

Accelerating Wave-Propagation Algorithms on Adaptive Mesh with the Graphics Processing Unit (GPU)



Introduction

Clawpack is a library for solving nonlinear hyperbolic partial differential equations using high-resolution finite volume methods based on Riemann solvers and limiters. It supports Adaptive Mesh Refinement (AMR), which is essential in solving multi-scale problems. Recently, we added capabilities to accelerate the code by using the Graphics Process Unit (GPU). Routines that manage CPU and GPU AMR data and facilitate the execution of GPU kernels are added. Customized and CPU thread-safe memory managers are designed to manage GPU and CPU memory pools, which is essential in eliminating the overhead of memory allocation and de-allocation. A global reduction is conducted every time step for dynamically adjusting the time step based on Courant number restrictions. Some small GPU kernels are merged into bigger kernels, which greatly reduces kernel launching overhead. A speed-up between 2 and 3 for the total running time is observed in an acoustics benchmark problem.

Wave Propagation Algorithm

Hyperbolic PDEs in form:

$$q_t + f(q)_x + g(q)_y = 0$$

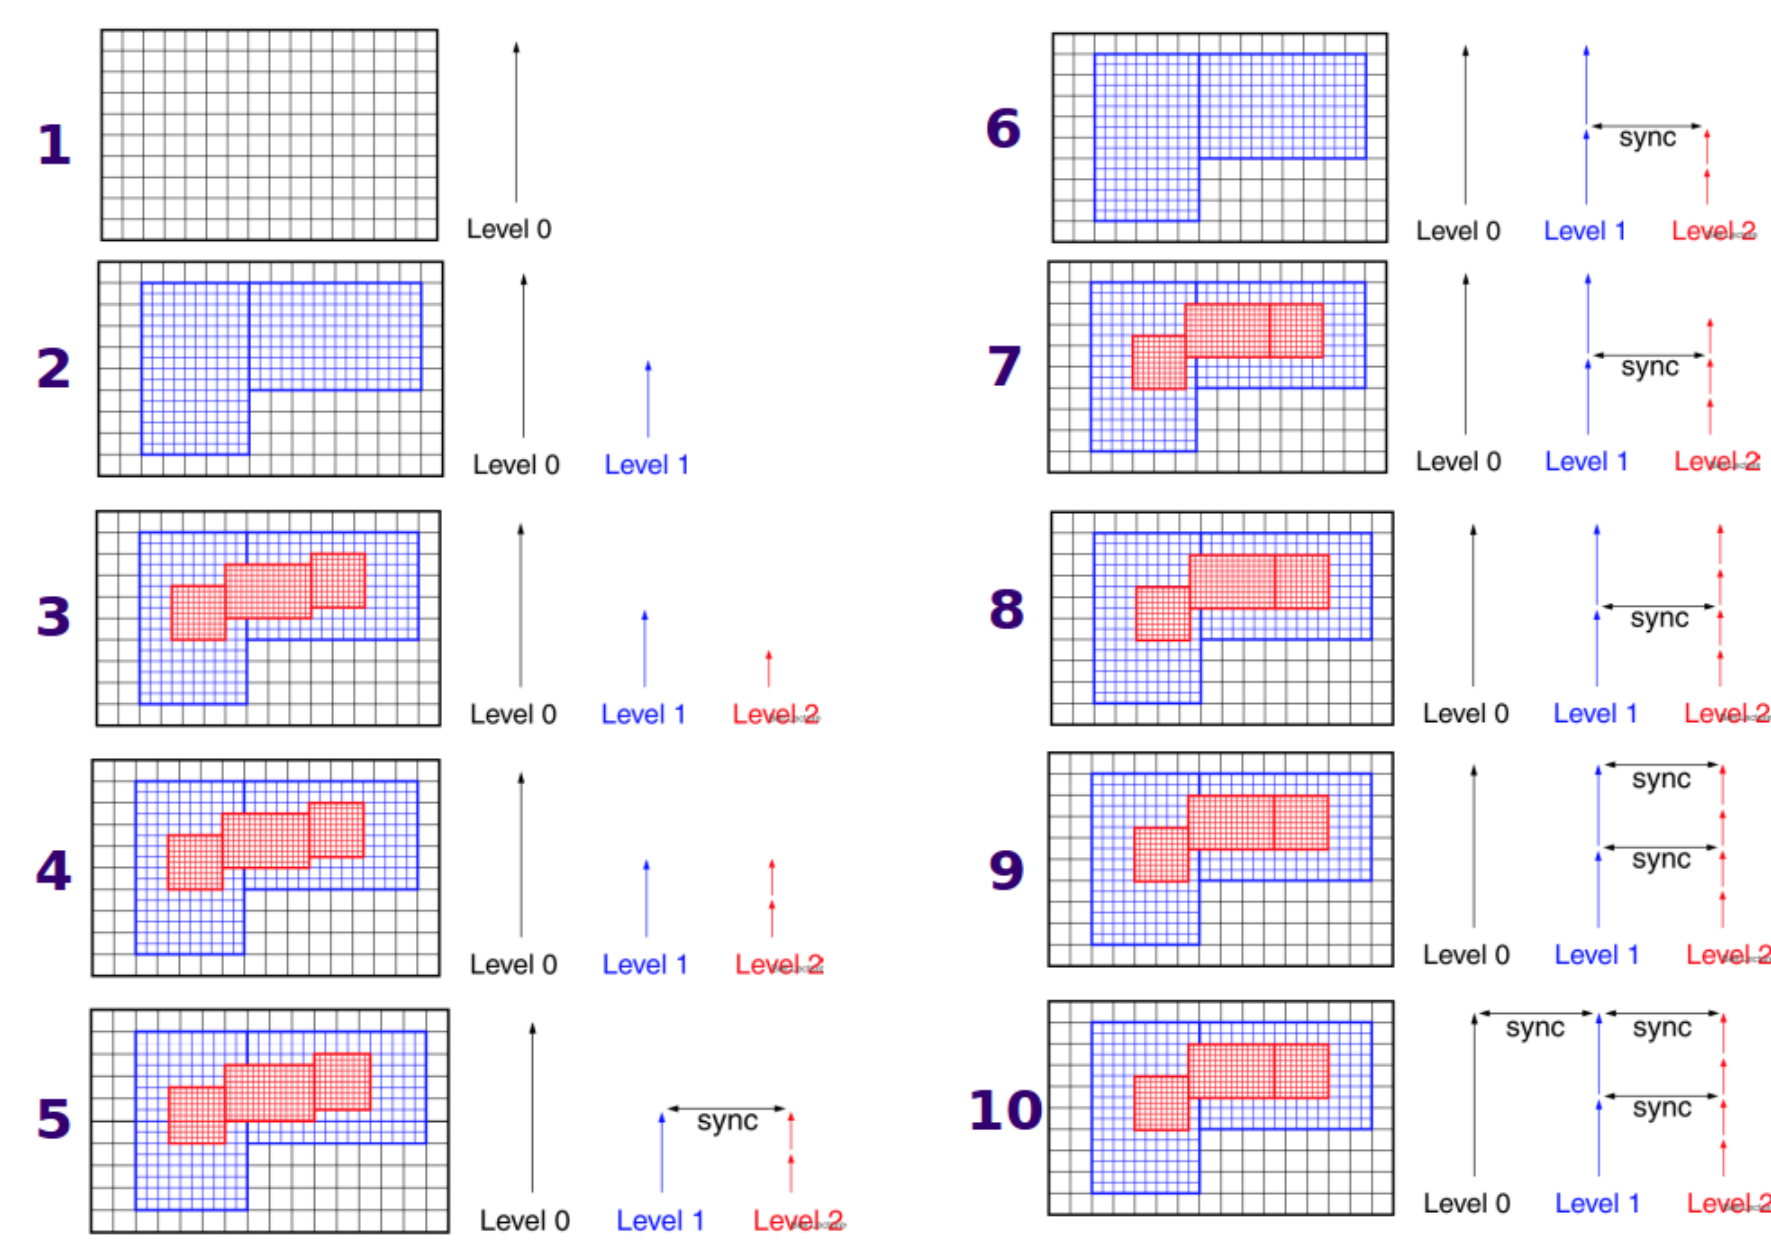
Cell approximation:

$$Q_{ij} = \frac{1}{\Delta x \Delta y} \int_{y_{j-1/2}}^{y_{j+1/2}} \int_{x_{i-1/2}}^{x_{i+1/2}} q(x, y, t^n) dx dy$$

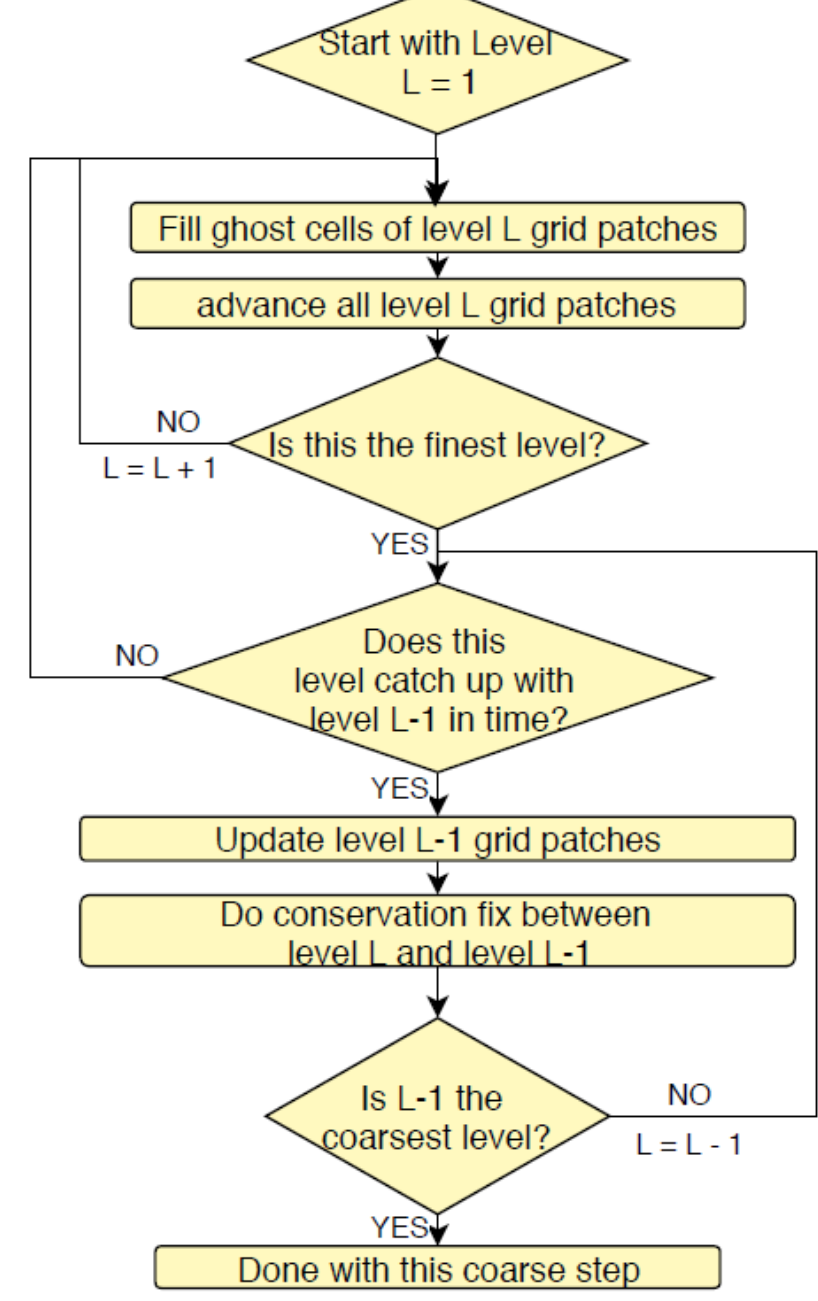
Wave propagation algorithm:

$$Q_{ij}^{n+1} = Q_{ij}^n - \frac{\Delta t}{\Delta x} (A^+ \Delta Q_{i-1/2,j} + A^- \Delta Q_{i+1/2,j}) - \frac{\Delta t}{\Delta y} (B^+ \Delta Q_{i,j-1/2} + B^- \Delta Q_{i,j+1/2}) - \frac{\Delta t}{\Delta x} (\tilde{F}_{i+1/2,j}^n - \tilde{F}_{i-1/2,j}^n) - \frac{\Delta t}{\Delta y} (\tilde{G}_{i,j+1/2}^n - \tilde{G}_{i,j-1/2}^n)$$

Adaptive Mesh Refinement (AMR)



Time stepping with AMR grid patches.



A flow chart of AMR algorithm including 1) advancing solution, 2) updating process 3) and conservation fix. Assume at least two levels of grid patches exist.

Updating process:

When the fine grid patches catch up with the coarse grid patches in time then the solution in coarse grid cells covered by any fine grid cell will be over-written by the average solution in fine grid cells.


$$Q_{i-1,j} := \frac{1}{4} (\hat{Q}_{m-1,b} + \hat{Q}_{m,b} + \hat{Q}_{m-1,b+1} + \hat{Q}_{m,b+1})$$

Conservation fix:

After the updating process, solution in coarse cells that are next to the coarse-fine grid patch interface must be modified to preserve global conservation of state variables.

$$Q_{ij}^{n+1} := Q_{ij}^{n+1} + C1 + C2 + C3$$

where C1, C2 and C3 are modification terms and are functions of waves and fluxes at the coarse-fine grid patch interface.




A 2D coarse grid patch and a nested fine grid patch. i and j are index for the coarse grid in x- and y-directions. m and b are index for the fine grid in x- and y-directions

Code Design and Implementations

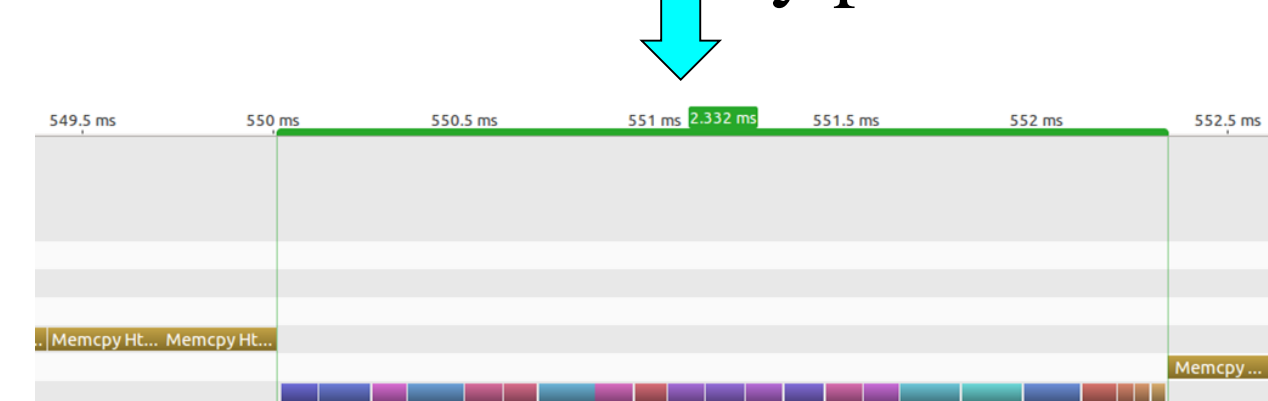
Customized GPU and CPU Memory Allocators

- Allocating system CPU/GPU memories by calling CUDA runtime API directly has a very large overhead (can take up to 80% of the computational time).
- We wrote memory allocator that manages pre-allocated CPU and GPU memory pools.
- This type of overhead is almost completely removed.



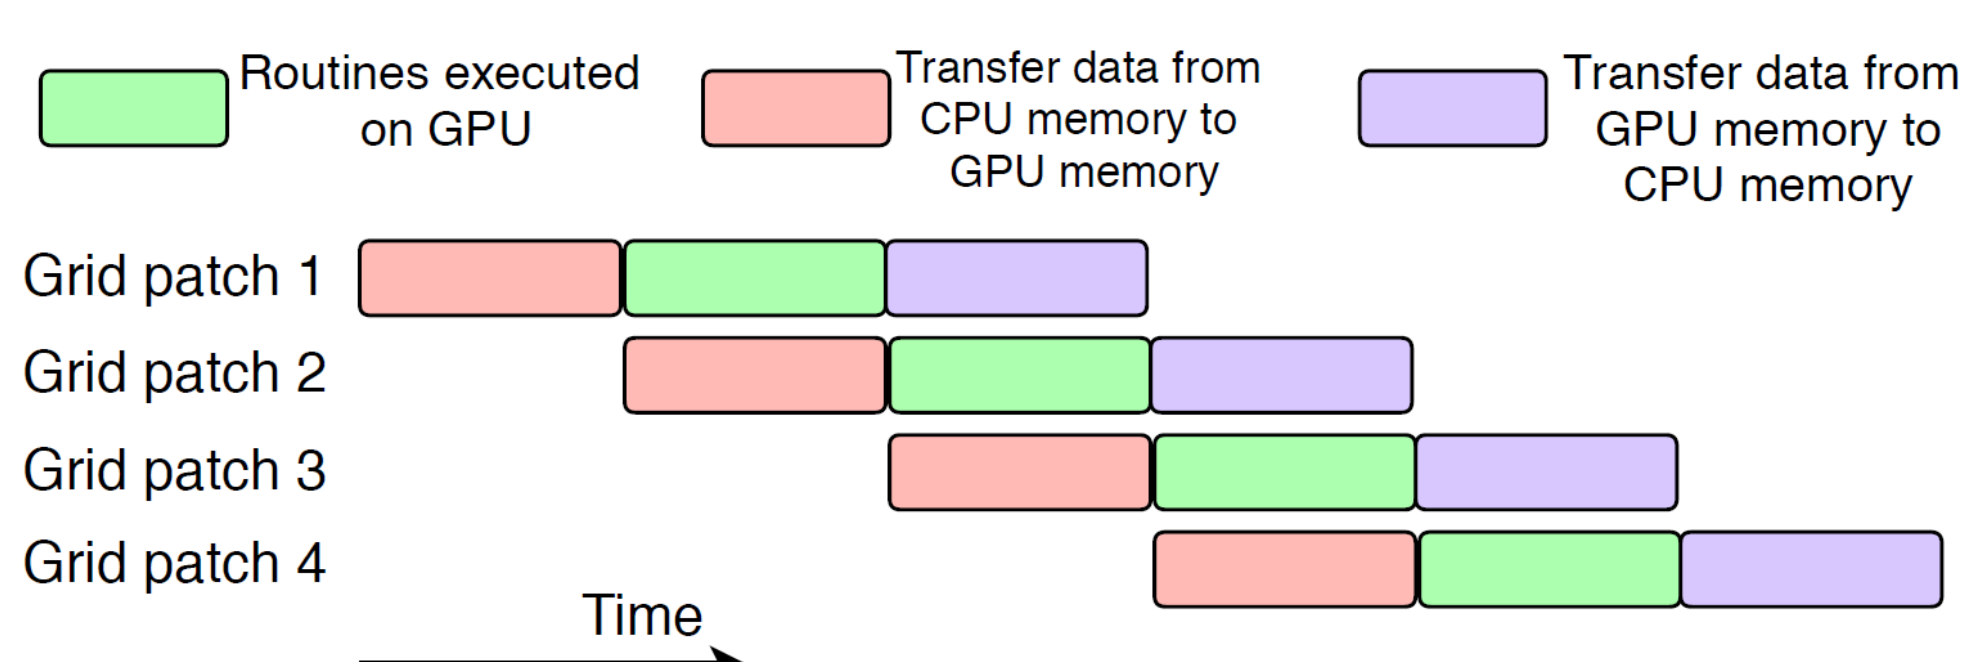
Profiling results BEFORE using memory pool (All gaps between blocks in colors represent time spent on allocating GPU memory. The blocks represent time spent on actual computation on GPUs)

use memory pool



Profiling results AFTER using memory pool (There is almost no gap between blocks)

Pipeline for Hiding Data Transfer



Benchmark and Results

PDEs solved:

A linearized 2D acoustic equation:

$$q_t + Aq_x + Bq_y = 0 \quad \text{where } q = \begin{bmatrix} p \\ u \\ v \end{bmatrix}, \quad A = \begin{bmatrix} 0 & K_0 & 0 \\ \frac{1}{\rho_0} & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix}, \quad B = \begin{bmatrix} 0 & 0 & K_0 \\ 0 & 0 & 0 \\ \frac{1}{\rho_0} & 0 & 0 \end{bmatrix}$$

Machines used:

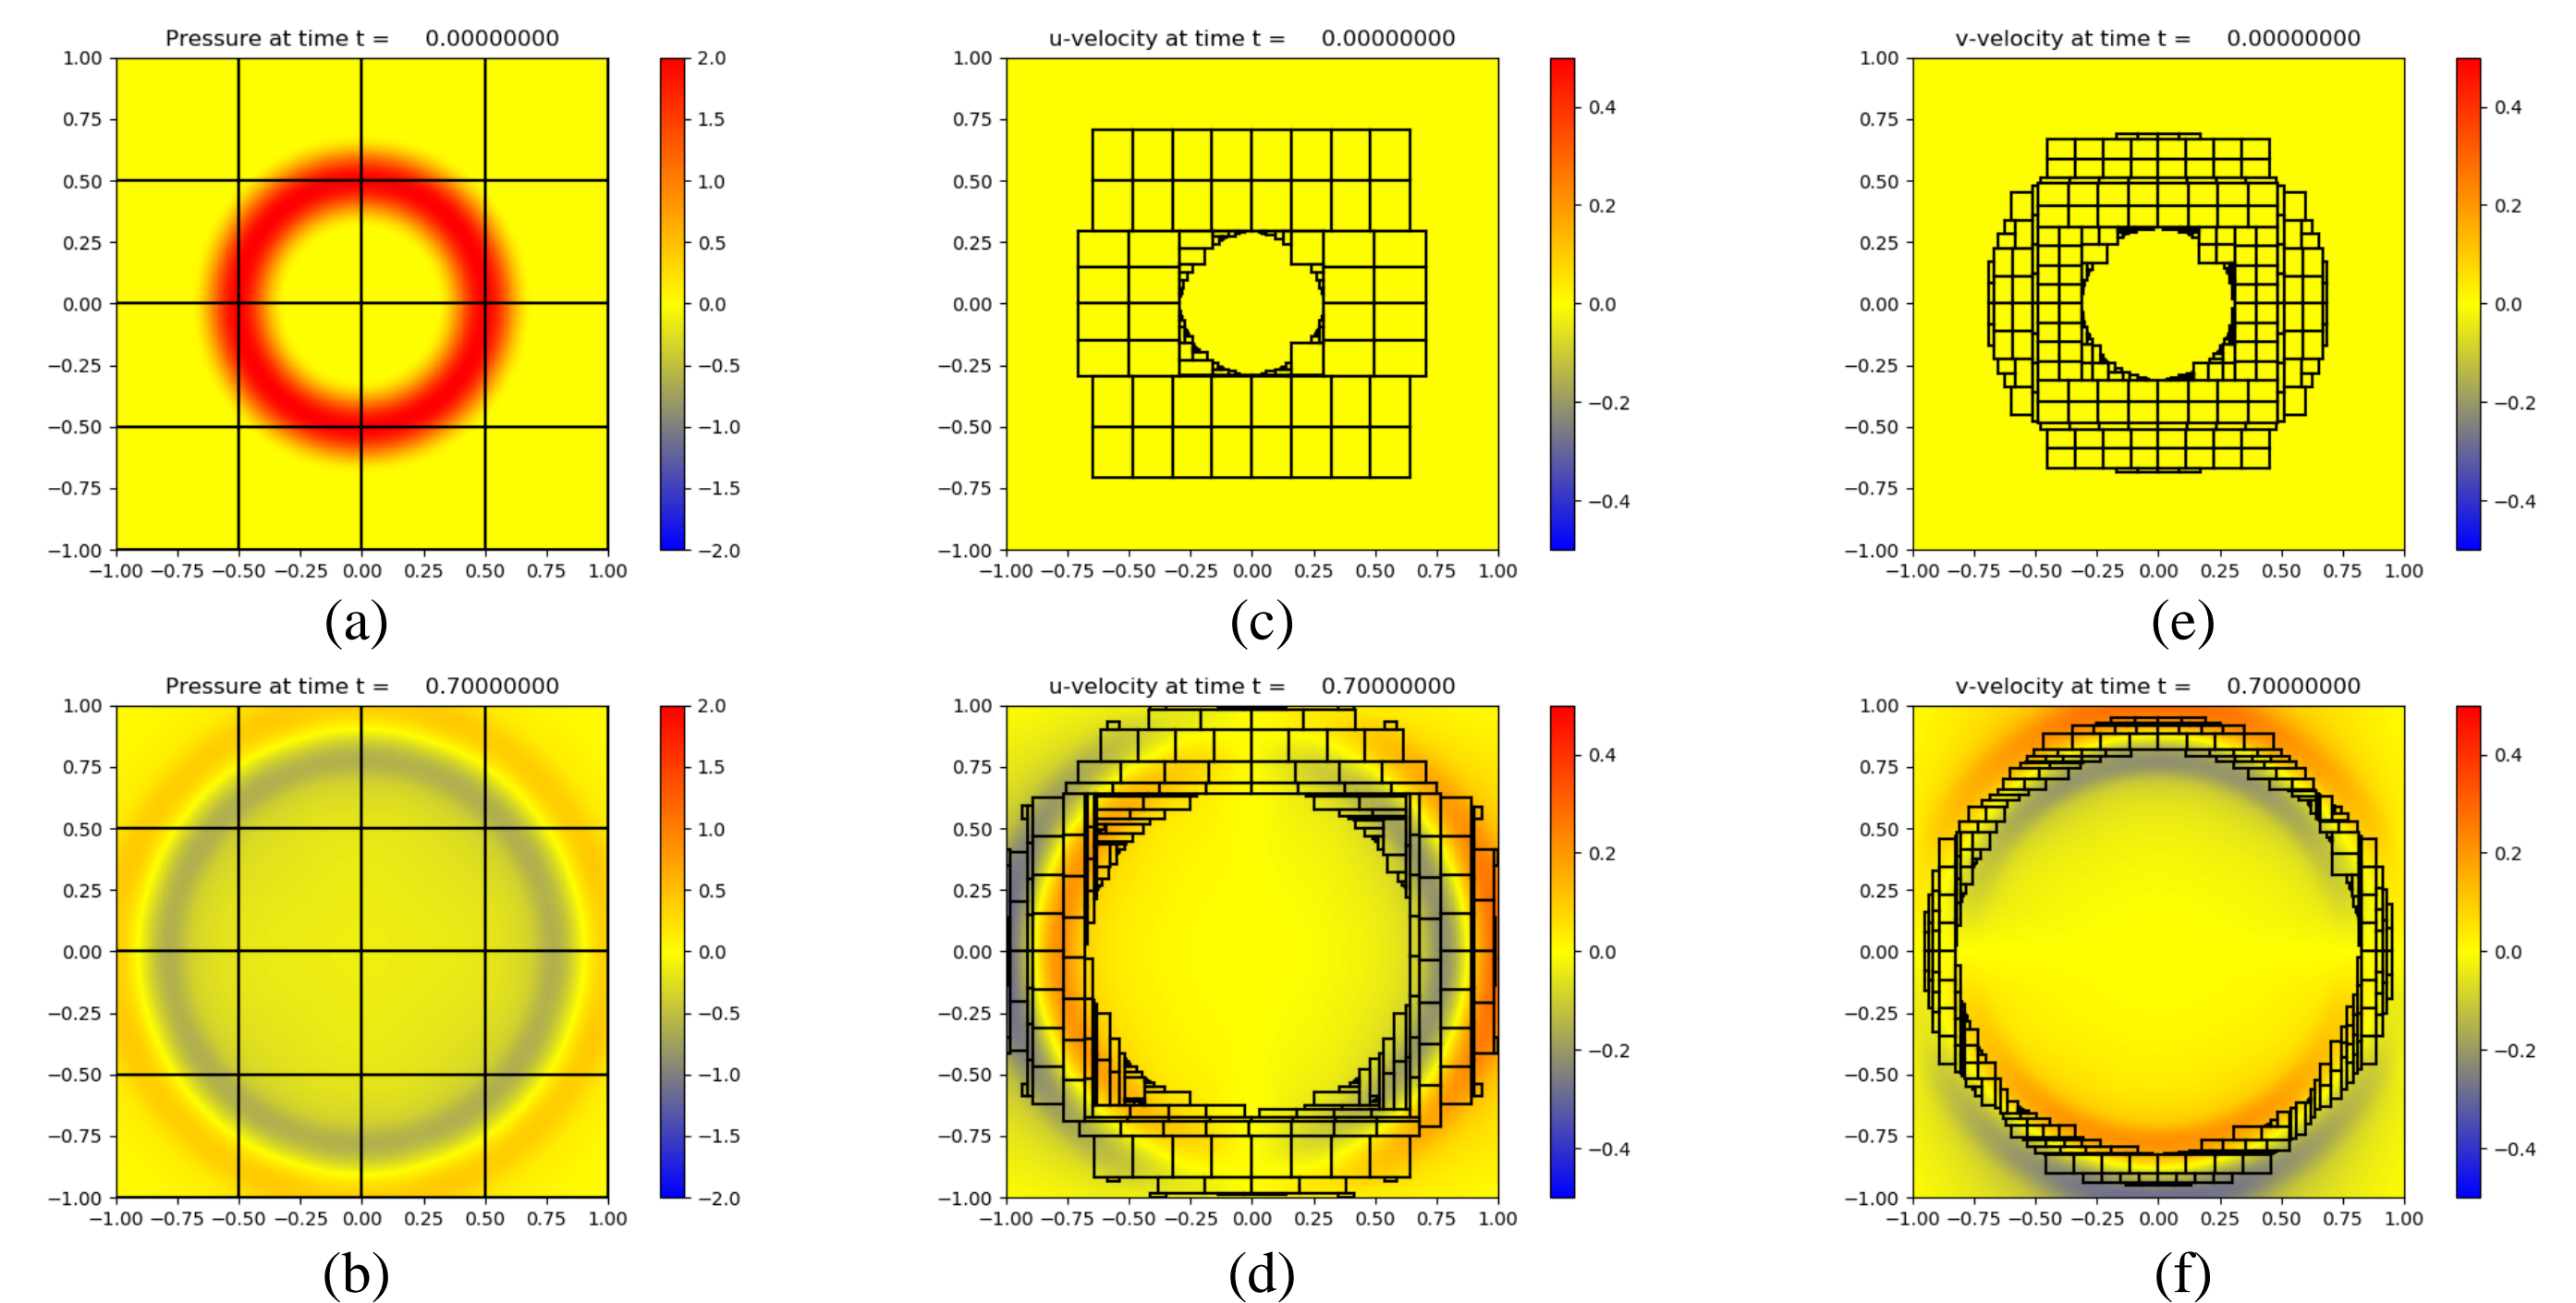
- A single NVIDIA Kepler K20x GPU with a 16-core AMD Opteron 6274 CPU running at 2.2 GHz as the host;
- A single NVIDIA Pascal 100 GPU with a 20-core Intel E5-2698 CPU running at 2.2 GHz as the host (but only 16 CPU threads are used for fair comparison with others);
- A single 16-core AMD Opteron 6274 CPU running at 2.2 GHz;
- A single 16-core Intel Xeon E-2650 CPU running at 2.0 GHz;

Benchmark setup:

A total of 3 levels of AMR grids are used with refinement ratios of 2 between each two levels.

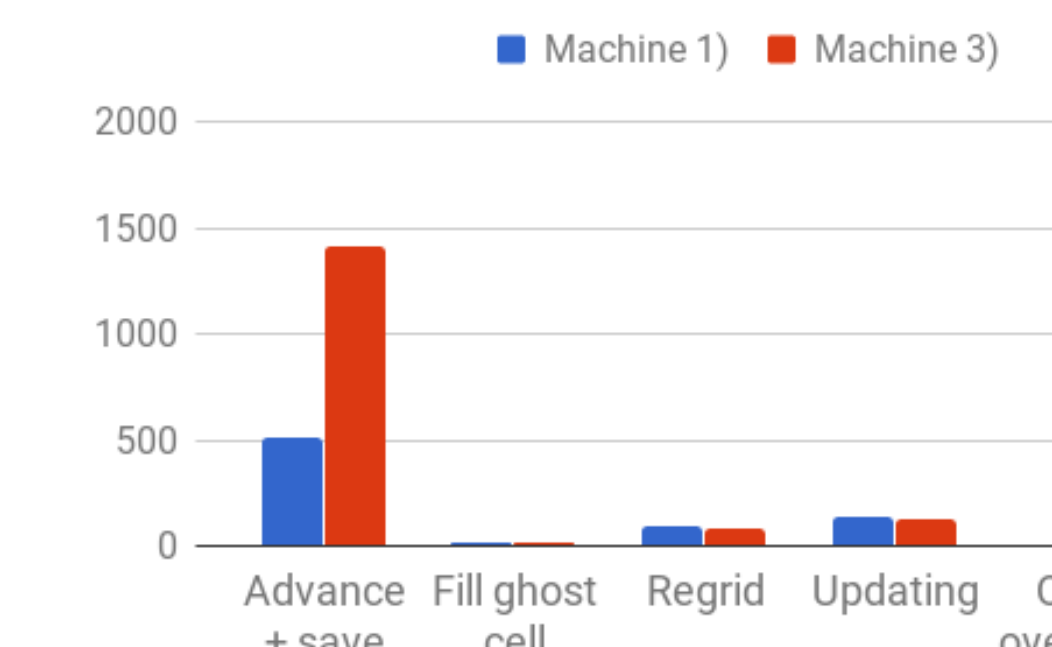
The computational domain is square with 1000 by 1000 cells on the base level, leading to an effective 4000-by-4000-cell resolution on the finest level

Simulation results and performance:

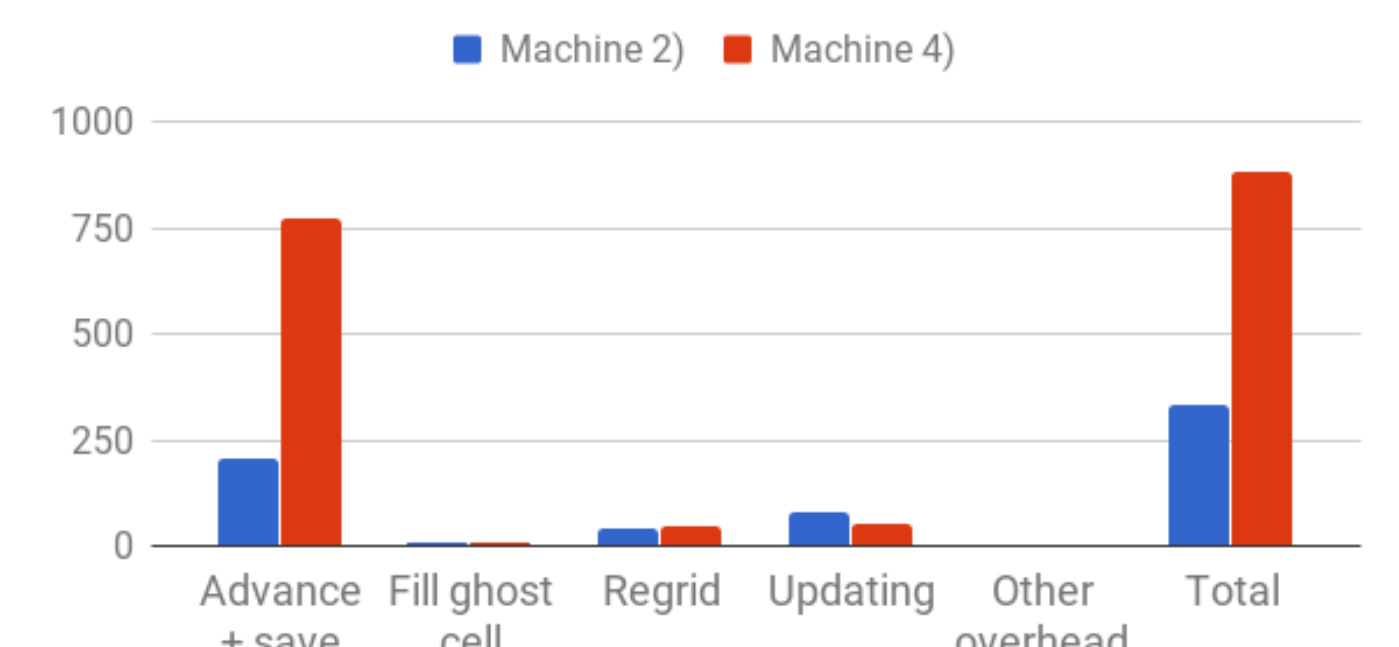


Snapshots of the simulation at $t = 0$ (top) and $t = 0.7$ (bottom). (a) and (b) show pressure field with only level 1 AMR grid patches; (c) and (d) show u-velocity field with only level 2 AMR grid patches; (e) and (f) show v-velocity field with only level 3 AMR grid patches. Note that only grid patch edges, not grid cells, are showed.

Comparison of Running Time



Comparison of Running Time

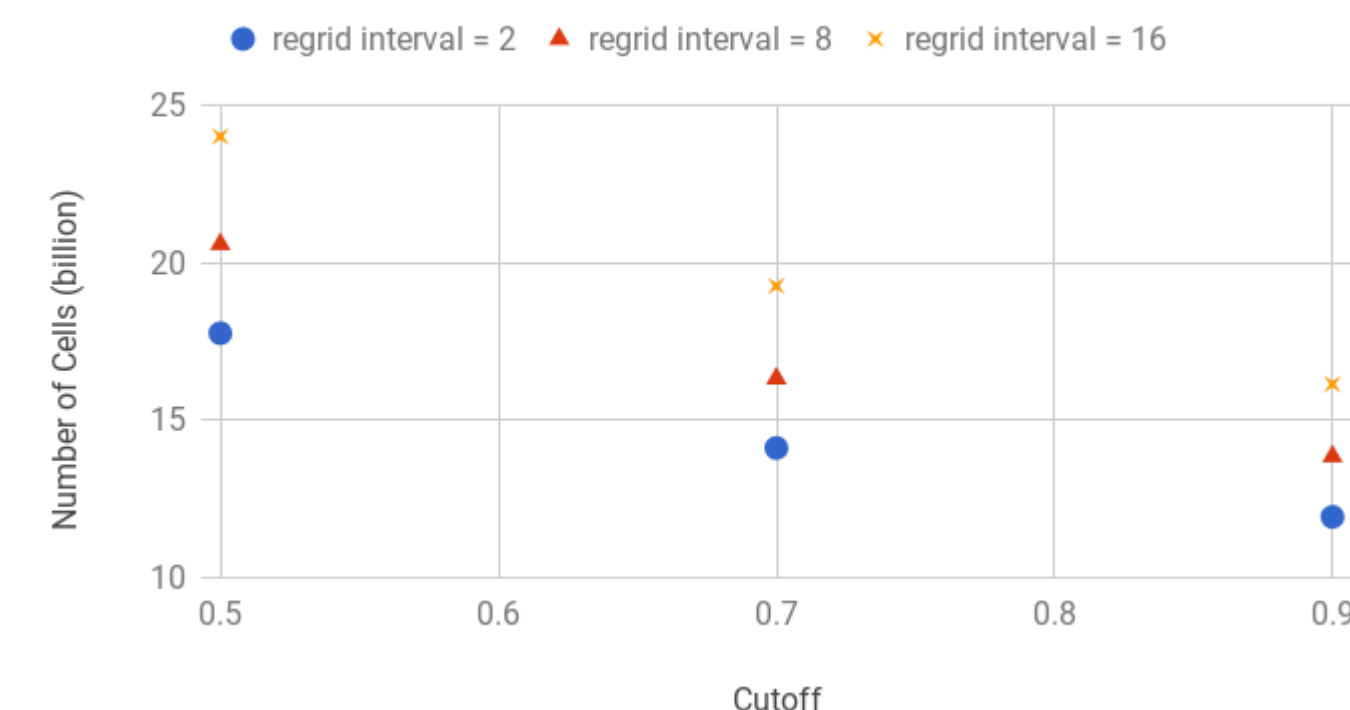


Running Time (in second) of different sections of the code. From left to right: 1) Time on advancing solution and saving fluxes for later use in conservation fix; 2) Time on filling ghost cells; 3) Time on regridding; 4) Time on the updating process; 5) Time on all other overhead; 6) Total running time

The Influence of AMR Parameters:

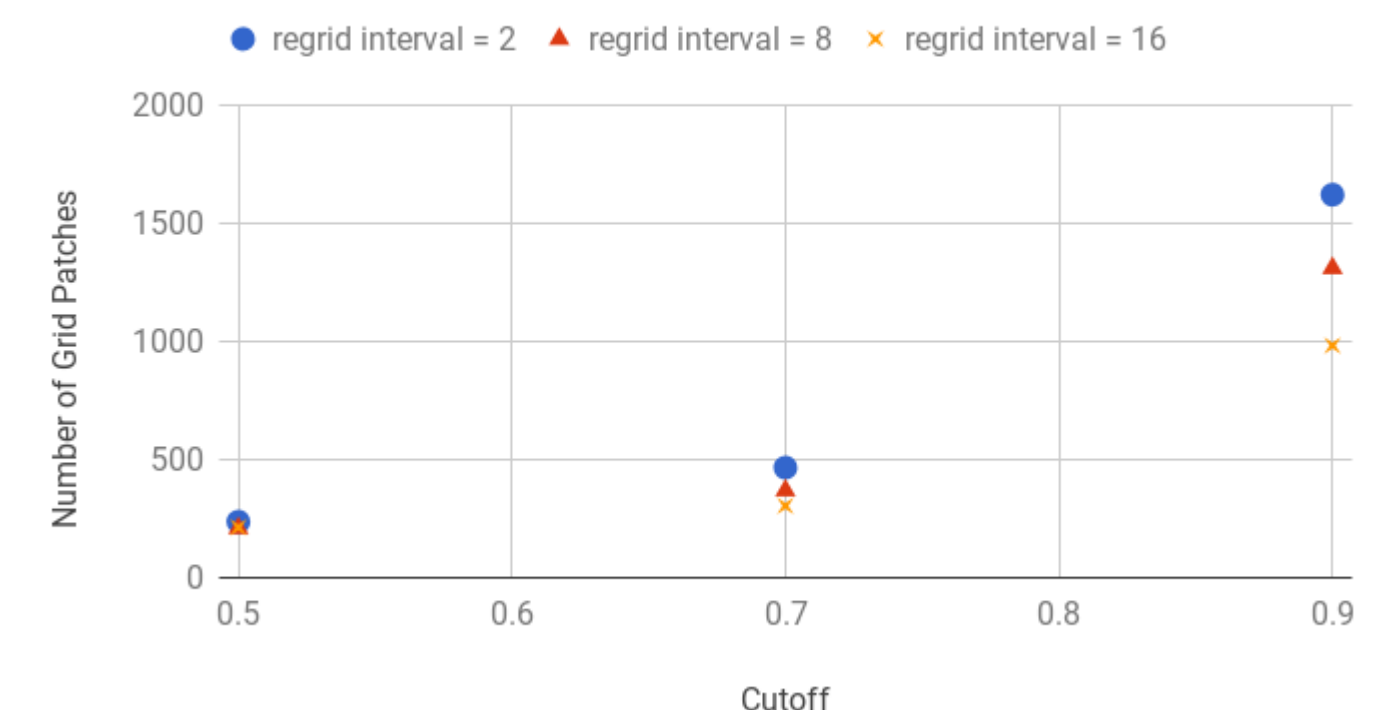
- Cutoff:** specifies the minimum ratio of the number of flagged cells to all cells on a grid patch. A larger cutoff results in a larger collection of smaller grid patches on each level, chosen such that each grid patch does not contain too many unflagged cells. As a result, there are fewer cells in total.
- Regrid interval:** specifies how frequent the regridding process is conducted. If regrid interval is a , usually an extra layer of a cells surrounding the original flagged cells will be flagged during the regridding process, such that the waves in the solution do not propagate beyond the refined region before the next regridding process. As a result, each grid patch is approximately a cells wider at each side and there are more cells in the entire domain.

Total Number of Cells Advanced During the Simulation



Total Number of Cells Advanced on All Levels During the Entire Simulation.

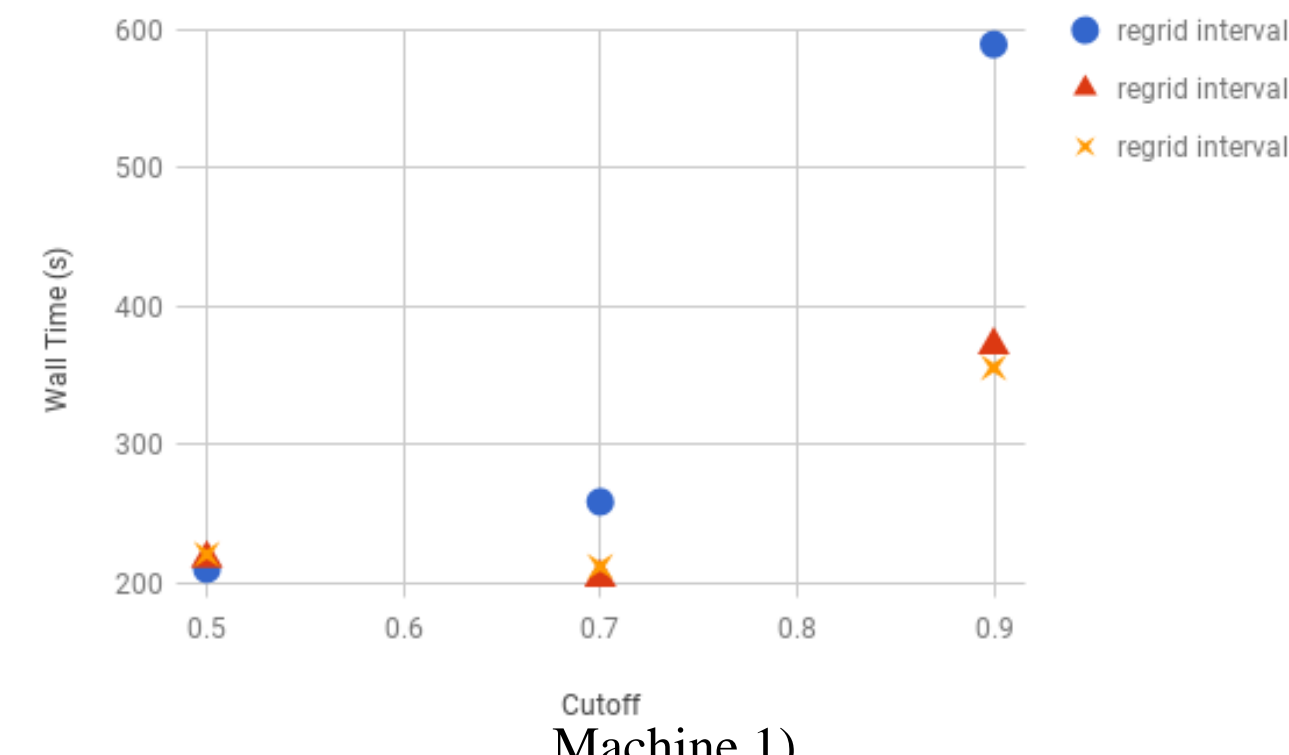
Average Number of Grid Patches at Each Time Step



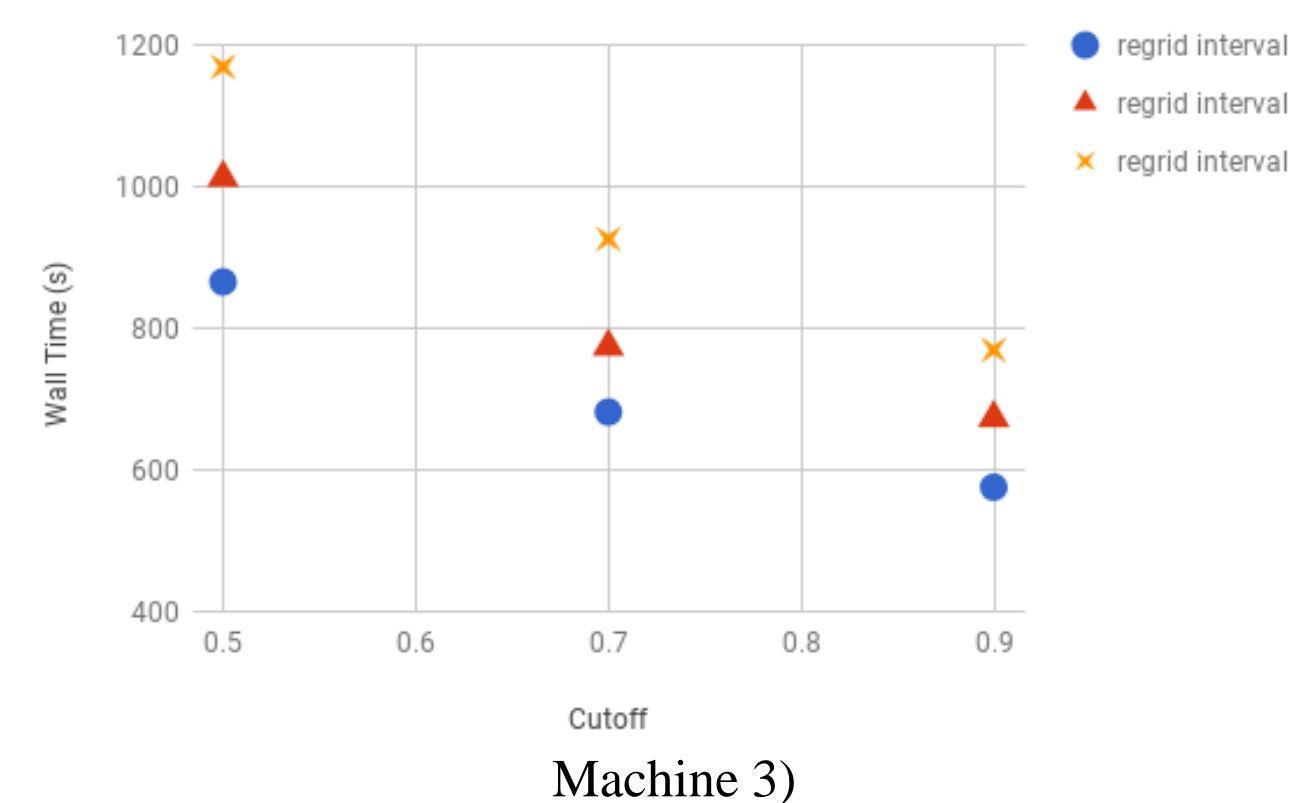
Average Number of Grid Patches on All Levels at Each Time Step.

Running time for different choices of the two parameters:


Wall Time on Advancing Solution and Computing C1, C2 and C3



Wall Time on Advancing Solution and Computing C1, C2 and C3



Wall Time on Advancing Solution and Computing C1, C2 and C3



Wall Time on Advancing Solution and Computing C1, C2 and C3

