TSUBAME Grand Challenge program
Adopted projects in the spring 2016 program

TSUBAME Grand Chalenge Summary

This program is only chance to use all nodes of TSUBAME2.5 exclusively, because TSUBAME2.5 is shared by thousands of users. There are two categories:

Category A The large scale application aims high peak-performance. All of TSUBAME2.5 nodes are available.

Category B The large scale application aims scientifically meaningful results. A large portion (1/3) of TSUBAME2.5 is available.

Table Number of Adopted Projects in the TSUBAME Grand-Challenge Program

Category	FY2016		FY2015		FY2014		FY2013		FY2012		FY2011		Total
	Fall	Spr.	Fall	Spr.	Fall	Spr.	Fall	Spr	Fall		Fall	Spr.	
A	1	1	1	2	1	2	0	1	2	2	3	4	20
B	0	1	1	3	2	2	1	1	0	0	2	-	12
Total	1	2	2	5	3	4	1	2	2	2	5	4	33

We started this program since FY2011, and keep on carrying out twice in each year.

Under this program, we have adopted total 30 fruitful projects, some of which were awarded prizes as below.

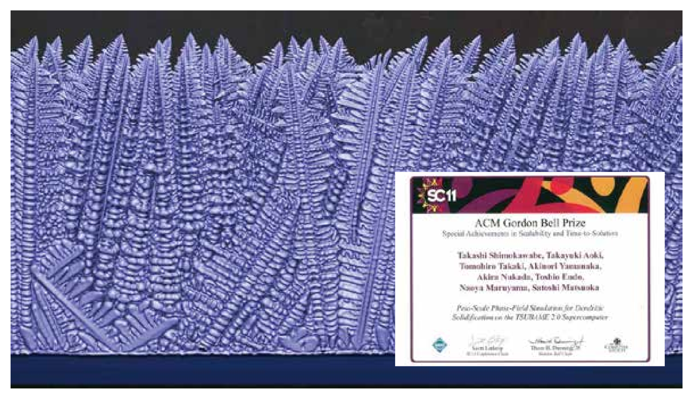

2011 ACM Gordon Bell Prize: Special Achievements in Scalability and Time-to-Solution "Peta-scale Phase-Field Simulation for Dendritic Solidifcation on the TSUBAME 2.0 Supercomputer"

2011 ACM Gordon Bell Prize: Honorable Mention Large scale biofluidics simulations on TSUBAME2

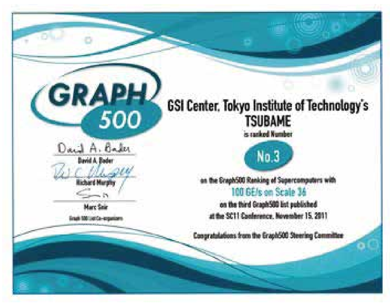

2011 Graph500 Challenge on TSUBAME 2.0

Daino: A High-level Framework for Parallel and Efficient AMR on GPUs

Mohamed Wahib (AICS, RIKEN) Come see our paper at SC16: Wed. 16th 16:30-17:00 at 355-D (Best Paper Finalist)

Motivation

Adaptive Mesh Refinement (AMR) is a model for reducing computation by adapting the resolution of a stencil meshes locally. However, producing efficient AMR code is hard, especially for GPUs. As a result, AMR frameworks require the user to write his own optimized code for the target architecture.

Our approach to achieve target

- A compiler-based framework for producing efficient AMR code (for GPUs)
- Architecture-independent interface provided to the user
- A performance model for quantifying the efficiency

Key results

Our framework generates code comparable in speedup and scalability to hand-written optimized GPU AMR implementations using up to 3,600 GPUs.

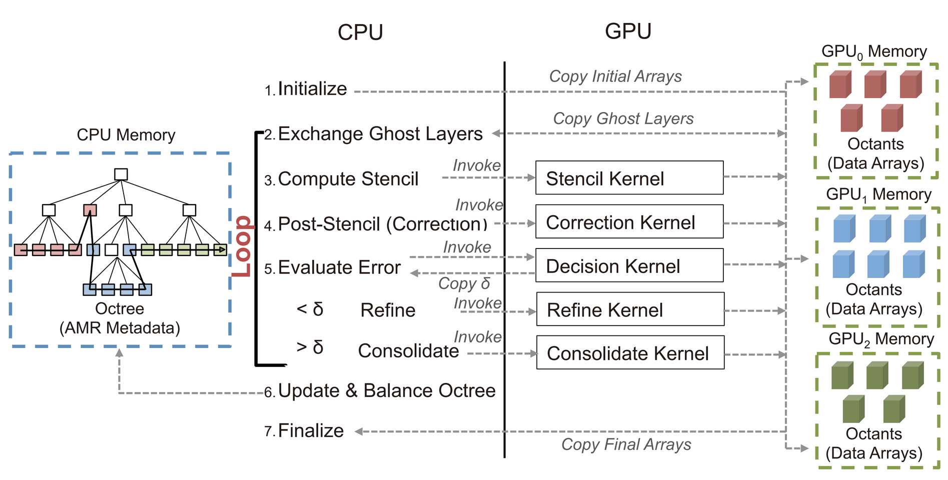

Figure Conceptual overview of the proposed model. All operations touching the data arrays are done by GPU kernels. Accordingly, CPU specializes in operations applied on the octree while GPU specializes in operations applied on the data arrays.

Applications

- Hydrodynamics Solver: We model a hydrodynamics application using Euler equations extending the GAMER implementation.
- Shallow-water Solver: We model shallow water simulations by depth-averaging the Navier–Stokes equations.
- Phase-field Simulation: We evaluate an AMR version of a phase-field simulation for modeling 3D dendritic growth.

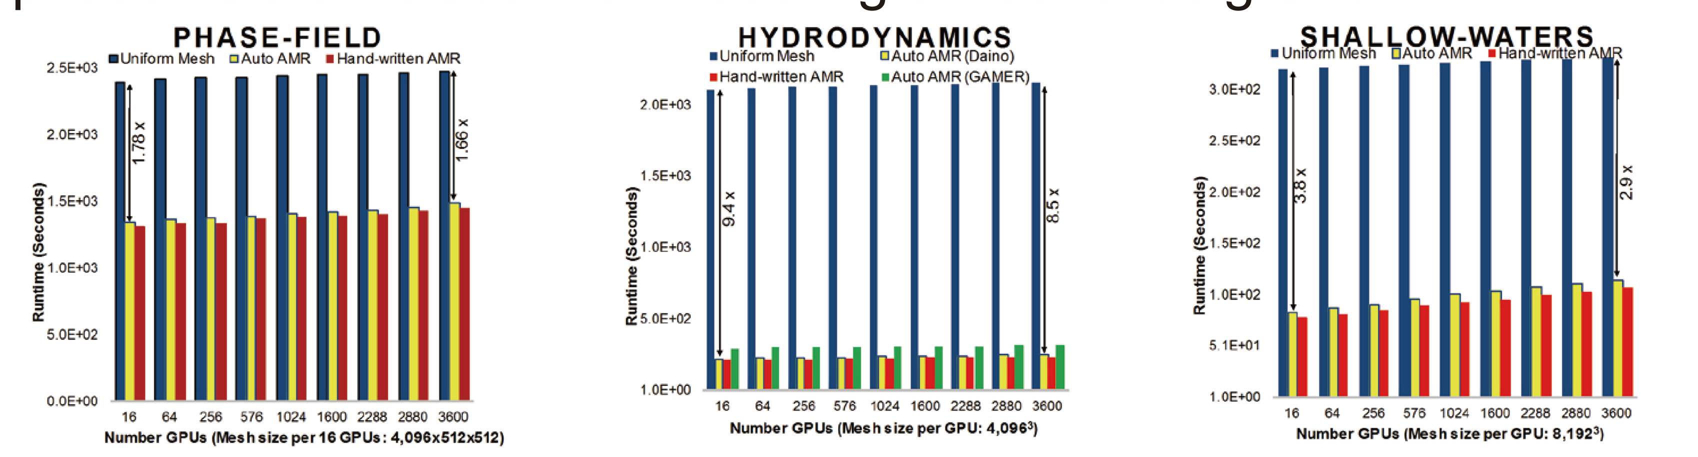

Figure Weak scaling of uniform mesh, hand-written and automated AMR (GAMER-generated AMR included in hydrodynamic)

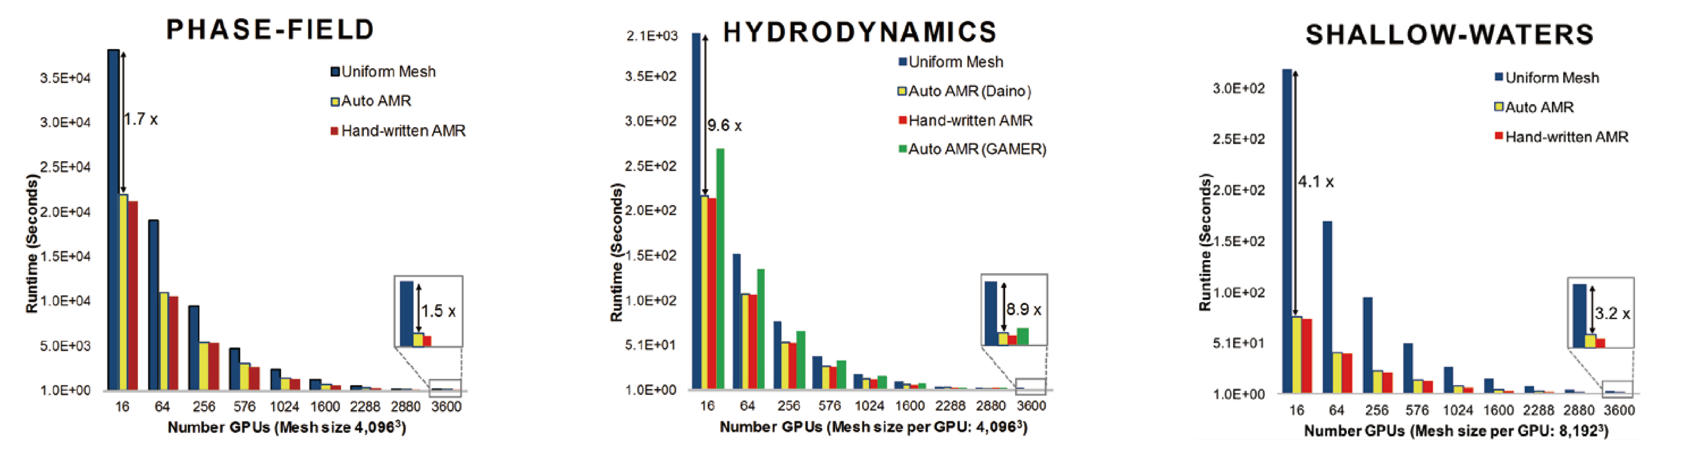

Figure Strong scaling of uniform mesh, hand-written and automated AMR (GAMER-generated AMR included in hydrodynamic)

Particle-In-Cell Simulation of Fusion Plasmas

Rio YOKOTA (GSIC, Tokyo Tech)

Gyrokinetic Particle-In-Cell Simulation

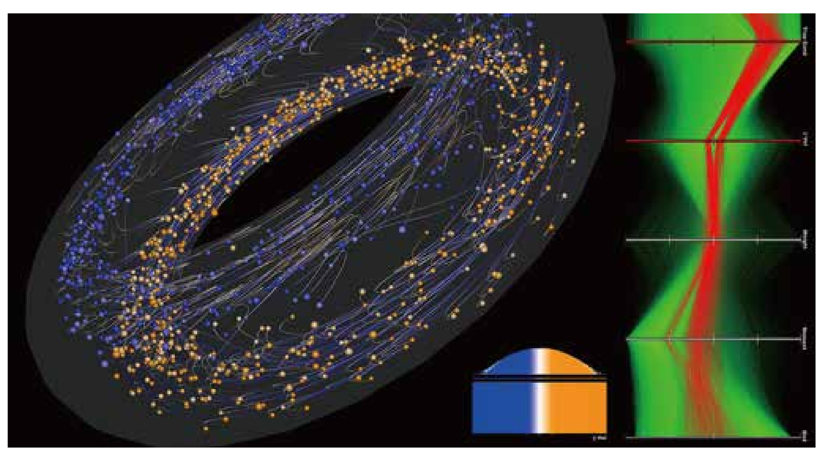

The Gyrokinetic Toroidal Code at Princeton (GTC-P) is a highly scalable and portable particle-in- cell (PIC) code. It solves the 5D Vlasov-Poisson equation featuring efficient utilization of modern parallel computer architectures at the petascale and beyond. Motivated by the goal of developing a modern code capable of dealing with the physics challenge of increasing problem size with sufficient resolution, new thread-level optimizations have been introduced as well as a key additional domain decomposition. GTC-P's multiple levels of parallelism, including inter-node 2D domain decomposition and particle decomposition, as well as intra-node shared memory partition and vectorization have enabled pushing the scalability of the PIC method to extreme computational scales. New discovery science capabilities in the magnetic fusion energy application domain are enabled, including investigations of Ion-Temperature-Gradient (ITG) driven turbulence simulations with unprecedented spatial resolution and long temporal duration.

Comparison with Other GPU Supercomputers

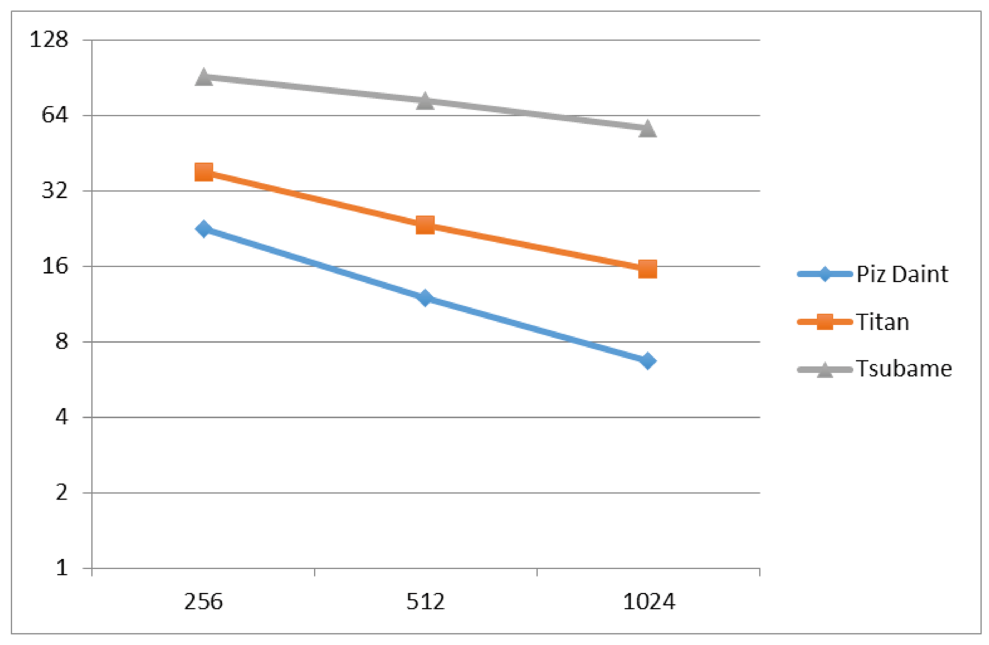

The TSUBAME Grand Challenge Category A was used to perform a full system run. GTC-P has been benchmarked on other leadership class GPU systems such as Piz Daint (CSCS) and Titan (ORNL). These performance comparisons using a realistic discovery- science-capable domain application code provide valuable insights on optimization techniques across one of the broadest sets of current high-end computing platforms worldwide. A strong scalability plot is shown above for 256, 512, and 1024 MPI processes. Piz Daint is the newest machine while Titan is slightly older and TSUBAME2.5 is just about to be upgraded to TSUBAME3.0 so the difference in the absolute performance is somewhat expected. Though, we are currently investigating exactly what is causing the large difference in the performance between these machines. These investigations beneficial to not only the current cross-machine studies, but also to the eventual deployment of GTC-P as a meaningful science-capable code on TSUBAME3.0.

http://www.gsic.titech.ac.jp/sc16