

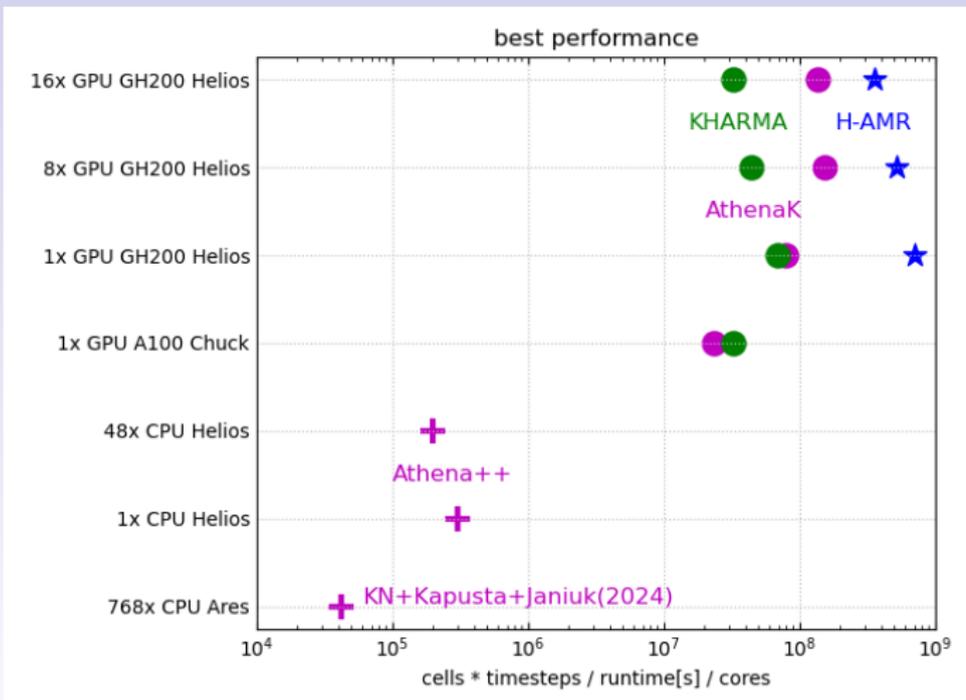
Magic numbers in Kharma and H-AMR GPU scaling

Jakub Szpila

Supervisor: Krzysztof Nalewajko

CAMK Annual Meeting, 2026

CPU vs GPU performance



CPU: 10^4 – 10^6 /core vs GPU: 10^7 – 10^9 /device

Configuration pitfalls

- MPI setup
- Meshblocks: number and size
- Memory - both CPU and GPU

H-AMR, Torus problem size: 732x256x256, 4 LAT levels
Helios GH200 NVIDIA GPU

GPU count	1	4	8	16
number of meshblocks	8	8	16	16
effective zone-cycles/s	6.2×10^8	3.7×10^8	2.9×10^8	1.9×10^8

Performance drops/gains:

- number of GPUs: 25-80% efficiency (depends on LAT), some configuration (e.g. 12 GPU) significantly slower
- number of blocks: by up to a factor of 1.5 (depends on LAT)
- which dimension (N_r , N_θ , N_ϕ) is subdivided for meshblocks: by up to a factor of 2