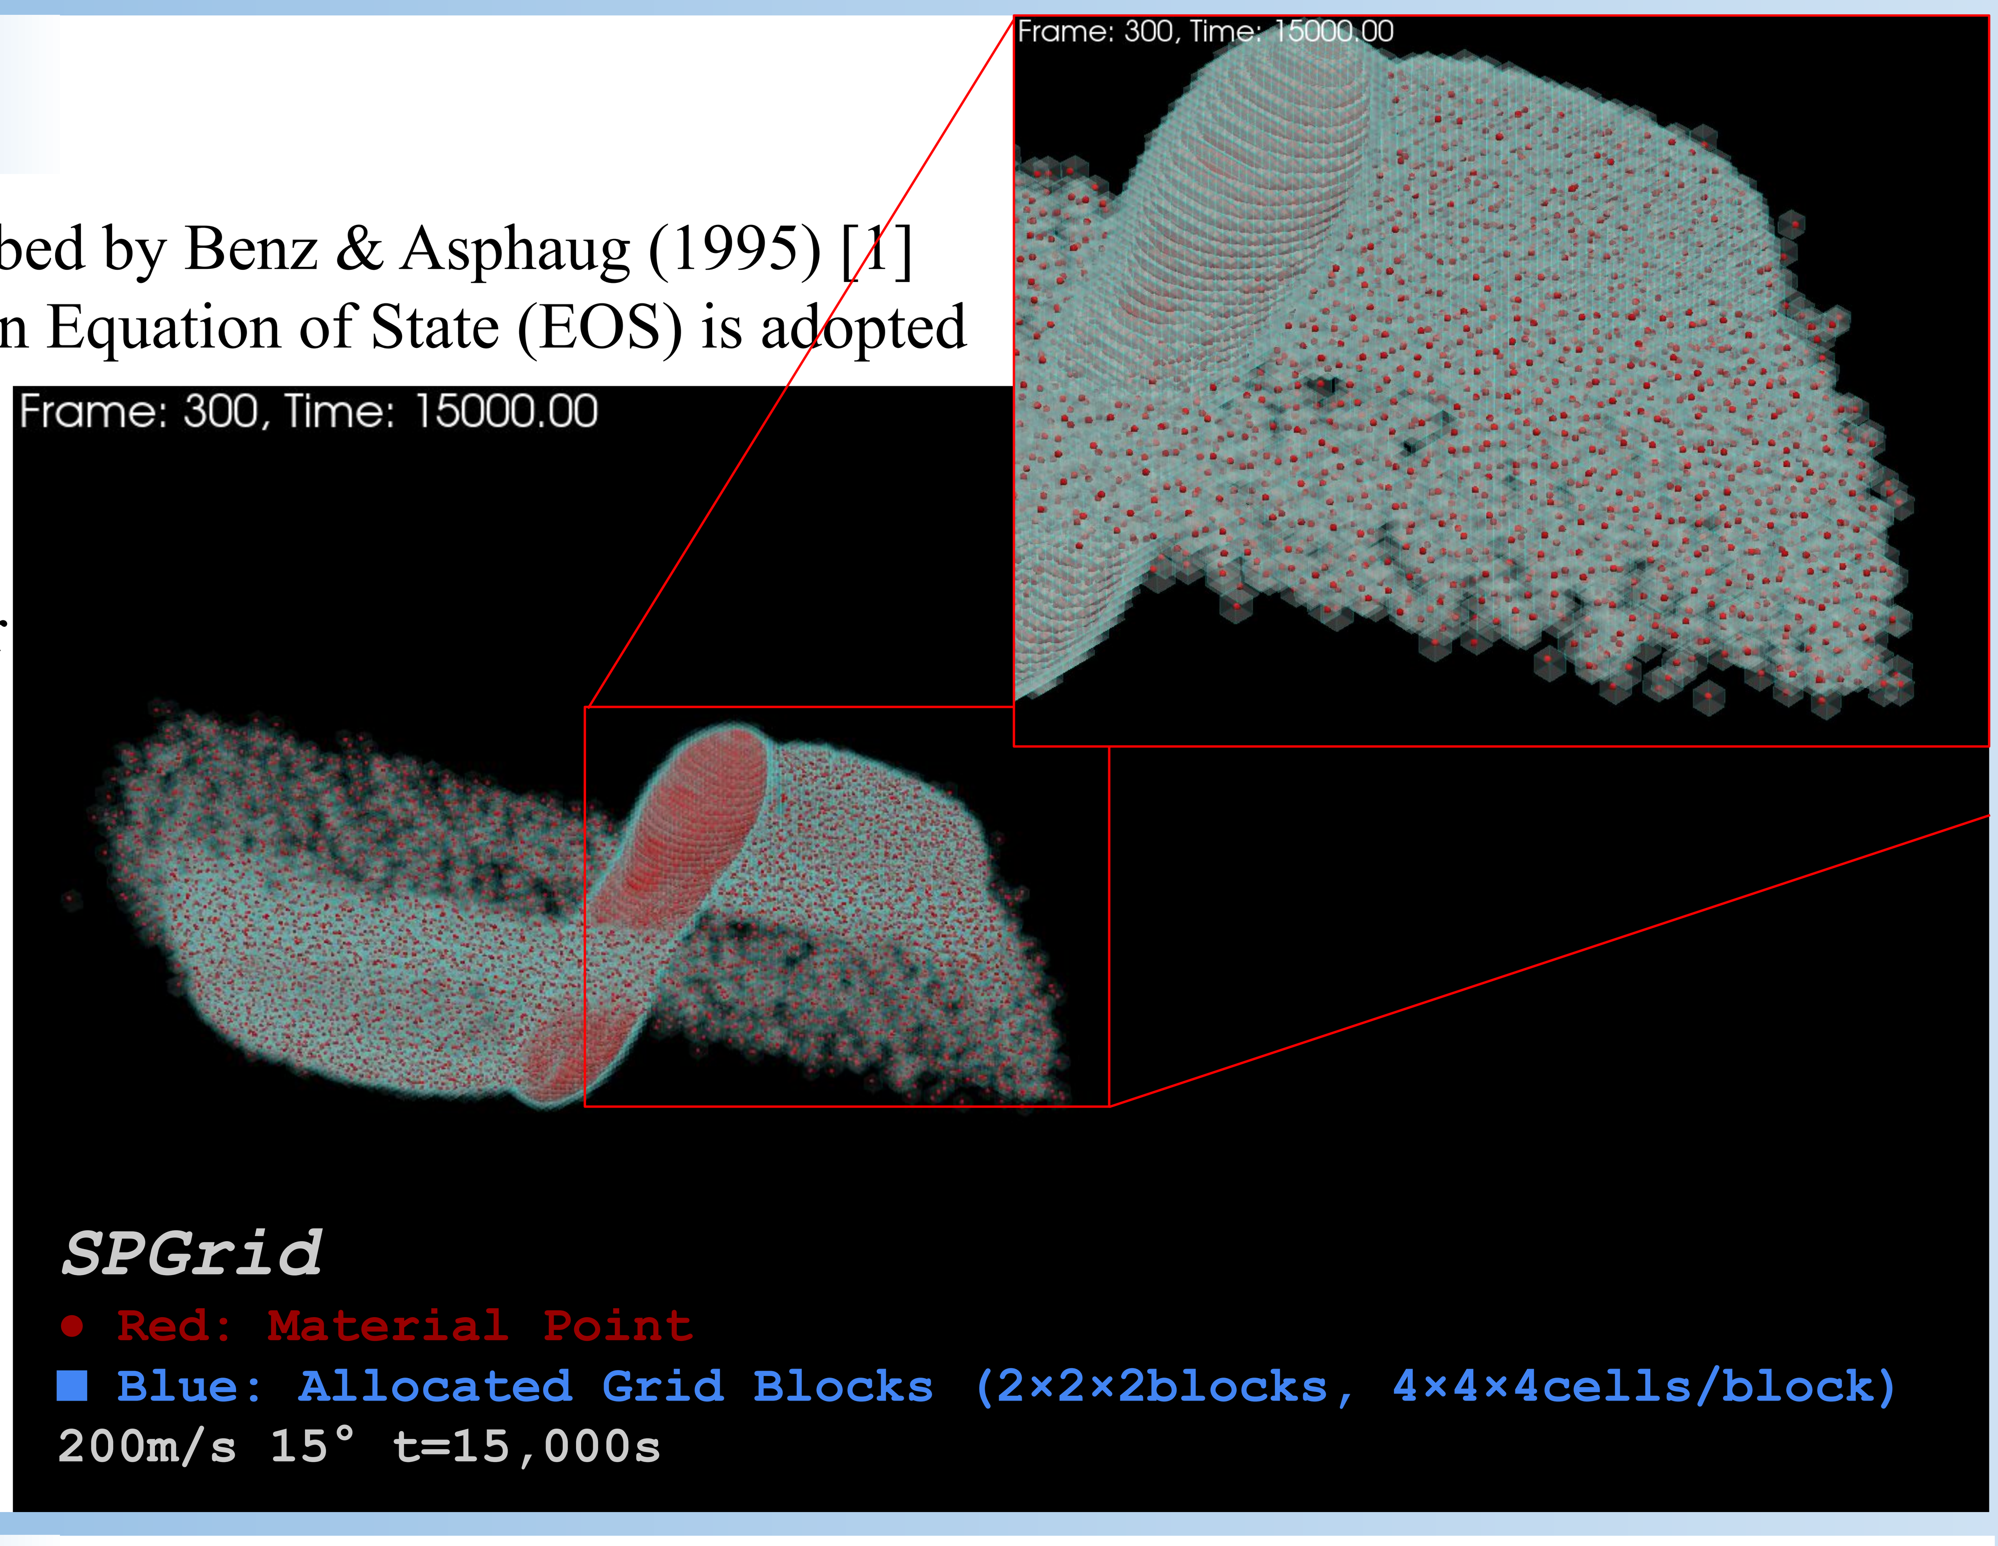

3. Result

- **Experimental Setup & Physics Validation**
  - Execution Environment: Single NVIDIA A100 GPU (40GB memory).
  - Simulation Setup: Initial impact conditions are adopted from Sugiura et al. 2018[2].
  - 5.4×10^5 particles (Grid size: 1.25km, 1 part./cell)
  - Physics: Symplectic Euler for MPM; 2nd-order RK for gravity (updated once per 100 steps).

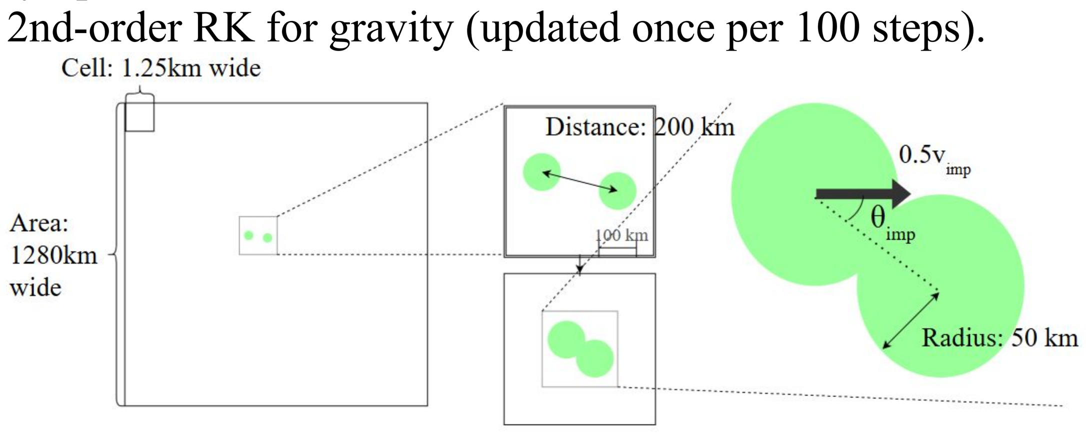

- To validate the code, we analyzed the collision outcomes. The figures below show the mass of the largest remnant and its shape properties.
- **Figure Caption:** (Left) Mass of the largest remnant normalized by an initial asteroid mass. Shape analysis showing axis ratios: intermediate/major (b/a, Middle) and minor/major (c/a, Right).

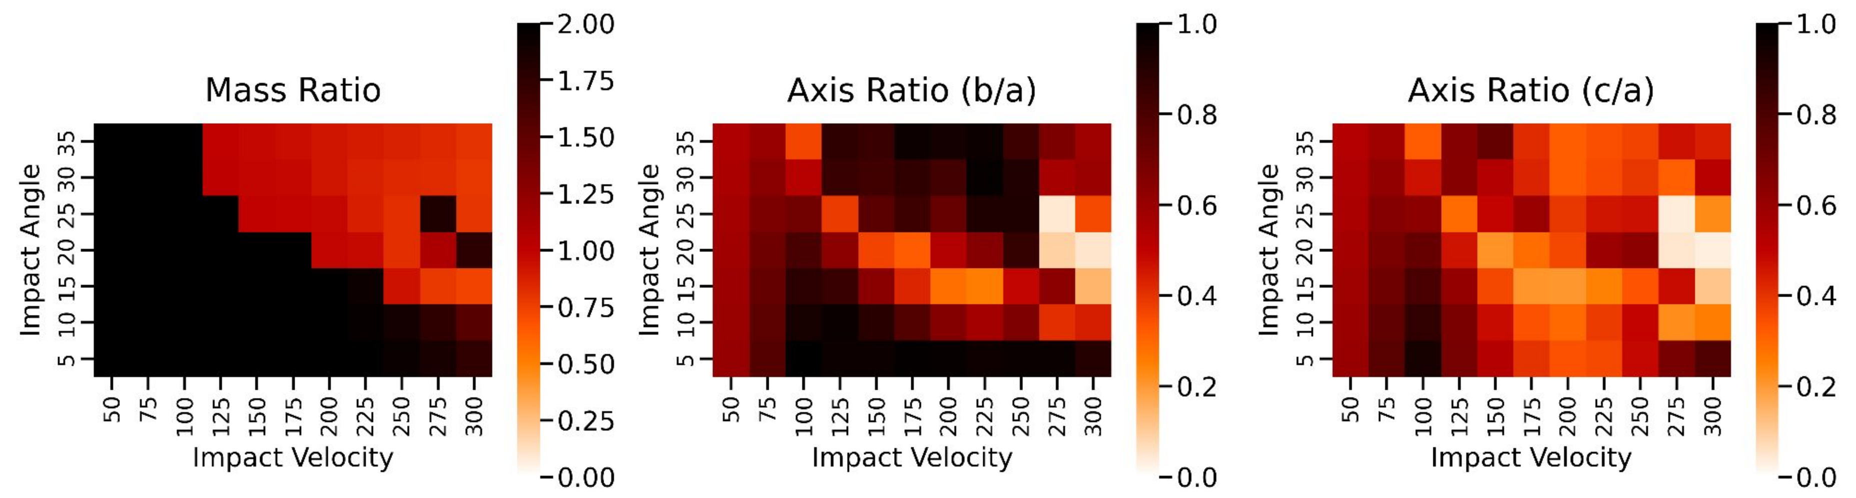

- **Performance Profiling**
  - **Profiling Tool:** Validated using **NVIDIA Nsight Systems (Nsys).**

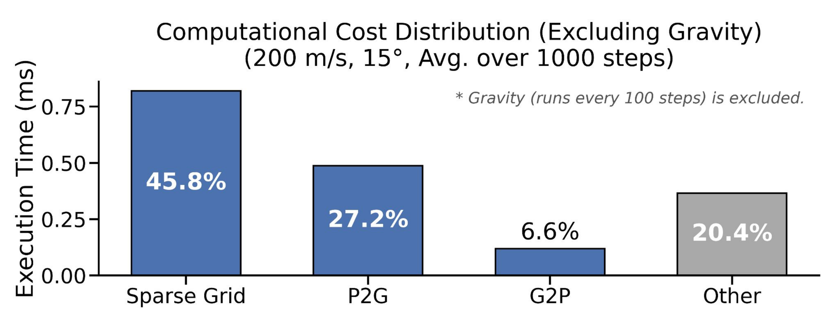

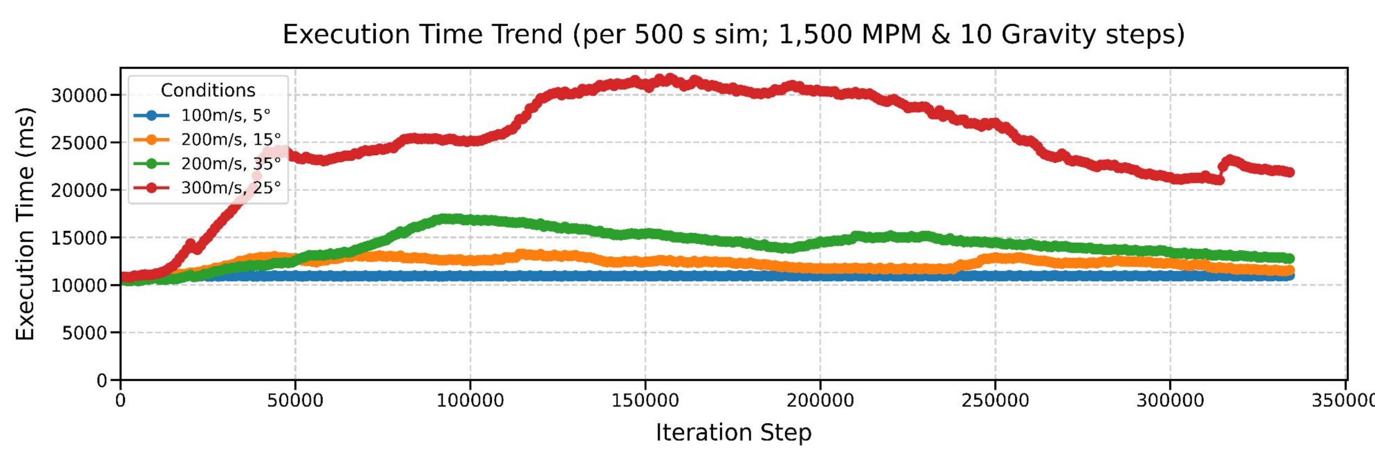

Kernel Tuning Results: We analyzed the computational cost of the dynamic grid construction relative to the physics kernels (P2G/G2P). Through intensive tuning, Sparse Grid construction time was reduced to a level **comparable** to the P2G transfer. On the A100, where P2G is highly accelerated by hardware atomics, this result ensures a **well-balanced performance profile**. Establishing this GPU-centric workflow serves as a critical prototype for the "**Fugaku Next**" era, which is expected to integrate massive GPU acceleration.

4.Future work

Next-Gen Optimization: Investigating kernel fusion of P2G and G2P phases utilizing **Thread Block Clusters** and **Distributed Shared Memory** on NVIDIA Hopper/Blackwell architectures.

Reference
[1] Benz, W. & Asphaug, E. (1995), Comput. Phys. Commun., 87, 253–265.
[2] Sugiura, K., Kobayashi, H. & Inutsuka, S. (2018), A&A, 620, A167.
[3] Gao, M. et al. (2018), ACM Trans. Graph., 37(6).
[4] Jutzi, M. (2015), Planet. Space Sci., 107, 3–9.
[5] Klár, G. et al. (2016), ACM Trans. Graph., 35(4).
[6] Hu, Y. et al. (2018), ACM Trans. Graph., 37(4).