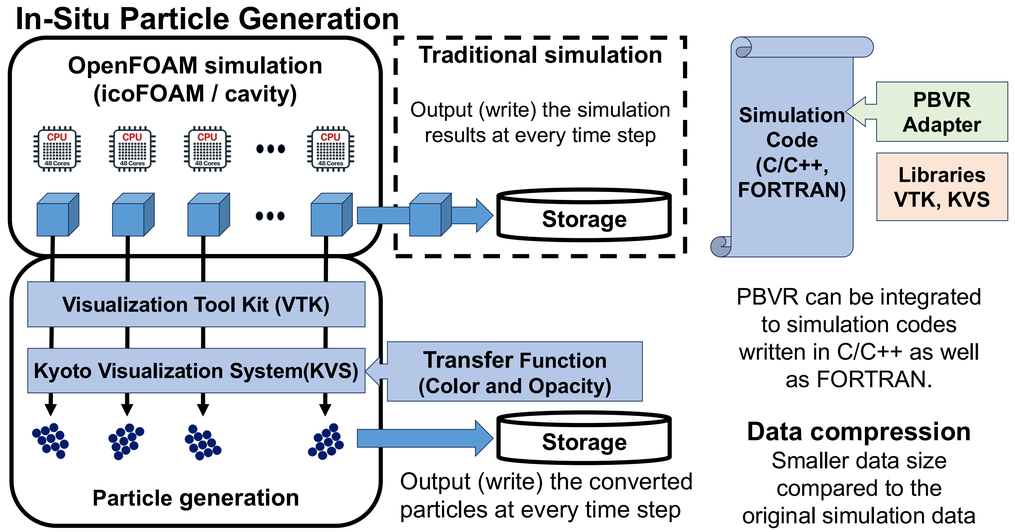
Revisiting Distributed PBVR on Supercomputer Fugaku and Beyond

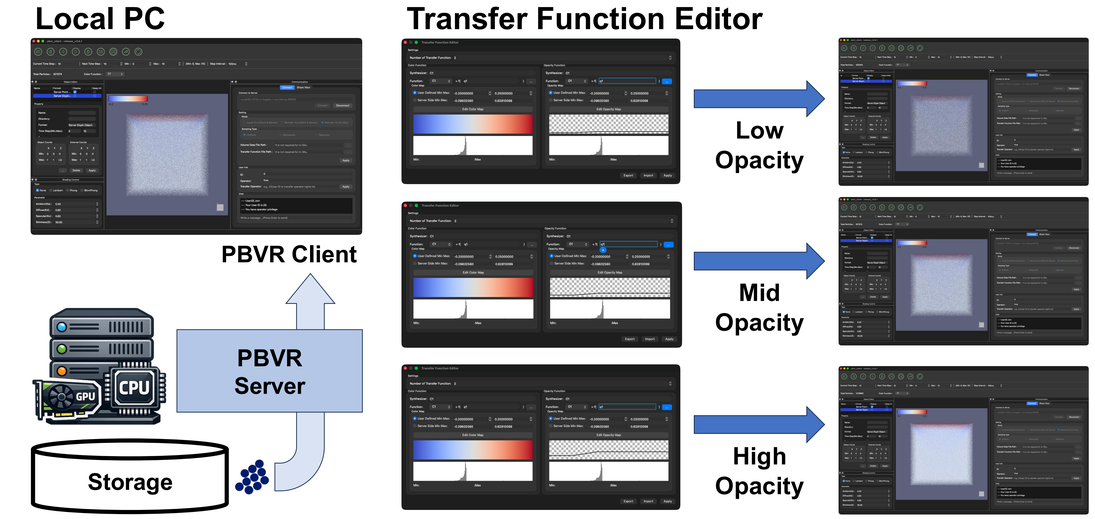
Tsukasa Osaki, Jorji Nonaka
RIKEN Center for Computational Science, Kobe, Japan

Takuma Kawamura
Japan Atomic Energy Agency (JAEA), Kashiwa, Japan

- **Motivation : Utilize available CPU resources for in-situ visualization-oriented particle generation to reduce simulation data I/O.**
- **Approach : Generate particles in situ on HPC CPUs using Particle-Based Volume Rendering (PBVR), and render them interactively on a GPU-equipped local PC.**
- **Results : Demonstrated high scalability up to 110,592 CPU cores on Fugaku (2,304 MPI processes / 48 OpenMP threads), while greatly reducing visualization data.**
- **Future Work : Evaluate the framework on GPU-centric HPC systems, utilizing available CPU resources for in-situ particle generation.**



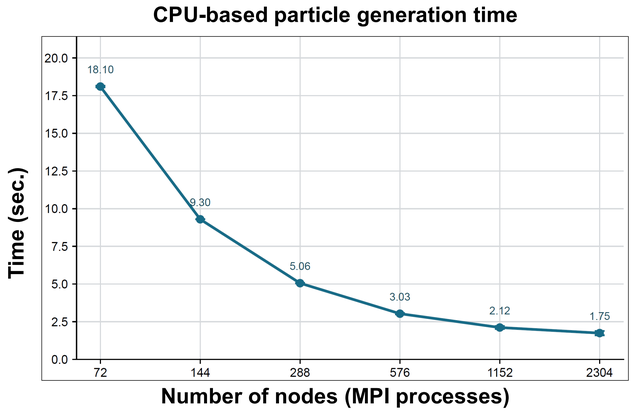




Experimental Evaluation (Fugaku)

We used OpenFOAM CFD simulation to generate particles in-situ on the Fugaku by using following model and parameters.

In-Situ Particle Generation	
OpenFOAM	v2206
Model	icoFoam/cavity
Grid size	432 x 432 x 432
MPI processes	72, 144, 288, 576, 1152, 2304
OpenMP threads	48
Compiler	Fujitsu Compiler 4.12.0



The CPU-based particle generation time includes the required data conversion using the VTK library and particle generation using the KVS library based on a user-defined transfer function. As a reference, the data reduction ratio is evaluated using a smaller simulation model (120 × 120 × 120). Since the opacity values specified by the transfer function directly affect the number of generated particles, data reduction ratios for aforementioned three different opacity settings are presented. Although particle generation introduces additional computational overhead, the resulting data reduction can reduce the overall I/O requirements compared with directly writing the original simulation data.

Data reduction (Model: 120 x 120 x 120)	
OpenFOAM Mesh Size	342 MB
OpenFOAM U・p（11 steps）	874 MB
OpenFOAM Total Data Size	1.22 GB
PBVR Particle files (t_*)	185 MB

Effect of Opacity (Model: 120 x 120 x 120)			
Opacity	# Particles	Particle Data Size	Reduction ratio
Low	293,304	7.9 MB	90.9%
Mid	621,074	16.8 MB	80.7%
High	1,419,960	38.3 MB	55.9%

Conclusions

We evaluated the proposed distributed PBVR framework, in which particle generation is performed on the HPC system using CPUs while interactive volume rendering is carried out on a GPU-equipped local PC. The framework demonstrated high scalability, scaling to 110,592 compute cores (2,304 MPI processes with 48 OpenMP threads each). We also observed substantial data reduction, making the approach well suited for interactive analysis of large-scale simulation data. As Fugaku is expected to remain in operation until around 2030, this distributed PBVR framework can help reduce simulation data I/O while enabling efficient large-scale visualization. Furthermore, the proposed architecture is readily applicable to emerging GPU-based HPC systems, where available CPU resources can be utilized for in-situ processing during GPU-dominated simulations.

Acknowledgements

This work used computational resources of the supercomputer Fugaku provided by the RIKEN R-CCS (Project ID: ra260007). AI tools (Grammarly and ChatGPT) were used to correct and improve the text.

Contact: Tsukasa Osaki <**tsukasa.osaki@riken.jp**>
Data Integration Technology Unit (DIT Unit)
Operations and Computer Technologies Division
RIKEN Center for Computational Science