

## GAMER: A GRAPHIC PROCESSING UNIT ACCELERATED ADAPTIVE-MESH-REFINEMENT CODE FOR ASTROPHYSICS

HSI-YU SCHIVE<sup>1,2</sup>, YU-CHIH TSAI<sup>1</sup>, AND TZIHONG CHIU<sup>1,2,3,4</sup>

<sup>1</sup> Department of Physics, National Taiwan University, 106, Taipei, Taiwan; [b88202011@ntu.edu.tw](mailto:b88202011@ntu.edu.tw)

<sup>2</sup> Leung Center for Cosmology and Particle Astrophysics (LeCosPA), National Taiwan University, 106, Taipei, Taiwan

<sup>3</sup> Center for Theoretical Sciences, National Taiwan University, 106, Taipei, Taiwan

<sup>4</sup> Center for Quantum Science and Engineering (CQSE), National Taiwan University, 106, Taipei, Taiwan

Received 2009 July 19; accepted 2010 January 4; published 2010 February 4

### ABSTRACT

We present the newly developed code, *GPU-accelerated Adaptive-Mesh-Refinement code (GAMER)*, which adopts a novel approach in improving the performance of adaptive-mesh-refinement (AMR) astrophysical simulations by a large factor with the use of the graphic processing unit (GPU). The AMR implementation is based on a hierarchy of grid patches with an oct-tree data structure. We adopt a three-dimensional relaxing total variation diminishing scheme for the hydrodynamic solver and a multi-level relaxation scheme for the Poisson solver. Both solvers have been implemented in GPU, by which hundreds of patches can be advanced in parallel. The computational overhead associated with the data transfer between the CPU and GPU is carefully reduced by utilizing the capability of asynchronous memory copies in GPU, and the computing time of the ghost-zone values for each patch is diminished by overlapping it with the GPU computations. We demonstrate the accuracy of the code by performing several standard test problems in astrophysics. *GAMER* is a parallel code that can be run in a multi-GPU cluster system. We measure the performance of the code by performing purely baryonic cosmological simulations in different hardware implementations, in which detailed timing analyses provide comparison between the computations with and without GPU(s) acceleration. Maximum speed-up factors of 12.19 and 10.47 are demonstrated using one GPU with  $4096^3$  effective resolution and 16 GPUs with  $8192^3$  effective resolution, respectively.

*Key words:* gravitation – hydrodynamics – methods: numerical

*Online-only material:* color figures

### 1. INTRODUCTION

Numerical simulations have played an indispensable role in modern astrophysics. They serve as powerful tools to probe the fully nonlinear evolutions in various problems, and provide connections between theoretical analyses and observation results. Moreover, thanks to the rapid development of the parallel computing techniques (e.g., the Beowulf clusters, Cell Broadband Engines, and graphic processing units (GPU)), the spatial and mass resolutions as well as computing performance have highly improved in the last decade.

The most essential ingredients in astrophysical simulations are the Newtonian gravity and hydrodynamics. In the last five decades, many studies have been devoted to improving both the accuracy and efficiency of numerical schemes. One of the simplest approaches is to discretize the simulation domain into a fixed number of grid cells, each of which occupies a fixed volume and position. The cell-averaged physical attributes are defined in each cell. The gravitational potential can be evaluated by several different schemes, for instance, the relaxation method, conjugate gradient method, and fast Fourier transform (FFT). As for hydrodynamic evolution, it can be described by the modern high-resolution shock-capturing algorithms, ranging from the first-order Godunov scheme (Godunov 1959), the second-order monotone upwind schemes for conservation laws (MUSCL; van Leer 1979), to the third-order piecewise parabolic method (PPM; Colella & Woodward 1984). In addition, for cosmological simulations, the particle-mesh (PM) scheme (e.g., Klypin & Shandarin 1983; Merz et al. 2005) is often adopted, in which the dark matter is treated as collisionless particles, and the mass density in each grid cell is estimated by the cloud-

in-cell (CIC) technique. Although this uniform-mesh method is relatively easy to implement, it suffers from the enormous memory and computing time requirements when the simulation size increases. Consequently, with the PM method, one must compromise between the size of simulation domain and the spatial resolution.

The Lagrangian particle-based approaches, in which both the collisionless dark matter and the collisional gaseous components are simulated using particles, are alternatives to the uniform-mesh method. The hydrodynamic evolution is generally solved by the smoothed particle hydrodynamics (SPH; Gingold & Monaghan 1977; Lucy 1977), and numerous algorithms have been adopted for solving the gravitational acceleration. The most straightforward scheme is the direct  $N$ -body method, in which all pairwise interactions are calculated. This method, while accurate, is extremely time consuming when the number of particles involved (denoted as  $N$ ) is too large, owing to its  $\mathcal{O}(N^2)$  scaling. Consequently, it is not suitable for the simulations with a large number of particles.

Several approximate algorithms have been developed to improve the computational performance of the gravitational force calculation for particle-based approaches. For example, the particle-particle/particle-mesh ( $P^3M$ ) method (e.g., Hockney & Eastwood 1981; Efsthathiou et al. 1985) improves the spatial resolution by calculating the short-range force with direct summation and adding the long-range PM force. The adaptive  $P^3M$  ( $AP^3M$ ) method (e.g., Couchman 1991) further improves the efficiency of force calculation by adding hierarchically refined submeshes in regions of interest and using the  $P^3M$  method locally to replace the direct pair summation. By contrast, the hierarchical tree algorithm (e.g., Barnes & Hut 1986; Springel et al.

2001) reduces the computational workload by utilizing multi-pole expansion. The force exerted by distant particles is calculated using low-order multi-pole moment, and the  $\mathcal{O}(N \log N)$  scaling is demonstrated. The tree-PM hybrid scheme (e.g., Xu 1995; Bagla 2002; Dubinski et al. 2004; Springel 2005) serves as a further optimization to the tree algorithm. The gravitational potentials are divided into long-range and short-range terms, evaluated by the PM and tree methods, respectively. Since the tree method is applied only locally, the total workload is greatly reduced.

The particle-based approaches offer high computational performance as well as large dynamical range, owing to the Lagrangian nature. However, the main disadvantage of these approaches is their relatively poor capability for simulating hydrodynamics using the SPH method. The resolution is relatively low in high-gradient regions, e.g., around the shock wave in the shock-tube test (Tasker et al. 2008). It also offers poor resolution in the low-density region where the number of particles is insufficient. In addition, the SPH method suffers from artificial viscosity and the inability to accurately capture the hydrodynamic instabilities in certain circumstances (Agertz et al. 2007).

The adaptive-mesh-refinement (AMR) scheme provides a promising approach to combine the accurate shock-capturing property of the uniform-mesh method and the high-resolution, large-dynamical-range property of the Lagrangian particle-based method. The simulation domain is first covered by uniformly distributed meshes (at the “root” level) with a relatively low spatial resolution, and hierarchies of nested refined meshes (at the “refinement” levels) with decreasing grid sizes are then allocated in regions of interest to provide the desired resolution. The gravitational potential can be computed by the multi-grid methods so that a high force resolution can be achieved in the dense region, and the grid-based shock-capturing algorithms can be applied to grids at different refinement levels to preserve the accuracy of hydrodynamic evolution. Since the simulation domain is only locally and adaptively refined, both the memory consumption and the computation time are highly reduced compared to the uniform-mesh method with the same effective resolution.

Detailed comparisons between the AMR and SPH methods have been addressed by several authors (e.g., Regan et al. 2007; Trac et al. 2007; Tasker et al. 2008; Mitchell et al. 2009). It is beyond the scope of this work. In cosmological simulations, the main drawback of the AMR method is, however, the requirement of sufficiently fine grids at the root level in order to provide adequate force resolution at the early epoch. Consequently, the memory consumption and the computation time can be larger than the Lagrangian particle-based method (O’Shea et al. 2005). Nevertheless, the superior capability of handling hydrodynamic properties and the better description for low-density regions make the AMR method a promising and competitive tool in astrophysical simulations, and it has been successfully adopted for large-scale cosmological simulations (e.g., Hallman et al. 2009; Teyssier et al. 2009). More recently, a moving-mesh scheme has been proposed by Springel (2010), aiming at integrating both the advantages of the AMR and particle-based methods.

Several approaches have been developed for the AMR implementations. The most commonly adopted approach is the block-structured AMR, which was first proposed by Berger & Olinger (1984) and Berger & Colella (1989). It has been implemented by many astrophysical codes, e.g., Enzo (Bryan & Norman 1997), AMRA (Plewa & Müller 2001), and CHARM

(Miniati & Colella 2007b). In this approach, the refined sub-domains (often referred as the mesh “patches”) are restricted to be geometrically rectangular, and hence reduce the complexity associated with the discontinuity of resolution across different refinement levels. The size of each patch is variable and adaptable, and patches at the same refinement level can be combined or bisected to fit the local flow geometry. However, the variable patch size also leads to difficulty in parallelizing the code efficiently and to sophisticated data management. In addition, the large patch size can increase the cache-miss rate and thus lower the computational performance.

An alternative approach has been implemented in, for example, the Adaptive Refinement Tree (ART) code (Kravtsov et al. 1997), MLAPM (Knebe et al. 2001), and RAMSES (Teyssier 2002). In this approach, instead of using the rectangular patches as the basic refinement units, the refinement is performed on a cell-by-cell basis. Compared to the block-structured AMR, it features a more efficient refinement configuration related to the local flow geometry, especially in the regions with complex geometry of structures. However, the main drawback of this method lies in the more sophisticated data management, owing to its irregular shape of domain refinement. The interface profiles between cells with different zone spacings are complex, and spatial interpolations must be frequently used to provide the boundary conditions for each cell. Moreover, since the size of the stencils required by the hydrodynamic and Poisson solvers for each cell is generally much larger than a single cell, the computational overhead is large and can lead to serious performance deterioration.

The FLASH code (Fryxell et al. 2000), which uses the PARAMESH AMR library (MacNeice et al. 2000), and the NIRVANA code (Ziegler 2005) have adopted a third approach for the AMR implementation. In this approach, the domain refinement is based on a hierarchy of mesh patches similar to the block-structured AMR, whereas each patch is restricted to have the same number of cells. The typically adopted patch sizes are  $8^3$  in FLASH and  $4^3$  in NIRVANA. Although this restriction will certainly impose the inflexibility of domain refinement and result in a relatively large refined volume compared to the two methods described above, it features several important advantages. First, the data structure and the interfaces between neighboring patches are considerably simplified, which can lead to significant improvements of performance and parallel efficiency. Second, since the additional buffer zones (often referred as the “ghost zones”) for finite-difference stencils are only needed for each patch instead of each cell, the computational overhead associated with the preparation of the ghost-zone data is greatly reduced compared to the cell-based refinement strategy. Finally, fixing the patch size allows for easier optimization of performance and also increases the cache-hit rate. All these features are essential for developing a high-performance code, especially for parallel computing such as using GPUs. Accordingly, in the *GPU-accelerated Adaptive-MESH-Refinement code (GAMER)*, we have adopted this approach as the refinement strategy.

Novel use of modern GPU for acceleration of numerical calculations has becoming a widely adopted technique in the past three years. The original purpose of GPU is to serve as an accelerator for computer graphics. It is designed to work with the Single Instruction, Multiple Data (SIMD) architecture, and processes multiple vertex and fragment data in parallel. The modern GPU, e.g., the NVIDIA Tesla C1060, has 240 scalar processor cores working at 1.3 GHz clock rate. It delivers a

peak performance of 933 Giga Floating Operations Per Second (GFLOPS), which is about an order of magnitude higher than the modern CPU. In addition, it has a 4 GB GDDR3 internal memory with memory bandwidth of  $102 \text{ GB s}^{-1}$ . The 240 scalar processor cores are grouped into 30 multiprocessors, each of which consists of 8 scalar processor cores and shares a 16 KB on-chip data cache. The NVIDIA Tesla S1070 computing system further combines four Tesla C1060 GPUs and offers nearly 4 TFLOPS computing power. Given the natural capability of parallel computing and the enormous computational power of GPU, using GPU for general-purpose computations (GPGPU<sup>5</sup>) have become an active area of research.

The traditional scheme in GPGPU works by using the high-level shading languages, which are designed for graphic rendering and require familiarity with computer graphics. It is therefore difficult and unsuitable for general-purpose computations. In 2006, the NVIDIA Corporation released a new computing architecture in GPU, the Compute Unified Device Architecture (CUDA; NVIDIA 2008). It was designed for a general-purpose usage and greatly lowers the threshold of using GPU for *non-graphic* computations. In CUDA, GPU is regarded as a multi-threaded coprocessor to CPU with a standard C language interface. To define the computational task for GPU, programmers should provide a C-like function called “kernel,” which can be executed by multiple “CUDA threads” in parallel. A unique thread ID is given to each thread in order to distinguish between different threads.

As an illustration, we consider the sum of two vectors, each of which has  $M$  elements. In CUDA, instead of writing a loop to perform  $M$  summation operations sequentially, we define a single summation operation in a kernel and use  $M$  threads. These  $M$  threads will execute the same kernel in parallel but perform the single summation operation on different vector elements. In this example, the thread ID may be used to define the targeted vector element for each thread.

Note that this scenario is analogous to the parallel computing in a Beowulf cluster using the message passing interface (MPI), in which a single program is simultaneously executed by multiple processors, and each process is given a unique ID (“MPI rank”). However, performance optimization in GPU is not straightforward and requires elaborate numerical algorithms dedicated to the GPU specifications. Especially, note that CUDA threads are further grouped into multiple “thread blocks.” Threads belonging to the same thread block are executed by one multiprocessor and can share data through an on-chip data cache (referred as the “shared memory”). This data cache is very small (typically only 16 KB per multiprocessor) but has much lower memory latency than the off-chip DRAMS (referred as the “global memory”). Accordingly, the numerical algorithms must be carefully designed to store common and frequently used data in this fast shared memory so that the memory bandwidth bottleneck may be removed.

Nowadays, the most successful approach to utilize the GPU computing power in astrophysical simulations is the direct  $N$ -body calculation (e.g., Belleman et al. 2008; Schive et al. 2008; Gaburov et al. 2009). Schive et al. (2008) have built a multi-GPU computing cluster named GraCCA (Graphic-Card Cluster for Astrophysics), and have demonstrated its capability for direct  $N$ -body simulations in terms of both high computational performance and high parallel efficiency. The direct calculation of all  $N^2$  pairwise interactions is extremely computational

intensive, and thus is relatively straightforward to obtain high performance in GPU. However, the direct  $N$ -body simulations can address only a limited range of problems. Aubert et al. (2009) have proposed a GPU-accelerated PM integrator using a single GPU. It remains considerably challenging and unclear whether the performance of other kinds of astrophysical simulations with complex data structure and relatively low arithmetic intensity, such as the AMR simulations, can be highly improved by using GPUs, especially in a multi-GPU system.

In this paper, we present the first GPU-accelerated, AMR, astrophysics-dedicated, and parallelized code, named *GAMER*. We give a detailed description of the numerical algorithms adopted in this code, especially focusing on the GPU implementations. The accuracy of the code is demonstrated by performing various test problems. Detailed timing analyses of individual GPU solvers as well as the complete program are conducted with different hardware implementations. In each timing test, we further compare the performances of runs with and without GPU(s) acceleration.

The paper is organized as follows. In Section 2, we describe the numerical schemes adopted in *GAMER*, including the AMR method, both the hydrodynamic scheme and the Poisson solver, and the parallelization strategy. We then focus on the GPU implementations of different parts in the code, along with individual performance measurements in Section 3. In Section 4, we present the simulation results of several test problems to demonstrate the accuracy. Detailed timing analyses of the complete program in purely baryonic cosmological simulations are presented in Section 5. Finally, we summarize the work and discuss the future outlooks in Section 6.

## 2. NUMERICAL SCHEME

In this section, we describe in detail the numerical schemes adopted in *GAMER*, including the AMR implementation, the algorithms of both hydrodynamics and self-gravity, and the parallelization strategy. To provide a more comprehensible description, here we focus on the generic algorithms that are unrelated to the hardware implementation. Important features related to the GPU implementation will be emphasized and a more detailed description will be given in Section 3.


### 2.1. Adaptive Mesh Refinement

The AMR scheme implemented in *GAMER* is similar to that adopted by FLASH, in which the computational domain is covered by a hierarchy of grid patches with similar shape but different spatial resolutions. In *GAMER*, a grid patch is defined to have a fixed number of grid cells in each spatial direction. The computational domain is first covered by root patches with the lowest spatial resolution. Then, according to the user-defined refinement criteria, each root patch may be refined into eight child patches with a spatial resolution twice that of their parent patch. The same refinement operation may further be applied to all patches in different refinement levels, in which a patch at level  $\ell$  has a spatial resolution  $2^\ell$  times higher than that of a root patch at level zero. Accordingly, a hierarchy of grid patches with oct-tree data structure is dynamically and adaptively constructed during the simulation. In Figure 1, we show a two-dimensional example of the refinement map.

Patches are the basic units in *GAMER*. Owing to the oct-tree data structure, eight patches are always allocated or deallocated

<sup>5</sup> <http://gpgpu.org/>





**Figure 1.** Two-dimensional example of the refinement map. Each patch is composed of  $8 \times 8$  cells. The borders of patches are highlighted by bold lines. One patch may be refined into four child patches with a spatial resolution twice that of their parent patch. A jump of more than one refinement level between nearby patches is forbidden (even in the diagonal direction).

simultaneously. The data stored in each patch include its own physical variables, the absolute coordinates in the computational domain, the indices of the parent, child, and 26 sibling patches, the pointers of flux arrays corresponding to six patch boundary surfaces, and the flag recording its refinement status.

Restricting all patches to be geometrically similar to each other greatly simplifies the AMR framework, with respect to both the structure of the program and the GPU implementation. A single GPU kernel can be applied to all patches, even in different refinement levels. Moreover, since the amount of computation workload of each patch is the same, there will be no synchronization overhead when multiple patches are evolved in parallel by GPU. However, it does impose certain inflexibility of spatial refinement. The region being refined will be larger than necessary, especially when the volume of a single patch is too large. On the other hand, having a small-volume patch will introduce higher computational overhead associated with the preparation of the ghost-zone data. In *GAMER*, the optimized size of a single patch is set to  $8^3$ . It will be demonstrated in Sections 3 and 5 that by exploiting the feature of parallel execution between CPU and GPU, the ghost-zone filling time can be overlapped with the execution time of the GPU solvers, and yields considerable performance enhancement.

*GAMER* can be used as either a purely hydrodynamic or a coupled self-gravity and hydrodynamic code. When only the hydrodynamic module is activated, the code supports both the uniform and individual time step algorithms. The time step of level  $\ell$  may be either equal to or twice smaller than that of level  $\ell - 1$ . However, when the gravity module is also activated, the code currently only supports the uniform time step algorithm. The same time step is applied to all levels, and the evolution of patches at level  $\ell$  proceeds in the steps as follows:

1. Update physical quantities for all patches at level  $\ell$ .
2. Begin the evolution of the next refinement level if there are patches at level  $\ell + 1$ .


3. Correct the physical quantities at level  $\ell$  by using the updated results at level  $\ell + 1$ .
4. Rebuild the refinement map at level  $\ell$ .

Since the fine-grid values are presumably more accurate than the coarse-grid values, there are two cases where the data of a coarse patch require further correction in the step 3. First, if a coarse patch is overlaid by its child patches, its values are simply replaced by the spatial average of the fine-grid values. Second, if the border of a coarse patch is near the boundary of refinement, the flux correction operation (Berger & Colella 1989) is applied. Then, a corresponding flux array is allocated. This array will store the difference between the coarse-grid flux and the fine-grid flux across the coarse–fine boundary, and will be used to correct the coarse-grid values adjacent to this boundary. This flux correction operation ensures that the flux out of the coarse-grid patch is equal to the flux into the fine-grid patch, and therefore the conservation of hydrodynamic variables is preserved (assuming no self-gravity).

Rebuilding the refinement map in step 4 takes two sub-steps: a flag step, followed by a refinement step. A patch is flagged for refinement if any cell inside the patch satisfies the refinement criteria. In *GAMER*, both the hydrodynamic variables and their gradients can be taken as the refinement criteria. There is however a situation requiring special treatment during the flag step. Since the refinement map is always rebuilt from finer levels to coarser levels, a patch at level  $\ell$  may not be flagged even if its child patches at level  $\ell + 1$  have already been flagged. In this case, the patch at level  $\ell$  is also flagged to ensure that the fine-grid data are preserved. Finally, a proper-nesting constraint is applied to all patches. It prohibits the spatial refinement from jumping more than one level across two adjacent patches. Patches failing to satisfy this constraint are unflagged.

In the refinement step, eight child patches at level  $\ell + 1$  are constructed for each flagged patch at level  $\ell$ . The hydrodynamic data of child patches are either directly inherited from existing data or filled via conservation-preserving interpolation from their parent patches. The Min-Mod limiter is used to ensure the monotonicity of interpolation. The indices of parent, child, and sibling patches are stored, and null values are assigned to them if the corresponding child or sibling patches do not exist. Finally, the flux arrays are properly allocated for patches adjacent to the coarse–fine boundaries.

The frequency of rebuilding the refinement map is also a free parameter provided by users. The guideline is that the refinement map must be rebuilt before the regions of interest propagate away from fine-grid patches into coarse-grid patches. Although in the extreme case we may rebuild the refinement map in every step, it is too expensive in time and not practical in general situations. Therefore, in order to reduce the frequency of performing the refinement operation, we follow the scheme suggested by Berger & Colella (1989). A free parameter  $N_b$  is provided to define the size of the *flag buffer*. If a cell exceeds the refinement threshold during the flag check,  $(1 + 2N_b)^3 - 1$  cells surrounding this cell are regarded as the flag buffers. If any of these flag buffers extends across the patch border, the corresponding sibling patch is also flagged for refinement (as long as it satisfies the proper-nesting condition). Figure 2 shows an example of the refinement result with  $N_b = 3$ . The extreme case is to have  $N_b$  equal to the size of a single patch, in which case all 26 sibling patches will always be flagged if the central patch is flagged. Generally speaking, the larger the number  $N_b$ , the longer the period adopted between two refinement steps might be.



**Figure 2.** Two-dimensional example of the flag operation. The size of the flag buffer ( $N_b$ ) is set to 3. The borders of the patches are highlighted by bold lines. The filled circles represent the cells satisfying the refinement criteria, and the open circles represent their corresponding flag buffers. The numbers at the center of each patch stand for the patch indices. In this example, although there are no cells fulfilling the refinement criteria in patches 1–3, they are still flagged due to the extension of the flag buffers.

The procedure of patch construction during the initialization is different from that during the simulation. As illustrated by the evolution procedure described previously, the patch construction during the simulation is always performed from finer levels to coarser levels. It ensures that the fine-grid data are predominant of patch refinement. By contrast, the spatial resolution of the initial condition is solely provided by users. Accordingly, in *GAMER*, three kinds of initialization methods are supported.

First, a user-defined initialization function can be applied to set the initial value of each physical quantity. The patch construction starts from level 0 up to the maximum level. If any patch at level  $\ell$  satisfies the refinement criteria, eight child patches at level  $\ell + 1$  are allocated and initialized by the same initialization function. After patch construction, a restriction operation is performed, starting from the maximum level down to the root level, in order to ensure that the physical quantities of a cell at level  $\ell$  is always equal to the spatial average of its eight child cells at level  $\ell + 1$ .

Second, the code can load an array storing the uniform-mesh data as the initial condition. Assuming that the input data have spatial resolution equal to the refinement level  $\ell$ , patches at level  $\ell$  are first constructed. After that, a restriction operation is performed from level  $\ell$  down to level 0 to construct patches of levels  $< \ell$ . Any patch failing to satisfy the refinement criteria at the current level will be removed. Since we assume that the input uniform-mesh data possess the highest resolution during the initialization, no patch at level  $> \ell$  is allocated at this stage. However, patches of higher resolution can still be constructed during the run, and hence the highest resolution is not limited by the initial input data.

The third initialization procedure loads any of the previous data dumps as the restart file. It is essential when the program is terminated unexpectedly, and it also provides an efficient way for tuning parameters and analyzing simulation results.

## 2.2. Hydrodynamics

In *GAMER*, the Euler equations are solved in conservative forms:

$$\frac{\partial \rho}{\partial t} + \frac{\partial}{\partial x_j}(\rho v_j) = 0, \quad (1)$$

$$\frac{\partial(\rho v_i)}{\partial t} + \frac{\partial}{\partial x_j}(\rho v_i v_j + P \delta_{ij}) = -\rho \frac{\partial \phi}{\partial x_i}, \quad (2)$$

$$\frac{\partial e}{\partial t} + \frac{\partial}{\partial x_j}[(e + P)v_j] = -\rho v_j \frac{\partial \phi}{\partial x_j}, \quad (3)$$

where  $\rho$  is the mass density,  $v$  is the flow velocity,  $P$  is the thermal pressure,  $e$  is the total energy density, and  $\phi$  is the gravitational potential. The relation between pressure  $P$  and total energy density  $e$  is given by

$$e = \frac{1}{2}\rho v^2 + \epsilon, \quad (4)$$

$$P = (\gamma - 1)\epsilon, \quad (5)$$

where  $\epsilon$  is the internal thermal energy density and  $\gamma$  is the ratio of specific heats. The self-gravity is included in the Euler equations as a source term, and will be addressed in more detail in the following subsection.

The hydrodynamic scheme adopted in *GAMER* is based on the algorithm proposed by Trac & Pen (2003). It is a second-order accurate relaxing total variation diminishing (TVD) scheme (Jin & Xin 1995), which has been implemented and well tested in both the hydrodynamic simulation (Trac & Pen 2003) and the magnetohydrodynamic simulation (Pen et al. 2003). In the following, we first review the one-dimensional relaxing TVD scheme, and then follow the generalization to the three-dimensional case.

Consider the one-dimensional Euler equation in the vector form:

$$\frac{\partial \mathbf{u}}{\partial t} + \frac{\partial \mathbf{F}(\mathbf{u})}{\partial x} = 0, \quad (6)$$

where  $\mathbf{u} = (\rho, \rho v, e)$  is the flow-variable vector and  $\mathbf{F}(\mathbf{u}) = (\rho v, \rho v^2 + P, e v + P v)$  is the corresponding flux vector. First, a free positive function  $c(x, t)$ , which is referred to as the freezing speed, is evaluated, and an auxiliary vector is defined by  $\mathbf{w} \equiv \mathbf{F}(\mathbf{u})/c$ . To guarantee the TVD condition, the freezing speed  $c$  must be greater than the speed of information propagation. For the one-dimensional Euler equation, this requirement is satisfied by having  $c(x, t) = |v(x, t)| + c_s(x, t)$ , where  $c_s$  is the sound speed.

The flux term in Equation (6) is then decomposed into two terms,

$$\frac{\partial \mathbf{u}}{\partial t} + \frac{\partial \mathbf{F}^R}{\partial x} - \frac{\partial \mathbf{F}^L}{\partial x} = 0, \quad (7)$$

where

$$\mathbf{F}^R \equiv c \left( \frac{\mathbf{u} + \mathbf{w}}{2} \right) \text{ and } \mathbf{F}^L \equiv c \left( \frac{\mathbf{u} - \mathbf{w}}{2} \right) \quad (8)$$

are referred to as the right-moving and left-moving fluxes with advection speed  $c$ , respectively. Since these two fluxes have well-defined directions, the MUSCL scheme can be applied straightforwardly. Let  $\mathbf{u}_n^t$  denote the cell-centered value of the cell  $n$  at time  $t$ , and  $\mathbf{F}_n^t$  denote the corresponding cell-center flux. To integrate Equation (7) in a conservative form, the fluxes  $\mathbf{F}_{n\pm 1/2}^{R,t}$  and  $\mathbf{F}_{n\pm 1/2}^{L,t}$  defined at the boundaries of the cell  $n$  must be evaluated. In the following, we describe the algorithm to evaluate  $\mathbf{F}_{n+1/2}^{R,t}$  as an illustration.  $\mathbf{F}_{n-1/2}^{R,t}$  and  $\mathbf{F}_{n\pm 1/2}^{L,t}$  can be derived in a similar way.

In the first step, the upwind scheme is used to assign a value to the boundary flux as a first-order approximation. Since  $\mathbf{F}_{n+1/2}^{R,t}$  has a positive advection velocity, we can simply set  $\mathbf{F}_{n+1/2}^{R,t} = \mathbf{F}_n^t$ . The second-order correction  $\Delta \mathbf{F}_{n+1/2}^{\text{TVD},t}$  satisfying

the TVD condition is obtained by applying a flux limiter  $\phi$  to two second-order flux corrections,

$$\Delta \mathbf{F}_{n+1/2}^{\text{TVD},t} = \phi(\Delta \mathbf{F}_{n+1/2}^{(1),t}, \Delta \mathbf{F}_{n+1/2}^{(2),t}), \quad (9)$$

where

$$\Delta \mathbf{F}_{n+1/2}^{(1),t} = \frac{\mathbf{F}_n^t - \mathbf{F}_{n-1}^t}{2} \quad \text{and} \quad \Delta \mathbf{F}_{n+1/2}^{(2),t} = \frac{\mathbf{F}_{n+1}^t - \mathbf{F}_n^t}{2}. \quad (10)$$

The flux limiter adopted in the current implementation is the van Leer limiter (van Leer 1974), which takes the harmonic average of two second-order flux corrections:

$$\phi_{\text{vanLeer}}(\Delta \mathbf{F}^{(1)}, \Delta \mathbf{F}^{(2)}) = \begin{cases} \frac{2\Delta \mathbf{F}^{(1)}\Delta \mathbf{F}^{(2)}}{\Delta \mathbf{F}^{(1)} + \Delta \mathbf{F}^{(2)}}, & \text{if } \Delta \mathbf{F}^{(1)}\Delta \mathbf{F}^{(2)} > 0, \\ 0, & \text{if } \Delta \mathbf{F}^{(1)}\Delta \mathbf{F}^{(2)} \leq 0. \end{cases} \quad (11)$$

Note that, as indicated by Equation (11), no second-order correction is applied to  $\mathbf{F}_{n+1/2}^{R,t}$  if  $\mathbf{F}_n^t$  assumes a local extreme value, and hence the hydrodynamic scheme is locally reduced to only first-order accurate.

Finally, the second-order accurate right-moving flux is given by

$$\mathbf{F}_{n+1/2}^{R,t} = \mathbf{F}_n^t + \Delta \mathbf{F}_{n+1/2}^{\text{TVD},t}. \quad (12)$$

$\mathbf{F}_{n-1/2}^{R,t}$  and  $\mathbf{F}_{n\pm 1/2}^{L,t}$  can be evaluated in the way similar to Equation (12).

To achieve second-order accuracy in time as well, the second-order Runge–Kutta method (also known as the midpoint method) is adopted for the time integration. First, the temporal midpoint value  $\mathbf{u}_n^{t+\Delta t/2}$  is evaluated by

$$\mathbf{u}_n^{t+\Delta t/2} = \mathbf{u}_n^t - \left( \frac{\mathbf{F}_{n+1/2}^t - \mathbf{F}_{n-1/2}^t}{\Delta x} \right) \frac{\Delta t}{2}, \quad (13)$$

where  $\mathbf{F}_{n+1/2}^t = \mathbf{F}_{n+1/2}^{R,t} - \mathbf{F}_{n+1/2}^{L,t}$  is computed by the first-order upwind scheme. The midpoint fluxes  $\mathbf{F}^{t+\Delta t/2}$  are then computed by applying the second-order TVD scheme to  $\mathbf{u}^{t+\Delta t/2}$ . Eventually, the full-step value  $\mathbf{u}_n^{t+\Delta t}$  is given by

$$\mathbf{u}_n^{t+\Delta t} = \mathbf{u}_n^t - \left( \frac{\mathbf{F}_{n+1/2}^{t+\Delta t/2} - \mathbf{F}_{n-1/2}^{t+\Delta t/2}}{\Delta x} \right) \Delta t. \quad (14)$$

It is straightforward to generalize the one-dimensional TVD scheme described above to three dimensions by applying the dimensional-splitting method (Strang 1968). The three-dimensional Euler equations are solved by first applying a forward sweep in the order  $xyz$  and following a backward sweep in the order  $zyx$ . The same time step must be employed by these two sweeps to maintain the second-order accuracy. The dimensional splitting method also makes it easy to parallelize the computation of the three-dimensional Euler equations, as addressed by Trac & Pen (2003). By taking advantage of this feature, a high-performance GPU hydrodynamic solver based on the above TVD scheme has been implemented in *GAMER*. It will be described in detail in Section 3.1.

The one-dimensional TVD scheme uses a seven-point stencil (one cell on each side for evaluating the midpoint values by the upwind scheme plus two cells on each side for evaluating the full-step values by the TVD scheme). Therefore, three

ghost zones are required on each side in each spatial direction to update the hydrodynamic variables in a single patch. The ghost-zone values are filled in two ways. If a desired sibling patch exists, they are filled by a direct memory copy. Otherwise, they are filled by linear interpolation with the Min-Mod limiter from patches one level coarser. Since the proper-nesting condition is fulfilled everywhere in the simulation domain, interpolation from patches two (or more) levels coarser is prevented. Note that computing the ghost-zone values can lead to a significant computational overhead in almost all kinds of AMR implementations. Nevertheless, this issue is well handled in *GAMER* and will be addressed in Sections 3 and 5.

*GAMER* can work in both the physical coordinates and the comoving coordinates. For the cosmological hydrodynamic simulation, the forms of Euler equations (Equations (1)–(3)) may remain unchanged by applying the following changes of variables:

$$\tilde{\mathbf{x}} = \frac{\mathbf{x}}{a}, \quad d\tilde{t} = \frac{dt}{a^2}, \quad \tilde{\rho} = a^3 \rho, \quad (15)$$

$$\tilde{\mathbf{v}} = a(\mathbf{v} - H\mathbf{x}), \quad \tilde{P} = a^5 P, \quad \tilde{\phi} = a^2(\phi - \phi_b), \quad (16)$$

where  $a$  is the cosmological scale factor,  $H$  is the Hubble parameter, and  $\phi_b$  is the gravitational potential related to the background density (here, we have assumed that the ratio of specific heats  $\gamma = 5/3$ ). It makes the *GAMER* code more flexible and hence can be applied to different aspects of astrophysical applications.

### 2.3. Self-gravity

The gravitational potential is evaluated via solving the Poisson equation

$$\nabla^2 \phi(\mathbf{x}) = 4\pi G \rho(\mathbf{x}), \quad (17)$$

where  $G$  is the gravitational constant. In its discrete form, the Laplacian operator  $\nabla^2$  can be replaced by a seven-point finite difference operator:

$$\frac{1}{\Delta h_\ell^2} (\phi_{i+1,j,k}^{t,\ell} + \phi_{i,j+1,k}^{t,\ell} + \phi_{i,j,k+1}^{t,\ell} + \phi_{i-1,j,k}^{t,\ell} + \phi_{i,j-1,k}^{t,\ell} + \phi_{i,j,k-1}^{t,\ell} - 6\phi_{i,j,k}^{t,\ell}) = 4\pi G \rho_{i,j,k}^{t,\ell}, \quad (18)$$

where  $\Delta h_\ell$  is the zone spacing at level  $\ell$ .

In *GAMER*, two numerical methods have been implemented to solve Equation (18). At the root level, where the coarsest patches always cover the entire computational domain, we adopt the standard FFT method. A Green's function associated with the discretized Laplacian operator is used, and the periodic boundary condition is assumed. At the refined levels, where in general the computational domain is only partially refined and hence Equation (18) cannot be solved globally, we adopt the successive overrelaxation method (SOR; Press et al. 2007) with the Dirichlet boundary condition. Only one ghost zone is required on each side in each spatial direction to evaluate the potential in a single patch, and the ghost-zone values are filled by interpolation from the patches one level coarser.

The SOR scheme starts with evaluating the residual of each cell by

$$R_{i,j,k}^\ell = (\phi_{i+1,j,k}^{\text{old},\ell} + \phi_{i,j+1,k}^{\text{old},\ell} + \phi_{i,j,k+1}^{\text{old},\ell} + \phi_{i-1,j,k}^{\text{old},\ell} + \phi_{i,j-1,k}^{\text{old},\ell} + \phi_{i,j,k-1}^{\text{old},\ell} - 6\phi_{i,j,k}^{\text{old},\ell}) - 4\pi G \rho_{i,j,k}^\ell \Delta h_\ell^2, \quad (19)$$

where the superscript “old” indicates the values at a previous step. The updated values are then given by

$$\phi_{i,j,k}^{\text{new},\ell} = \phi_{i,j,k}^{\text{old},\ell} + \frac{1}{6}\omega R_{i,j,k}^{\ell}, \quad (20)$$


where  $\omega$  is the overrelaxation parameter. Equations (19) and (20) are solved iteratively until the 1-norm of the residual has been diminished to the floating-point precision limit (compared to the source term  $4\pi G\rho\Delta h^2$ ) to ensure the solution of potential has converged.

The odd–even ordering is adopted to determine the order in which different cells in a patch are updated by Equations (19) and (20). A cell with position indices  $(i, j, k)$  is regarded as an “even” cell if  $i + j + k$  is even, and an “odd” cell if  $i + j + k$  is odd. This scheme is also referred as the “red–black ordering” since the odd and even cells are defined in the way just like the red and black squares of a checkerboard in the two-dimensional case. During one iteration, the even cells are first updated, and then these *updated* values are used to update the odd cells. As can be seen from Equation (19), the updates of even cells depend only on the odd cells, and vice versa. Therefore, by exploiting the odd–even ordering, all even cells can be calculated independently at the first half-step (hereafter referred to as the even step), and all odd cells can also be calculated independently at the second half-step (hereafter referred to as the odd step). The independent operations reveal the possibility of parallel computing. This property plays a crucial role in the development of an efficient and memory-saving GPU kernel for the SOR scheme.

In real astrophysical simulations, generally the number of root-level patches takes only a small fraction ( $< 10\%$ ) of the total number of patches in all levels, and hence the execution time of the root level is much less than the total simulation time. Accordingly, for the root-level Poisson solver, we execute on the CPU the free available package FFTW (Frigo & Johnson 1998), which is a highly optimized, parallelized, and portable package for solving the discrete Fourier transform (DFT). On the contrary, for the refinement levels where the Poisson equation is solved via the SOR scheme, a high-performance GPU Poisson solver has been implemented in *GAMER*. It will be described in detail in Section 3.2.

Solving the Poisson equation in an AMR framework requires additional attention. The primary issue is that the boundary condition for a fine-grid patch is always obtained by interpolation from the coarse-grid values. The potential is continuous across the coarse–fine interface; however, the normal derivative of the potential is not necessarily so. The discontinuity in normal derivative of the potential acts as a *pseudomass sheet* on the coarse–fine interface, and will eventually contaminate the solution of potential in finer levels.

Several methods have been proposed to improve the solution of the Poisson equation in adaptively refined meshes (e.g., Martin & Cartwright 1996; Huang & Greengard 2000; Ricker 2008). In *GAMER*, the two-level potential correction is in work with the following procedure. First, we estimate the pseudomass sheet on all interfaces between a parent patch and each of its eight child patches. A two-dimensional example of the mesh structure adjacent to a coarse–fine interface is shown in Figure 3, in which  $\phi_{i_c,j_c}^c$  denotes a coarse-grid potential to the left of the interface, and  $\phi_{i_f,j_f}^f$  and  $\phi_{i_f,j_f+1}^f$  denote the corresponding two fine-grid potentials (defined in the ghost



**Figure 3.** Two-dimensional example of the mesh structure adjacent to a coarse–fine interface. The filled triangles represent the coarse-grid potentials, the filled circles represent the fine-grid potentials, and the open circles represent the fine-grid potentials in the ghost zones. The values used for approximating the normal derivatives of potentials are connected by two-way arrows.

zones). The pseudomass sheet  $\xi$  is then defined as

$$\xi(\mathbf{x}) = \frac{1}{4\pi G\Delta h^c} \left( \left. \frac{\partial\phi}{\partial n} \right|_{\epsilon^-} - \left. \frac{\partial\phi}{\partial n} \right|_{\epsilon^+} \right), \quad (21)$$

in which the normal derivatives of potentials are approximated by

$$\left. \frac{\partial\phi}{\partial n} \right|_{\epsilon^-} = \frac{\phi_{i_c,j_c}^c - \phi_{i_c+1,j_c}^c}{\Delta h^c} \quad (22)$$

on the coarse-grid side and

$$\left. \frac{\partial\phi}{\partial n} \right|_{\epsilon^+} = \frac{(\phi_{i_f,j_f}^f + \phi_{i_f,j_f+1}^f - \phi_{i_f+1,j_f}^f - \phi_{i_f+1,j_f+1}^f)/2}{\Delta h^f} \quad (23)$$

on the fine-grid side.

To compensate the mass-sheet potentials, we first place the *negative* of the pseudomass sheet in the coarse meshes adjacent to the coarse–fine boundary, and evaluate the correction to the coarse-grid potential (denoted as  $\zeta^c$ ) by solving the correction equation

$$\nabla^2\zeta(\mathbf{x}) = -4\pi G\xi(\mathbf{x}). \quad (24)$$

Afterward, the fine-grid correction can also be solved by Equation (24), in which the boundary condition is provided through the interpolation of the coarse-grid correction. Finally, the corrected potentials in both coarse and fine grids are obtained by

$$\phi_{\text{corrected}}(\mathbf{x}) = \phi_{\text{uncorrected}}(\mathbf{x}) + \zeta(\mathbf{x}). \quad (25)$$

Note that for AMR implementations that demand patches in every refinement level to have identical shape (e.g., the FLASH and *GAMER* codes), the Poisson equation is solved on a patch-by-patch basis in order to diminish the data transfer between adjacent patches (Ricker 2008). Consequently, similar numerical errors are also introduced from the interfaces between patches at the same refinement level. Nevertheless, since Equations (21)–(23) are used at *all interfaces* of a parent patch and each of its eight child patches, this numerical error can be made diminished in Equation (25).

To further improve the solution across the inter-patch boundaries, a *sibling relaxation step* is employed immediately after



solving Equations (17) and (24). In this step, instead of using the interpolated values as the fine-grid boundary conditions, we exchange the values near the inter-patch boundaries between neighboring patches, store them in the corresponding ghost zones, and apply another SOR iteration. Finally, when necessary, a few times of the sibling relaxation steps may also be employed after Equation (25). Nevertheless, it only gives rise to minor corrections to the final results. In *GAMER*, typically we apply one sibling relaxation step after solving Equations (17) and (24), and two to four sibling relaxation steps after Equation (25).

The procedure of the two-level solution described above can be straightforwardly generalized to the multi-level case. It works in the steps as follows.

1. In each level  $\ell$  from the root level ( $\ell = 0$ ) up to the maximum level  $\ell_{\max}$ , evaluate the uncorrected potential  $\phi_{\text{uncorrected}}$  by solving Equation (17). Immediately after that, perform one sibling relaxation step.
2. In each level  $\ell$  from level  $\ell_{\max} - 1$  down to level 0, evaluate the pseudomass sheet  $\xi$  via Equations (21)–(23), and add up with the pseudomass sheet derived from patches at level  $\ell + 1$ .
3. In each level  $\ell$  from level 0 up to level  $\ell_{\max}$ , evaluate the correction  $\zeta$  by solving Equation (24). Immediately after that, perform one sibling relaxation step.
4. For each level  $\ell$  from level 0 up to level  $\ell_{\max}$ , obtain the corrected potential  $\phi_{\text{corrected}}$  by Equation (25). Afterward, if necessary, perform a few times of sibling relaxation steps.

Note that this multi-level procedure is similar to the scheme proposed by Ricker (2008). However, instead of using the residuals as the source function to solve the potential correction (as applied in Ricker's scheme), our method evaluates the potential correction via estimating the pseudomass sheet and features the conservation of mass. This feature is described below in more detail.

The potential correction scheme adopted in *GAMER* has two important nice features. First, since the Poisson equation is solved in individual patches with inhomogeneous Dirichlet boundary conditions, the inter-patch communication is minimized, and all patches at a given level can be solved independently. By taking advantage of this feature, we can implement a high-performance GPU Poisson solver, in which hundreds of patches can be solved in parallel. The second and most critical feature is the conservation of mass. Since the pseudomass sheet is defined by the jump of the normal derivative of potential, our scheme enforces the summation of local pseudomass sheet introduced by the interface between a parent patch and a child patch to be zero (to the machine precision).

The feature of mass conservation can easily be seen from the fact that both the coarse-grid and fine-grid potentials satisfy the Poisson equation. Therefore, the Gauss's theorem states that

$$\oint_{\Gamma} \left( \frac{\partial \phi}{\partial n} \Big|_{\epsilon^-} \right) ds = 4\pi G M^c \quad \text{and} \quad \oint_{\Gamma} \left( \frac{\partial \phi}{\partial n} \Big|_{\epsilon^+} \right) ds = 4\pi G M^f, \quad (26)$$

where  $\Gamma$  is the closed boundary surface of a child patch, and  $M^c$  and  $M^f$  are the total coarse-grid mass and fine-grid mass enclosed by  $\Gamma$ , respectively. Now, in *GAMER*, we always have  $M^c = M^f$  thanks to the constrained operation, and hence

$$\oint_{\Gamma} \xi(\mathbf{x}) ds = \frac{1}{4\pi G \Delta h^c} \oint_{\Gamma} \left( \frac{\partial \phi}{\partial n} \Big|_{\epsilon^-} - \frac{\partial \phi}{\partial n} \Big|_{\epsilon^+} \right) ds = 0. \quad (27)$$

Clearly, Equations (26) and (27) still hold in their discrete forms, in which the normal derivatives of potentials are approximated by Equations (22) and (23).

The physical picture for reinforcing the summation of the local pseudomass sheet to vanish is as follows. Beside the truncation error in Equation (18), the numerical error in the fine-grid region is caused by the insufficiently accurate boundary condition, which is obtained by interpolation on the coarse-grid values. In other words, even though the refined region can provide higher resolution, it does not contribute to the coarse-grid solution. The absence of local density distribution in the fine-grid region when evaluating the coarse-grid potential can only provide the coarse-grid resolution, and the numerical error will propagate into the solution of fine-grid potential through the setting of fine-grid boundary condition (even with high-order interpolation). It is the reason why we want to estimate the pseudomass sheet to correct the coarse-grid potential, and thereby provide a more accurate boundary condition for solving the fine-grid potential. However, since in *GAMER* the total mass within a given volume is the same in different refinement levels, there should be *no mass monopole correction* for the coarse-grid potential. Therefore, the total pseudomass introduced by the coarse-fine interface between a parent patch and a child patch is zero.

Finally, the cell-centered gravitational accelerations are evaluated by the three-point finite difference approximation of the gradient operator. The flow variables are advanced by solving Equations (2) and (3), in which the flux terms are ignored. By utilizing the same operator-splitting method described in Section 2.2, the Euler equations with self-gravity can be solved in the order  $xyzGGzyx$  (Trac & Pen 2003), in which the operator  $xyz$  represents the order of directions to update flow variables by the flux differences, and the operator  $G$  represents the updating of flow variables by gravity. Note that the continuity equation (Equation (1)) has no self-gravity term. Therefore, the two successive self-gravity operators  $GG$  can be combined together with a twice larger time step.

In cosmological simulations, the Poisson equation can be rewritten as

$$\tilde{\nabla}^2 \tilde{\phi}(\tilde{\mathbf{x}}) = 4\pi G a [\tilde{\rho}(\tilde{\mathbf{x}}) - \tilde{\rho}_b(\tilde{\mathbf{x}})], \quad (28)$$

where each comoving variable is defined in Equations (15) and (16), and the subscript  $b$  indicates the background density. The gravitational constant  $G$  can be replaced by

$$G = \frac{3H_0^2 \Omega_{m,0}}{8\pi \rho_{b,0}}, \quad (29)$$


where  $\Omega_{m,0}$  is the matter density and the subscript 0 indicates the values at the present time.

#### 2.4. Parallelization

For astrophysical simulations, the spatial resolutions are generally limited by the amount of total memory. It is therefore essential to develop a parallel code that distributes the workload to multiple processors. Accordingly, *GAMER* is developed to work in a multi-CPU system with multi-GPU acceleration. Each CPU manages one GPU, and the data transfer between different CPUs is accomplished by using the MPI library.

Developing a parallel GPU-accelerated program requires elaborate treatments. First of all, even though the computation time may be highly reduced by using GPUs, *the communication*





**Figure 4.** Two-dimensional example of the allocation of buffer patches with two refinement levels. The left panel shows the refinement map within a fixed domain. Each grid represents a patch, and the bold contour represents the boundaries of a sub-domain. The right panel shows the patches actually allocated to the processor in charge of the sub-domain. The solid and dashed lines represent the real patches and buffer patches, respectively. The crosses indicate the patches that do not require to store physical quantities.

*time is not reduced at all.* The network bandwidth may easily become the performance bottleneck, and results in a limited overall performance improvement. Moreover, since the performance of the GPU solver also relies on the massively parallel architecture inside GPU, our parallel algorithm must preserve this capability.

In *GAMER*, the parallelism is based on a rectangular domain decomposition. All patches within a rectangular sub-domain are calculated by one CPU/GPU combination. These patches are referred to as the real patches. Boundary conditions of each sub-domain are provided by allocating the buffer patches surrounding the sub-domain. Figure 4 shows a two-dimensional example of the allocation of buffer patches. The physical data stored in each buffer patch are always filled by transferring data between processors. Note that each buffer patch also stores the correct indices of the parent, child, and 26 sibling patches, and a null value is assigned if the corresponding patch does not exist. Accordingly, all coarse–fine boundaries can be correctly identified even if they coincide with the sub-domain boundaries. Moreover, by doing so, we do not need to store all the oct-tree data structure redundantly in all processors.

To avoid the bottleneck of network communication, the amount of data transfer between different CPUs is carefully minimized in *GAMER*. For the hydrodynamic solver, since the TVD scheme requires three ghost-zone values, only a three-cell-wide array is transferred and stored in each buffer patch. The data stored in the buffer patches at level  $\ell$  are used to set the ghost-zone values for both the patches at level  $\ell$  via direct memory copy and the patches at level  $\ell + 1$  via interpolation. If a coarse–fine interface coincides with the sub-domain boundaries, the corresponding flux data are also transferred for the flux correction operation. For the Poisson solver, since it requires only one ghost-zone value, a two-cell-wide array is transferred for setting the boundary conditions via interpolation, and a one-cell-wide array is transferred for the sibling relaxation step.

We note that not all buffer patches are necessary to be filled with the hydrodynamic and potential data. To be more precise, a buffer patch is required to receive data only if any of its 26 sibling patches corresponds to a real patch. A two-dimensional example is also illustrated in Figure 4, in which the buffer patches marked with a cross are not required to store physical quantities and are used only to provide the correct oct-tree data structure for real patches adjacent to the sub-domain boundaries. Also note that the additional memory overhead associated with the allocation

of buffer patches is generally negligible as long as the number of real patches is sufficiently large.

Adopting the FFTW library as the root-level Poisson solver requires additional works. The parallelization strategy of FFTW is based on the slab decomposition, which is incompatible with the rectangle decomposition adopted by *GAMER*. Therefore, a rectangle-to-slab transformation must be performed before the root-level Poisson solver, and a slab-to-rectangle transformation must be performed after the root-level Poisson solver. Both these two operations require global communications. However, since in general the amount of data in the root level is much less than that in higher levels, this additional communication time is usually negligible.

The parallelization algorithm described above requires a recurrent search for all patches within the sub-domain. It can also be time consuming in a large-scale simulation, especially when the performance of single-patch solvers is highly improved by using GPUs. In *GAMER*, we get around this problem by first constructing a table recording the indices of root-level border patches, which is defined as the real patches adjacent to each side of the sub-domain boundaries. The table recording the indices of “higher-level” border patches can then be constructed hierarchically, as a consequence of the fact that a border patch at level  $\ell$  must be the child patch of a border patch at level  $\ell - 1$ . After that, the table listing the indices of patches to send and to receive data can be built by only searching over the border-patch table. Therefore, the global search is performed only at the root level, at which the number of patches is much smaller than that of higher levels. Also note that these tables are only re-constructed every time after rebuilding the refinement map; afterward they can be reused until the next domain refinement.

As the number of processors increases, applying the rectangular domain decomposition in the AMR implementation can lead to an issue of load unbalance, where different computing nodes have different loads. Generally, this issue is solved by adopting the method of a space-filling curve (e.g., Campbell et al. 2003) to redistribute the computing loads. This is currently being implemented into the *GAMER* code. Note that the concept of allocating the buffer patches should still be adopted, as long as we impose the constraint that each child patch is placed in the same node (more precisely, the same MPI rank) as its parent patch. Preliminary tests have shown that this constraint results in only a minor influence of the load balance. The unbalance is found to be less than 2% in a purely baryonic cosmological simulation with 32 CPU/GPUs.

### 3. GPU IMPLEMENTATION

In *GAMER*, the GPU implementation is inspired by the two parallelism levels naturally embedded inside the AMR structure. First, each mesh patch can be calculated independently as long as its ghost-zone values are provided. Therefore, we can use one thread block to calculate one patch. Second, all cells inside a patch can be calculated in parallel as long as there is a synchronization mechanism to coordinate the data update of each cell. This can be accomplished by using multiple CUDA threads to calculate different cells within the same patch, and store the updated results in the shared memory. In the following, we describe in detail the GPU implementations of different parts in the code, and give the results of timing analyses.

#### 3.1. GPU Hydrodynamic Solver

The GPU hydrodynamic solver implemented in *GAMER* involves three basic steps as follows.

1. Send the input array, which stores mass density, momentum density, and energy density, downstream from the CPU memory to the GPU global memory.
2. Invoke the GPU hydrodynamic kernel to advance the equations for one time step.
3. Send the output array, which stores mass density, momentum density, and energy density, upstream from the GPU global memory back to the CPU memory.

Before invoking the GPU kernel, a preparation step is performed by the CPU to prepare the input array which stores the interior data of each patch and the ghost-zone values. For the TVD scheme, three ghost zones are required on each side in each spatial direction. The ghost-zone values are obtained by direct memory copies if the sibling patches exist. Otherwise, the values are obtained by linear interpolation with the Min-Mod limiter.

Note that eight nearby patches are always allocated simultaneously thanks to the oct-tree data structure. Therefore, to reduce the amount of workload associated with the preparation of the ghost-zone values, we can group these eight patches into a larger array (hereafter referred to as the patch group) before sending into GPU. For the hydrodynamic solver, each patch group contains  $(8 \times 2 + 3 \times 2)^3 = 22^3$  cells, where 8 is the size of a single patch and 3 is the size of the ghost zones. Since in this approach the ghost zones are prepared for the exterior region of each patch group instead of each individual patch, it reduces 63% of the computational overhead associated with the ghost-zone preparation.

After sending the input array into GPU, each patch group is advanced by one GPU thread block. The single-block GPU hydrodynamic solver is based on the TVD scheme described in Section 2.2. The three-dimensional evolution is achieved by using the dimensional-splitting method, in which the solution is obtained by first applying a forward sweep followed by a backward sweep within the same step. During one GPU kernel execution, either a forward sweep or a backward sweep is performed, and a Boolean parameter is sent into GPU to indicate the direction of sweeping.

Since the dimensional-split Euler equations are equivalent to a set of one-dimensional conservation equations, the data of one patch group can be regarded as a set of data columns, and each of which can be evolved independently. Accordingly, for the single-block GPU hydrodynamic solver, we advance the solutions of a fixed number of data columns in parallel. Each thread is responsible for advancing a single cell. We then iterate through the remaining data columns until the whole patch group is updated. Data that need to be accessed by more than one threads within the same data column (e.g., the fluid fluxes) are stored in the GPU shared memory. Otherwise they are stored in the per-thread registers.

Briefly, the GPU hydrodynamic kernel executed by each thread works with the steps as follows.

1. Get the index of cell being calculated. Fetch the corresponding data from the global memory and store in the per-thread registers.
2. Calculate the freezing speed  $c(x, t)$  and construct the corresponding auxiliary vector  $\mathbf{w}$ .
3. Calculate the left-moving and right-moving fluxes (Equation (8)) defined at the boundaries of each cell by the first-order upwind scheme. Store the fluxes in the shared memory.
4. Obtain the midpoint solutions by Equation (13). Store the solutions in the per-thread registers.

5. Recalculate the freezing speed using the midpoint values. Construct the corresponding midpoint auxiliary vector.
6. Calculate the midpoint left-moving and right-moving fluxes defined at the boundaries of each cell by the second-order TVD scheme. Store the fluxes in the shared memory.
7. Obtain the full-step solutions by Equation (14). Store the solutions in the per-thread registers.
8. Store the full-step solutions as well as the fluxes across the patch boundaries back to the global memory.
9. Repeat steps 1–8 for the next targeted data column until the entire patch group is updated.
10. Repeat steps 1–9 for the next one-dimensional sweep until either a forward sweep or a backward sweep is complete.

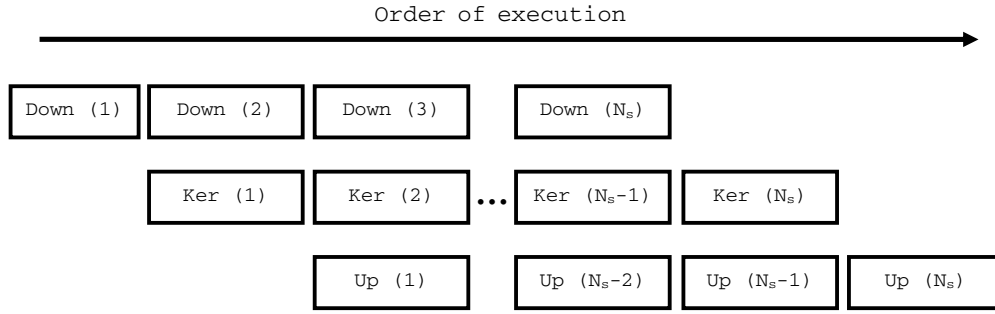
The output array stores the updated solutions of each patch group as well as the fluxes across the boundaries of each patch. To reduce the amount of data transfer, no ghost-zone values are stored in the output array. After the GPU kernel execution, the output array is transferred upstream to the CPU memory, and followed by a closing step performed by the CPU.

The closing step involves two operations. First, it copies the updated solutions back to each corresponding patch pointer. Second, for a patch adjacent to a coarse–fine interface, the data of fluxes across this interface are copied into its own flux array (to be corrected afterward by the fine-grid fluxes) if this patch is in the coarse side of the interface. On the contrary, if this patch is in the fine side of the interface, the data of fluxes are copied into the flux array of the corresponding neighboring coarse patch (to correct the coarse-grid fluxes).

We note that it is unnecessary to simultaneously advance all patch groups in GPU, since the GPU computing power can be fully exploited as long as there are sufficient arithmetic operations. Typically, we advance 128–240 patch groups in GPU in parallel. The input and output arrays are allocated only for the patch groups being updated, and a single array can be reused for different sets of patch groups. The additional memory requirement in CPU for storing the ghost-zone data is therefore nearly negligible. Moreover, the memory requirement in GPU for the hydrodynamic kernel is less than 200 MB, and hence the limited amount of DRAM memory in GPU is not an issue in the current implementation of *GAMER*.

In a practical point of view, the performance comparisons between CPU and GPU must include the data transfer time between the CPU memory and the GPU memory through the PCI Express bus. This data transfer time can be greatly reduced by utilizing the capability of asynchronous memory copies in GPU, in which the memory copies between CPU and GPU can be overlapped with the kernel executions (NVIDIA 2008). In CUDA, the concurrency between different operations is managed by creating the stream objects, which contain a sequence of memory copy operations and kernel invocations. To simplify the discussion, here we assume that a stream identification number (hereafter referred to as the stream ID) is assigned to each stream object. For a memory copy operation and a kernel launch with different stream IDs, they can be performed concurrently. Otherwise, they will be performed sequentially in the order they are declared in the same stream object.

In *GAMER*, since different patch groups at the same refinement level can be evolved in an arbitrary order, we can create several stream objects inside one GPU solver. Each stream object contains one downstream memory copy, one kernel invocation, and one upstream memory copy for a fixed number of patch groups. Accordingly, the data transfer time of the patch



**Figure 5.** Illustration of the concurrent execution between the memory copy and the kernel launch. The abbreviations “Down,” “Up,” and “Ker” stand for the downstream memory copy, upstream memory copy, and kernel execution, respectively. The numbers in the parentheses indicate the stream IDs. The operations at the same column are performed concurrently. Therefore, the memory copies can be overlapped with the kernel executions.

groups belonging to one stream object can be overlapped with the kernel execution time of the patch groups belonging to a different stream object.

As an illustration, let  $N_s$  denote the number of stream objects and  $N_g$  denote the number of patch groups associated with one stream object. The total number of patch groups advanced by one launch of the GPU solver is then given by  $N_s \times N_g$ . At the first step, the patch groups 1 to  $N_g$  are sent to the GPU memory. At the second step, a GPU kernel is executed to advance the patch groups 1 to  $N_g$ , while at the same time the patch groups  $N_g + 1$  to  $2N_g$  are sent to the GPU memory. At the third step, a GPU kernel is executed to advance the patch groups  $N_g + 1$  to  $2N_g$ , while at the same time the updated solutions of the patch groups 1 to  $N_g$  are sent back upstream to the CPU memory, and the prepared data of the patch groups  $2N_g + 1$  to  $3N_g$  are sent downstream to the GPU memory. Figure 5 illustrates the complete procedure.  $N_s + 2$  steps are required to complete one launch of the GPU solver. Theoretically, by using  $N_s$  streams, the data transfer time between CPU and GPU can be reduced to  $(1/N_s)$ th of the time without using streams.

For comparison, a CPU hydrodynamic solver with the same TVD scheme has also been implemented. In order to have more reliable timing measurements, we have measured the performance of the hydrodynamic solver in three different hardware implementations: the GraCCA system, the GPU system installed in the National Center for High-performance Computing of Taiwan (hereafter referred to as NCHC), and the GPU system installed in the Center for Quantum Science and Engineering of National Taiwan University (hereafter referred to as CQSE). Below, we give a brief description of the hardware implementations in different GPU systems.

The GraCCA system contains 18 nodes connected by Gigabit Ethernet. Each node is equipped with two GeForce 8800 GTX GPUs and one AMD Athlon 64 X2 3800 CPU. Two distinct configurations of GPU systems are implemented in NCHC. First, a GPU cluster consisting of 16 nodes connected by InfiniBand is installed in NCHC. In order to exploit the bandwidth of InfiniBand, currently each node is only equipped with two Tesla T10 GPUs and two Intel Xeon X5472 CPUs. In addition, an experimental node with four Tesla T10 GPUs and two Intel Xeon E5520 CPUs is also installed in NCHC. This node aims at exploring the computing power of four Tesla GPUs. Therefore, throughout this work, we perform the timing measurements in NCHC at this node. The CQSE GPU cluster contains 16 nodes connected by Gigabit Ethernet. Each node is equipped with four Tesla T10 GPUs and two Intel Xeon E5462 CPUs.

Note that in contrast to the Tesla T10 GPU, the GeForce 8800 GTX GPU installed in the GraCCA system does not support the

**Table 1**  
Detailed Timing Analysis of the GPU Hydrodynamic Solver

Platform	Downstream ( $\mu$ s)	Upstream ( $\mu$ s)	Kernel ( $\mu$ s)	Total <sup>a</sup> ( $\mu$ s)	Total <sup>b</sup> ( $\mu$ s)	CPU <sup>c</sup> ( $\mu$ s)	Speed-up
GraCCA	71	46	252	369	...	5369	14.55
NCHC	51	64	144	262	177	2366	13.37
CQSE	44	29	150	223	177	2701	15.26

**Notes.** Timing results shown here are the execution times per patch group.

<sup>a</sup> Timing results with only one stream.

<sup>b</sup> Timing results with four streams.

<sup>c</sup> Timing results using CPU only.

capability of asynchronous memory copies. Consequently, the memory copies and the kernel invocations must be performed sequentially, which is equivalent to have only one stream. For the Tesla T10 GPU, we have compared the performances between tests using only one stream ( $N_s = 1$ ) and four streams ( $N_s = 4$ ). Also note that each thread block is calculated by one multiprocessor in GPU, and there are 16 and 30 multiprocessors in GeForce 8800 GTX GPU and Tesla T10 GPU, respectively.

Figure 6 shows the performance speed-up of one GPU over one CPU as a function of the number of patch groups associated with each stream ( $N_g$ ) for the hydrodynamic solver. The performance measurements include the downstream and upstream data transfers as well as the kernel executions. The computing times for the preparation step and the closing step are not included. It can be seen that the performance of GPU solver is linear proportional to  $N_g$  when  $N_g$  is smaller than the total number of multiprocessors in one GPU. As we have at least one patch group for each multiprocessor, the performance approaches the saturated values. Factors of 14.6, 13.4, and 15.3 performance speed-ups are demonstrated in the GraCCA system, NCHC, and CQSE, respectively, for GPU computation as opposed to CPU computation.

A detailed timing analysis for the hydrodynamic solver is listed in Table 1, in which we set  $N_g = 256$  in the GraCCA system and  $N_g = 240$  in NCHC and CQSE. Note that since the specifications of GPUs installed in NCHC and CQSE are the same, the difference of speed-up ratios between these two systems is mainly due to the different performances of CPUs as well as the different Northbridge chips. When using only one stream, the data transfer times are 46.4%, 79.9%, and 48.7% of the kernel execution times in the GraCCA system, NCHC, and CQSE, respectively. A relatively low bandwidth in PCI Express bus is found in NCHC, especially in the upstream bandwidth. Nevertheless, the data transfer times are reduced to 22.9% in NCHC and 18.0% in CQSE when the memory copies






Figure 6. Performances of the GPU hydrodynamic solver in different hardware implementations.

are partially overlapped with the kernel executions by using four streams. Also note that both GPUs and CPUs installed in NCHC and CQSE outperform those installed in the GraCCA system. However, the performance ratios between GPU and CPU measured in different hardware implementations do not vary significantly.

### 3.2. GPU Poisson Solver

The basic procedure of the GPU Poisson solver is similar to that of the GPU hydrodynamic solver. It works with the steps as follows.


1. Send the input array, which stores the mass density and potential, downstream from the CPU memory to the GPU global memory.
2. Invoke the GPU SOR kernel to evaluate the potential solutions.
3. Send the output array, which stores the potential only, upstream from the GPU global memory back to the CPU memory.

In order to reduce the computational overhead associated with the ghost-zone preparation as well as to reduce the interfaces between patches at the same refinement level, eight nearby patches are also grouped into a patch group at the preparation step. For the SOR method, the potential data require one ghost zone on each side in each spatial direction, while the mass density data require no ghost zones. Accordingly, for the GPU Poisson solver, each patch group stores  $(8 \times 2 + 1 \times 2)^3 = 18^3$  potential data and  $(8 \times 2)^3 = 16^3$  density data. After the preparation step, the input array is sent downstream to the GPU global memory, and the GPU SOR kernel is invoked to evaluate the potential solutions of each patch group. Afterward, the output array storing only the potential solutions is sent upstream to the CPU memory, and a closing step is performed by CPU to copy the solutions back to each corresponding patch pointer.

The single-block GPU Poisson solver is based on the SOR scheme with odd-even ordering as described in Section 2.3. Implementing the three-dimensional SOR scheme into GPU differs greatly from the implementation of the three-dimensional TVD scheme. For the GPU hydrodynamic kernel, the data of each patch group are decomposed into a set of data columns, and each of which can be evolved independently. Accordingly, only the data columns being calculated need to be stored in the shared memory, and a single shared memory array can be reused for many different data columns. The total amount of shared memory required in the GPU hydrodynamic kernel is therefore only 8.6 KB.

On the contrary, the SOR scheme requires the three-dimensional relaxation, where the data in each cell must be re-accessed during each iteration. In addition, the arithmetic intensity in each iteration of the SOR scheme is relatively low. Consequently, storing all data naively in the global memory will suffer from the high memory latency and result in marginal performance improvement. However, since the amount of potential data in each patch group is about 22.8 KB, it is impossible to store all potential data in the shared memory (which is only 16 KB per multiprocessor).

One of the solutions to reduce the data transfer is that during each iteration, we solve Equations (19) and (20) slice by slice. Since Equation (19) requires the data of six nearby cells, we can keep only three slices of data in the shared memory, and fetch a new slice into the shared memory each time when iterating to the next slice. It ensures that the reused data are stored in the shared memory during each iteration. However, this method still requires frequent data transfer between the global memory and the shared memory, and hence the achieved performance is far from optimized. On the contrary, in *GAMER* we have implemented a different scheme that minimizes the data transfer by utilizing both the shared memory and the per-thread registers, as detailed below.



**Figure 7.** Two-dimensional illustration of the configuration of odd and even cells in the GPU SOR kernel. To simplify the illustration, a patch group with only  $4 \times 4$  cells is shown. The border of patch group is highlighted by bold lines. The triangles and circles represent the even and odd cells, respectively. The interior cells are indicated by open symbols, and the ghost cells are indicated by filled symbols. At the first half-step for updating all even cells, the triangle data are stored in the per-thread registers, while the circle data are stored in the shared memory. A shared memory array with  $6 \times 3$  elements is allocated, and eight threads are used to solve the equation. Each thread stores the potentials of one interior cell and one ghost cell in its own registers. The dotted arrows indicate the direction of data exchange at the end of the even step.

To solve the issue of shared memory shortage, we first note that there are 8192 and 16,384 threads per multiprocessor in the GeForce 8800 GTX and Tesla T10 GPU, respectively. In principle, they provide another 32 KB and 64 KB storage through the per-thread registers, although the data stored in one register cannot be directly accessed by other registers. Next, by dividing all cells within a single-patch group into odd and even cells, the update of an even cell depends only on nearby odd cells, and vice versa. Accordingly, when updating all even cells at the even step in one iteration, *the data of all even cells can be stored in the per-thread registers* instead of the shared memory. By contrast, since the data of each odd cell need to be accessed several times by nearby even cells, *the data of all odd cells are stored in the shared memory* at the even step. After updating all even cells, we exchange the data stored in the registers and the shared memory, and proceed to the odd step for updating all odd cells. At the end of the odd step, another data exchange operation is applied, and a single iteration is complete.

Figure 7 illustrates the configuration of odd and even cells as well as the data exchange operation at the even step in two dimensions. Note that the ghost zones are also divided into odd and even cells, except that their values are fixed during iterations. In this scheme, the requirement for the amount of shared memory is nearly halved, and all potential data are stored in either the shared memory or the per-thread registers. The data transfers of the potential data between the global memory and the shared memory are only necessary before and after the relaxation loop, and thus the performance is highly improved. Also note that since the access rate of the mass density is much lower than that of the potential, the data of mass density are still stored in the global memory.

According to the discussions above, we store the potential data of each patch group in either the shared memory or the per-

thread registers, and update the potential solutions slice by slice. Since there are 128 interior even cells and 128 interior odd cells in each  $x$ - $y$  plane in one patch group, we use 128 threads for each patch group. Each thread will update the potential of one even cell in each slice (and store in the register) at the even step and update the potential of one odd cell in each slice (and store in the register) at the odd step. Since there are 16 slices of potential data to be updated in each patch group, each thread requires 16 registers for storing the temporary interior potentials. Besides, each thread also requires six registers for storing the ghost-zone potentials on each side of the patch group. The GPU SOR kernel executed by each thread works with the steps as follows.

1. Load the potentials of odd cells into the shared memory.
2. Load the potentials of even cells into the per-thread registers.
3. Initialize the parameters for the even step.
4. Evaluate the residual of an even cell at the targeted  $x$ - $y$  plane by Equation (19). Store the residual in a shared memory array.
5. Update the solution of an even cell by Equation (20).
6. Repeat steps 4–5 for the next  $x$ - $y$  plane until all even cells are updated.
7. Exchange the data stored in the shared memory and the per-thread registers.
8. Initialize the parameters for the odd step. Repeat steps 4–7 for all odd cells.
9. Perform a reduction operation to get the 1-norm of the residuals. Repeat steps 3–9 until the 1-norm of the residuals has been diminished to the floating-point precision limit (compared to the source term in the Poisson equation).
10. Store the solutions back to the global memory.

A CPU Poisson solver with the SOR scheme has been implemented for the sake of comparison. Note that the data exchange operation is not required in the CPU solver. Figure 8 shows the performance speed-up of one GPU over one CPU as a function of the number of patch groups associated with each stream ( $N_g$ ) for the Poisson solver. The performance measurements include the downstream and upstream data transfers as well as the kernel executions, whereas the computing times for the preparation step and the closing step are not included. In the NCHC and CQSE systems, four streams are used for the asynchronous memory copies. For comparison, the results using only one stream are shown as well. As for the GPU hydrodynamic solver, the speed-up ratios are linear proportional to  $N_g$  when  $N_g$  is smaller than the number of multiprocessors in one GPU, and they approach the saturated values when we have at least one patch group for each multiprocessor. Factors of 16.1, 23.4, and 24.1 performance speed-ups are demonstrated in the GraCCA system, NCHC, and CQSE, respectively.

Table 2 shows a detailed timing analysis for the Poisson solver, in which we set  $N_g = 256$  in the GraCCA system and  $N_g = 240$  in NCHC and CQSE. Again, a relatively low bandwidth of the PCI Express bus is measured in NCHC. When using only one stream, the data transfer times are 20.0%, 56.4%, and 33.3% of the kernel execution times in the GraCCA system, NCHC, and CQSE, respectively, and they are reduced to 15.4% in NCHC and 21.4% in CQSE as the asynchronous memory copies are activated by using four streams.

### 3.3. Gravitational Acceleration

In *GAMER*, the evaluation of the potential gradients as well as the updating of hydrodynamic variables by self-gravity is




Figure 8. Performances of the GPU Poisson solver in different hardware implementations.

Table 2  
Detailed Timing Analysis of the GPU Poisson Solver

Platform	Downstream ( $\mu\text{s}$ )	Upstream ( $\mu\text{s}$ )	Kernel ( $\mu\text{s}$ )	Total <sup>a</sup> ( $\mu\text{s}$ )	Total <sup>b</sup> ( $\mu\text{s}$ )	CPU <sup>c</sup> ( $\mu\text{s}$ )	Speed-up
GraCCA	13	8	105	126	...	2034	16.14
NCHC	10	12	39	61	45	1055	23.44
CQSE	8	6	42	56	51	1230	24.12

Notes. Timing results shown here are the execution times per patch group.

<sup>a</sup> Timing results with only one stream.

<sup>b</sup> Timing results with four streams.

<sup>c</sup> Timing results using CPU only.

also performed in GPU. The procedure is analogous to the ones adopted in the hydrodynamic and Poisson solvers. An input array storing both the hydrodynamic variables and the potentials is sent downstream to GPU, a GPU kernel is executed to update the solutions, and the updated hydrodynamic variables are sent upstream back to CPU.

Note that the arithmetic intensity is extremely low in this case, and hence the data transfer time between CPU and GPU dominates the execution time. Consequently, only marginal performance speed-up factors around 1.3–1.4 are achieved in the GraCCA system. In NCHC and CQSE, the performances are even slightly lower than using CPU. Nevertheless, in the cosmological tests using CPU only, the execution time associated with the calculation of gravitational acceleration only accounts for  $\sim 1\%$  of the total execution time. Therefore, the overall performance of the code is not constrained by the relatively poor performance of the computation of gravitational acceleration.

### 3.4. Courant–Friedrichs–Lewy Condition

To ensure the stability of numerical integration, the integration time step must satisfy the Courant–Friedrichs–Lewy (CFL) stability criterion. For the three-dimensional relaxing

TVD scheme, the CFL condition is given by

$$\Delta t \leq \frac{\Delta h_\ell}{c_s + \max(|v_x|, |v_y|, |v_z|)}, \quad (30)$$

where  $c_s$  is the sound speed and  $\Delta h_\ell$  is the zone spacing at level  $\ell$ . The denominator in Equation (30) gives the maximum speed of information propagation.

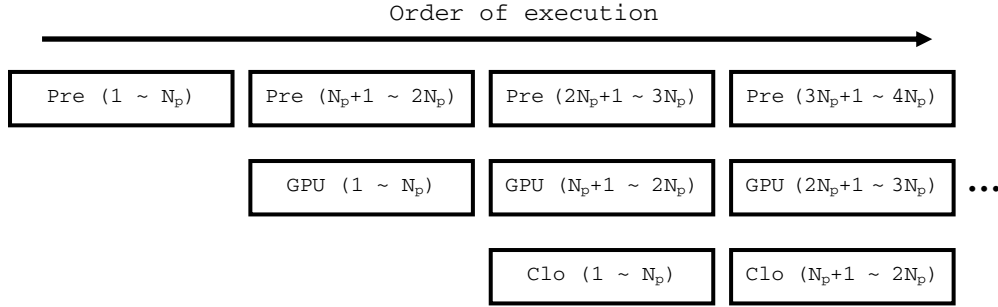
The evaluation of the CFL condition may, at first sight, seem to take negligible time compared to the total execution time. However, timing measurements show that, after the performances of both the hydrodynamic and Poisson solvers are highly improved by using GPUs, the calculation time for the CFL condition accounts for at most 8% of the total execution time if the CFL condition is computed in CPU. It is comparatively low but non-negligible.

Due to the fact that the evaluation of the CFL condition requires no ghost zones, we can calculate the CFL condition in GPU after the backward sweep of the GPU hydrodynamic solver. Since all data essential for the CFL condition already reside in the GPU memory, no extra downstream data transfer is required. To reduce the amount of upstream data transfer for the CFL condition, a reduction operation is performed in GPU in advance to obtain the maximum speed of information propagation among each patch group. By doing so, only a single floating-point value per patch group is required to be transferred from GPU to CPU. The reduction operation among different patch groups is then performed by CPU to evaluate the CFL condition. Accordingly, since the extra data transfer is minimized and the number of arithmetic operations of the CFL condition is much less than that of the hydrodynamic solver, the computation time of the CFL condition is reduced to less than 0.1% of the total.

### 3.5. Concurrent Execution Between CPU and GPU

In *GAMER*, two optimization strategies are adopted to minimize the computational overhead introduced by the preparation





**Figure 9.** Illustration of the concurrent execution between CPU and GPU. The abbreviations “Pre” and “Clo” stand for the preparation step and the closing step, respectively. The numbers in the parentheses indicate the indices of patch groups being calculated. The operations at the same column are performed concurrently. Accordingly, the preparation step and the closing step performed by CPU can be overlapped with the hydrodynamic solver or the Poisson solver performed by GPU.

step and the closing step. First, as described in Sections 3.1 and 3.2, we prepare the ghost-zone data for each patch group instead of each patch. The surface/volume ratio is reduced, and thus the computational overhead is reduced. However, since the number of the interior cells in one patch group is only  $16^3$ , the number of the ghost zones is still comparable to that of the interior cells. For example, for the hydrodynamic solver which requires three ghost zones, the total number of ghost zones for each patch group is 1.6 times more than the number of interior cells. Moreover, although the number of arithmetic operations involved in the preparation step and the closing step is much less than that of the hydrodynamic and Poisson solvers, *these two steps are performed by CPU* in the current scheme. On the contrary, the execution times of both solvers are reduced by at least a factor of 10 by using GPU. Consequently, timing analyses show that the execution times of the preparation step and the closing step are comparable to or even longer than that of the GPU solvers.

To further remove this bottleneck, a second optimization strategy is implemented in *GAMER* by taking advantage of the parallel execution between CPU and GPU. In CUDA, both the memory copy operations and the kernel invocations in GPU are asynchronous, meaning that the program returns from the function call before the requested task is completed. In other words, the CPU and GPU can work concurrently. With this insight, we can overlap the executions of the preparation step and the closing step in CPU with the executions of the GPU solvers.

As an illustration, let  $N_p$  denote the number of patch groups calculated by one invocation of the GPU solver. For using  $N_s$  stream objects, each of which contains the calculations of  $N_g$  patch groups, we have  $N_p = N_s \times N_g$ . At the first step, the patch groups 1 to  $N_p$  are prepared by CPU. At the second step, the patch groups 1 to  $N_p$  are calculated by the GPU solver, while at the same time the patch groups  $N_p+1$  to  $2N_p$  are prepared by CPU. At the third step, the patch groups  $N_p+1$  to  $2N_p$  are calculated by the GPU solver, while at the same time the patch groups  $2N_p+1$  to  $3N_p$  are prepared by CPU, and the solutions of the patch groups 1 to  $N_p$  are copied to the corresponding pointers by CPU in the closing step. This procedure continues until the solutions of all targeted patch groups are obtained, as illustrated in Figure 9.

We note the similarity between the concurrent executions with the downstream memory copy, the kernel launch, and the upstream memory copy in the GPU solver (Figure 5) and the concurrent executions with the preparation step, the GPU solver, and the closing step (Figure 9). In principle, having a larger  $N_s$  can further decrease the communication time in the PCI

Express bus. However, it will also deteriorate the efficiency of concurrent execution between CPU and GPU, because the amount of CPU workload that can be overlapped with GPU computing is reduced. Typically, we set  $N_s = 4$  to balance these two optimization approaches.


In order to test the efficiency of the concurrent execution between CPU and GPU, we execute each GPU solver in *GAMER* and simultaneously perform a matrix summation in CPU. We keep the workload of the GPU solver fixed while varying the size of the matrix, and measure the total execution time in each case. Ideally, the total execution time should be equal to the one that consumes more time.

Figures 10 and 11 show the test results for the GPU hydrodynamic solver and the GPU Poisson solver, respectively. The measurements are conducted in the NCHC system. As expected, in both solvers the total execution times are dominated by the tasks with higher workload, and good concurrent efficiencies are demonstrated. At the crossover point in Figure 10, where the execution times of GPU and CPU are approximately equal, the measured wall-clock times of performing only the GPU hydrodynamic solver, only the matrix summation, and the concurrent execution between CPU and GPU are 20.76 ms, 21.01 ms, and 22.02 ms, respectively. The overhead is 4.8%. At the crossover point in Figure 11, the measured wall-clock times of performing the GPU Poisson solver only, the matrix summation only, and the concurrent execution between CPU and GPU are 5.84 ms, 5.83 ms, and 6.57 ms, respectively. The overhead is 12.5%. For both solvers, the overheads are less than 1.5% when the GPU solvers dominate the execution times.


The performance improvement actually obtained by exploiting the concurrency between CPU and GPU is application dependent, since the execution times of the preparation step and the closing step are related to the structure of domain refinement. We demonstrate the significance of utilizing this feature by comparing the performances in purely baryonic cosmological simulations, which will be described in detail in Section 5.1.

#### 4. ACCURACY TESTS

*GAMER* is designed to achieve both high performance and high accuracy. All GPU kernels implemented in the code follow the same numerical algorithms as their CPU counterparts, except for certain optimized data managements which do not affect the accuracy. In other words, no numerical robustness is sacrificed for performance. In the following, we give the results of various standard tests to demonstrate the accuracy of the code.



**Figure 10.** Efficiency of the concurrent execution between CPU and GPU for the GPU hydrodynamic solver. The x-axis represents the size of the matrix summation performed by CPU.




**Figure 11.** Efficiency of the concurrent execution between CPU and GPU for the GPU Poisson solver. The x-axis represents the size of the matrix summation performed by CPU.

Single precision is assumed unless the floating-point precision is explicitly stated.

4.1. Acoustic Wave Test

As a first elementary test, we simulate the one-dimensional acoustic wave. This test is particularly useful for demonstrating the second-order accuracy of the hydrodynamic solver (Miniati

& Colella 2007a). We construct the initial condition as follows. The uniform background density and pressure are set to  $\rho_0 = 1$  and  $P_0 = 3/5$ , and the ratio of the specific heats is set to  $\gamma = 5/3$ . The sound speed is given by  $c_s = (\gamma P_0/\rho_0)^{1/2} = 1$ . A sinusoidal density perturbation is adopted with a very small amplitude  $\delta\rho/\rho_0 = 10^{-6}$  in order to avoid any nonlinear effect (e.g., the wave steepening). The flow is initially at rest. Double precision is adopted to reduce the round-off errors. To verify the



**Figure 12.** Acoustic wave test. The wave vector is aligned with the diagonal of simulation box. The data points represent the  $L1$  error norm of density as a function of numerical resolution ( $N$ ).

code accuracy in multi-dimensional cases, we choose the wave vector  $\mathbf{K}$  to be parallel to the diagonal of simulation box, i.e.,  $\mathbf{K} = (1/\sqrt{3}, 1/\sqrt{3}, 1/\sqrt{3})$ .

We perform simulations with increasing resolutions:  $N = 64^3$ – $512^3$ , and estimate the  $L1$  error norm of each case. The spatial coordinate is normalized to the grid size at  $N = 512^3$ . In each case, the  $L1$  error norm of density is defined as

$$L1(\rho) = \frac{1}{N} \sum_i |\rho_i - \rho(\mathbf{x}_i)|, \quad (31)$$

where  $\rho_i$  is the numerical solution of  $i$ th cell along the diagonal and  $\rho(\mathbf{x}_i)$  is the corresponding analytical solution. Figure 12 shows the  $L1$  error norm at  $t = 600$ , as a function of spatial resolution. The error converges as  $L1(\rho) \propto N^{-1.9}$ , which is slightly slower than second order. Error distribution shows that error peaks sharply near local extreme values of fluxes, at which the hydrodynamic solver locally reduces to only first-order accurate. It is found that these errors may contaminate the convergence rate of numerical solutions. Nevertheless, the global convergence still approaches second order.

#### 4.2. Shock Tube Test

The shock tube test (Sod 1978) gives a great insight into the capability and accuracy of our hydrodynamic AMR program. Initially, the flow is stationary, with mass density and pressure jumps across a discontinuous interface. Afterward, three characteristic waves with different propagation speeds are excited, namely, the rarefaction wave, the shock wave, and the contact discontinuity. While the hydrodynamic quantities in the Euler equations are continuous across the rarefaction wave, the mass density is discontinuous across the contact discontinuity, and the mass density, flow velocity, and pressure are discontinuous across the shock wave. A correct AMR scheme must fulfill the following criteria. The quantities to the left and the right of the shock front must satisfy the Rankine–Hugoniot shock jump conditions. All these features should be correctly captured by the simulation. Moreover, the domain refinement should be able to

follow the wave propagations closely, the discontinuities should be resolved with a small number of cells, and there should be no spurious oscillations resulting from the domain refinement.

The initial conditions are constructed as follows:

$$\rho_L = 1, \rho_R = 0.125, \quad (32)$$

$$P_L = 1, P_R = 0.1, \quad (33)$$

$$u_L = u_R = 0, \quad (34)$$

where the subscripts  $L$  and  $R$  indicate the states to the “left” and “right” of the initial discontinuous interface, respectively. The ratio of the specific heats  $\gamma$  is set to  $5/3$ . The interface normal is chosen to be parallel to the  $x$ -axis. We set the size of the root level equal to 64, with five refinement levels ( $\ell_{\max} = 5$ ). Accordingly, the effective resolution is 2048. The Dirichlet boundary condition is adopted to ensure no boundary effects to contaminate the wave propagations.

The refinement criterion is based on the density gradient. A cell is flagged for refinement if the relative change of density across the width of a single cell exceeds a given threshold  $C_\rho$ :

$$\frac{\Delta h_\ell}{\rho} \frac{\partial \rho}{\partial x} \geq C_\rho. \quad (35)$$

For the shock tube test, we have  $C_\rho = 0.03$ . The size of the flag buffer ( $N_b$ ) is set to 8.


Figure 13 shows the simulation results at  $t = 110$ , along with the analytical solutions for comparison. The distance is normalized to the grid size at the maximum level. Both the contact discontinuity and the shock wave are refined to the maximum level, whereas the rarefaction wave is refined to  $\ell = 3$ . Only 2.34%–7.81% of the computational domain is refined to the maximum level during the entire simulation, and note that the discontinuous interface is refined to the maximum level from the beginning. The simulation results demonstrate a good resolution at discontinuities. The shock front is well resolved within three to four cells, and the contact discontinuity, which is potentially more diffusive, is resolved within seven to nine cells. We find a good agreement between the numerical solutions and the analytical solutions. The maximum numerical error is found to be about 1%–3% in the rarefaction wave for the mass density, velocity, and pressure. No unphysical oscillations are observed throughout the computational domain, and the flow velocity and pressure remain uniform across the contact discontinuity.

We adopt the individual-time-step scheme for the shock tube test. At  $t = 110$ , it only takes 22 time steps to evolve the patches at the root level, whereas it takes  $22 \times 2^5 = 704$  steps to evolve the patches at the maximum level. A simulation with uniform time step is also performed for comparison, and no obvious differences in profiles are found between these two schemes.

#### 4.3. Sedov–Taylor Blast Wave Test

The Sedov–Taylor blast wave test features a strong spherical shock and a self-similar flow. Initially, the mass density is uniform and the flow is at rest. The explosion is triggered by instantaneously releasing a point-like energy source, which is assumed to be much larger than the background thermal energy, into a homogeneous medium. Afterward, a spherical shock propagates outward from the point of explosion. Since





**Figure 13.** Shock tube test. The top left, top right, bottom left, and bottom right panels show the velocity, mass density, pressure, and refinement level as a function of position at  $t = 110$ , respectively. The initial discontinuous interface is placed at  $x = 0$ . The solid lines depict the analytical solutions. (A color version of this figure is available in the online journal.)

the background thermal energy is negligible compared to the explosion energy, the shock wave is strong. The challenge of the blast wave test lies in solving a spherically symmetric problem on the Cartesian grid. It may also verify the capability of *GAMER* to accurately capture the propagation of strong shock waves.

A self-similar solution to the Sedov–Taylor blast wave is described in detail in Landau & Lifshitz (1987). Let  $\rho_1$  denote the initial background density, and  $E_0$  be the total amount of explosive energy introduced at  $t = 0$ . The explosion center is chosen to be the origin of the coordinate system. Then, from the dimensional analysis, the position and the propagation velocity of the shock front at  $t > 0$  are given by

$$r_s(t) = \beta \left( \frac{E_0 t^2}{\rho_1} \right)^{1/5}, \quad (36)$$

$$v_s(t) = \frac{dr_s}{dt} = \frac{2\beta}{5} \left( \frac{E_0}{\rho_1 t^3} \right)^{1/5}, \quad (37)$$

where  $\beta$  is a constant depending on  $\gamma$ . For  $\gamma = 5/3$ ,  $\beta \approx 1.15$ . The values immediately behind the shock front can be derived by applying the Rankine–Hugoniot shock jump conditions in the strong shock limit, which gives

$$\rho_2 = \left( \frac{\gamma + 1}{\gamma - 1} \right) \rho_1, \quad (38)$$

$$v_2 = \left( \frac{2}{\gamma + 1} \right) v_s, \quad (39)$$

$$P_2 = \left( \frac{2}{\gamma + 1} \right) \rho_1 v_s^2. \quad (40)$$

Finally, by introducing the dimensionless similarity variable

$$\xi = \frac{r}{r_s(t)} = \frac{r}{\beta} \left( \frac{\rho_1}{E_0 t^2} \right)^{1/5}, \quad (41)$$

the Euler equations can be transformed into a set of ordinary differential equations, and the post-shock solutions can be obtained by direct integration.


We simulate the blast wave by initializing the background environment as  $\rho_1=1$ ,  $v_1 = 0$ , and  $E_1 = 10^{-5}$ . The ratio of the specific heats  $\gamma$  is set to  $5/3$ . A total thermal energy  $E_0 = 8 \times 10^5$  is injected into the eight central cells at the finest level. Note that since an averaging operation is performed during the initialization, the explosive energy is also injected into the eight central cells at coarser levels with an energy density one-eighth of that of their parent cells. The size of the root level is set to  $64^3$  and the maximum level is chosen to be  $\ell_{\max} = 5$ , giving  $2048^3$  effective resolution.

We adopt the pressure gradient as the refinement criterion so that the central region can be refined from the beginning. A cell is marked for refinement if the relative change of pressure across the width of a single cell exceeds a given threshold  $C_P$ :

$$\frac{\Delta h_\ell}{P} \frac{\partial P}{\partial x} \geq C_P. \quad (42)$$

For the blast wave test, we set  $C_P = 1.0$ . The size of the flag buffer ( $N_b$ ) is chosen to be 8. Since the background quantities are continuous across the simulation box, the periodic boundary condition is adopted.

The simulation results at  $t = 2000$  are shown in Figure 14, along with the analytical solutions for comparison. The hydrodynamic variables are normalized to the values directly behind the shock front (Equations (38)–(40)), and the dimensionless radius ( $\xi$ ) is adopted for the spatial coordinate. The numerical



**Figure 14.** Sedov–Taylor blast wave test. The top left, top right, bottom left, and bottom right panels show the shell-averaged velocity, mass density, pressure, and refinement level as a function of radius at  $t = 2000$ , respectively. The error bars represent the standard deviations to the average values. The hydrodynamic variables are normalized to the values directly behind the shock front, and the radius is normalized to the position of the shock front. The solid lines depict the analytical solutions. In the right bottom panel, the refinement level is recorded along the  $x$ -axis.

(A color version of this figure is available in the online journal.)

results shown here are the spatial averages over radial bins with thickness equal to the resolution at the finest level. The standard deviations (asphericity) from the shell-averaged values in each bin are represented by error bars. We find a good agreement between the numerical solutions and the analytical solutions. The propagation of the strong shock front is accurately captured, with a compression ratio  $\rho(\xi = 0.998)/\rho_1 = 3.67$ . The shock front is well resolved within three cells. Moreover, despite using the Cartesian geometry to solve a spherically symmetric problem, the standard deviations from the shell-averaged values are small. The maximum deviation is found in the velocity profile at the radius  $\xi = 0.1$ – $0.3$ , where the effective resolution is only  $256^3$ .

The distribution of refinement levels along the  $x$ -axis is also shown in Figure 14 for illustration. Only the region surrounding the shock front is refined to the maximum level, which is 1.38% of the entire simulation box. Due to the proper-nesting constraint, the refinement levels gradually downgrade to  $\ell = 0$  and  $\ell = 2$  in the upstream and downstream regions, respectively. We adopt the individual time step for the blast wave test. At  $t = 2000$ , only 347 time steps are required for the root level, while 11,104 time steps are required for the finest level.

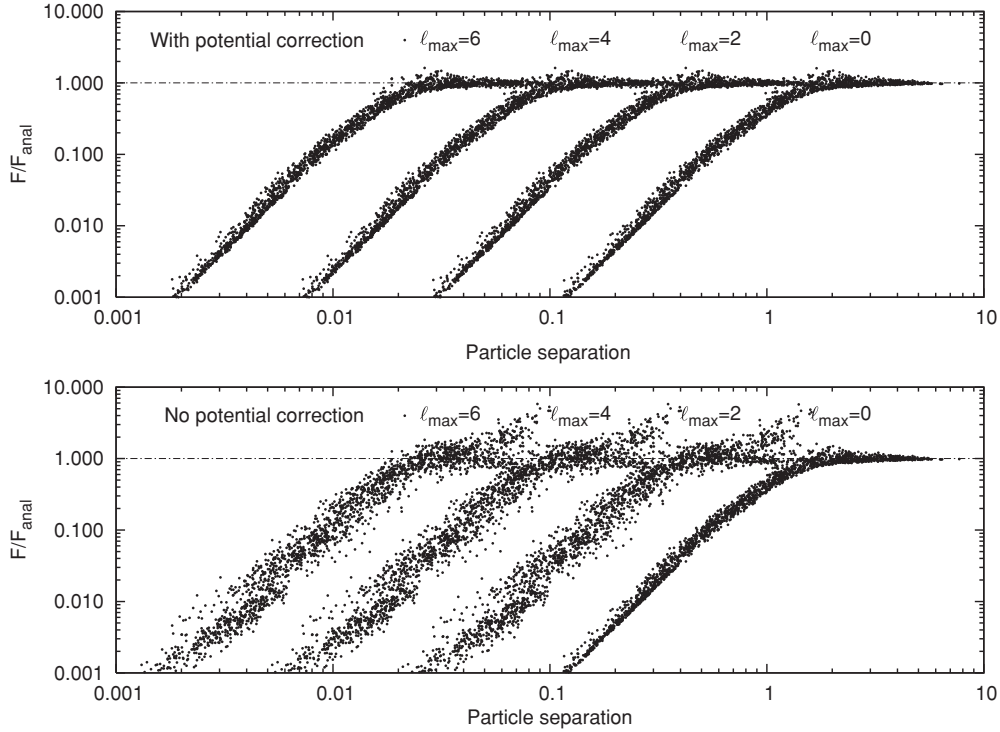
#### 4.4. Point-masses Poisson Solver Test

In order to verify the accuracy of the multi-level Poisson solver described in Section 2.3, we perform tests to measure the gravitational accelerations between point masses. Two unit-mass particles are randomly placed in a simulation box with  $32^3$  root-level cells. The maximum refinement levels are set to  $\ell_{\max} = 0$ – $6$ , giving  $32^3$ – $2048^3$  effective resolutions. The refinement criterion is set to  $\rho_\ell > 0$ , so that the eight cells

surrounding each particle are always refined to the maximum level. At each level, we first compute the mass density using the CIC interpolation scheme (Hockney & Eastwood 1981). Afterward, the gravitational potential and acceleration at each cell are evaluated, and the acceleration of each particle is obtained by the inverse CIC interpolation. Ideally, with an increasing refinement level, the force resolution should be improved and approaches the cell size of the current refinement level.

Figure 15 shows the evaluated gravitational accelerations divided by the analytical solutions. The contributions from image particles due to the periodic boundary condition are calculated using the Ewald summation technique (Hernquist et al. 1991). For comparison, we also show the results without applying any potential correction. These two cases should give identical results when only the root level is used and the multi-level Poisson solver reduces to a pure FFT solver. For  $\ell_{\max} > 0$ , we see that the results without potential correction significantly deviate from the analytical solutions. It reveals the fact that the interpolation errors at coarse–fine boundaries can seriously contaminate the solutions at finer levels. On the contrary, when the potential correction is applied, the force resolution can approach the cell size at the given maximum level. The effective force resolution remains approximately two cells in each case, and the scatters of solutions are highly suppressed. Note that the small solution scatters arise primarily from the fictitious anisotropic force of the square meshes. We also note that the error spatial distributions in the cases  $\ell_{\max} > 0$  are similar to those in the case  $\ell_{\max} = 0$ , indicating that the numerical errors are dominated by the CIC interpolation.

To provide a more quantitative analysis, we further compare in detail the solutions obtained from the hierarchical AMR



**Figure 15.** Poisson solver test. Two particles are randomly distributed in the simulation box, and the gravitational acceleration is evaluated using the multi-level Poisson solver. Here, we show the numerical results divided by the analytical solutions, using 0–6 refinement levels. The particle separation is normalized to the root-level grid size. The top and bottom panels show the results with and without potential correction, respectively.

grids with the uniform-grid solutions, and investigate the numerical errors introduced from the spatial interpolations at the coarse–fine boundaries after the potential correction. One of the approaches to verify the solution accuracy is to fix the particle distributions, compare the numerical solutions to the analytical solutions using different numbers of refinement levels, and estimate the solution convergence rate. However, this approach does not provide a good insight into the convergence behavior for the cases with particle separations close to the minimum cell size, because the errors will be dominated by the CIC interpolation when using less refinement levels. For example, for a particle pair with separation equal to 2 finest cells at  $\ell_{\max} = 6$ , the same particle separation only corresponds to 0.5 finest cell when using  $\ell_{\max} = 4$ , in which the numerical solution deviates significantly from the analytical solution due to the softening effect of the CIC method. Accordingly, it is inappropriate to compare the force solution of this particle pair using refined meshes of different levels. To resolve this issue and quantify the numerical errors for the cases that particle separations are close to the given minimum cell size, we proceed the analysis as follows.

First, for each particle pair, we fix its relative orientation, but normalize the particle separation by the current minimum cell size when adopting different maximum refinement levels. The multi-level Poisson solver is then applied to evaluate the gravitational force between each particle pair, and the forces introduced by image particles are subtracted. By doing so, since the minimum cell size provides the only length scale, the numerical solution obtained using a higher refinement level can easily be rescaled and compared with the root-level solution. Ideally, these two solutions should be identical, if the errors introduced from the coarse–fine boundaries are negligible. For example, for a given particle pair, the gravitational force obtained using  $\ell_{\max} = 1$  should be exactly four times larger than that obtained using  $\ell_{\max} = 0$ .

Figure 16 plots the ratio between the rescaled solutions at refinement levels and the solutions obtained using only the root level. Data are divided into bins that are equally spaced in logarithmic scale, and the standard deviation in each data bin is represented by an error bar. For the case  $\ell_{\max} = 6$ , the average ratios at different data bins range from 96.4% to 99.4%, and the maximum rms value is  $2.89 \times 10^{-2}$ . For particle separations larger than one cell, the average ratios are above 99.0%. Also note that the errors for  $\ell_{\max} = 4$  and  $\ell_{\max} = 6$  are nearly indistinguishable. It verifies that the errors introduced from the coarse–fine boundaries are suppressed after a few levels of refinement, and the accuracy of the point-masses Poisson solver test is always dominated by the CIC-induced errors at the finest level. In other words, it demonstrates that, within a fixed small error, the multi-level Poisson solver is able to give a force resolution equivalent to a uniform-mesh PM solver whose cell size is equal to that in the current maximum refinement level, regardless of the level of refinement.

#### 4.5. Jean’s Instability Test

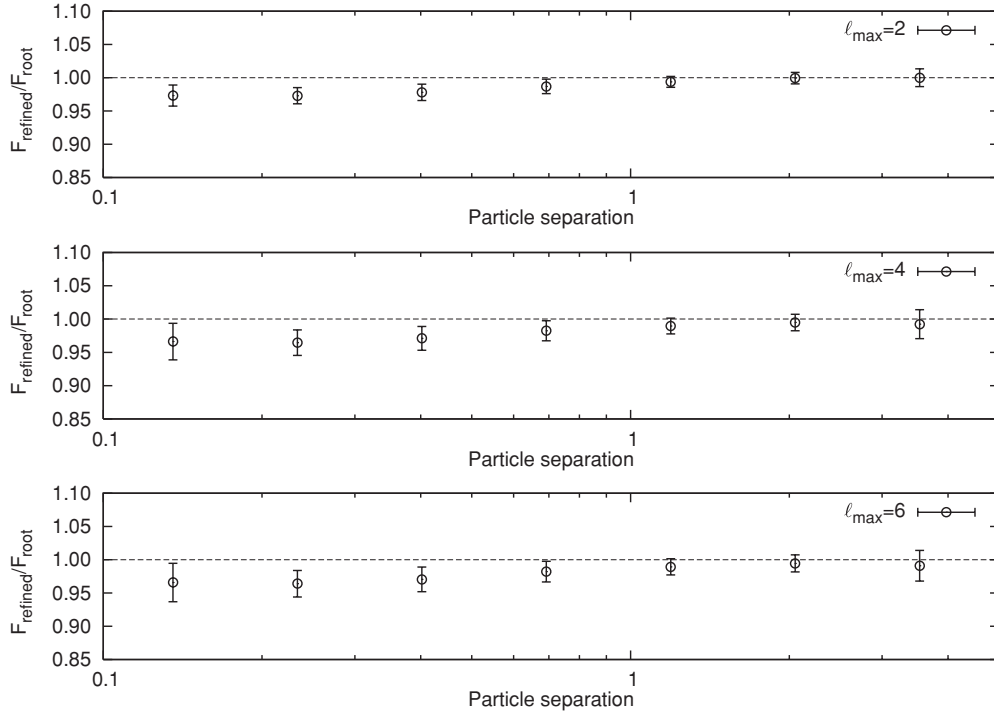
To further test the accuracy of Poisson solver and give a quantitative analysis of the integration scheme in a hydrodynamic + self-gravity system, we test the problem of Jean’s instability. For a small-amplitude perturbation, Equations (1)–(3) and (17) can be linearized, and the dispersion relation is given by

$$\omega^2 = c_s^2[k^2 - k_J^2], \quad (43)$$

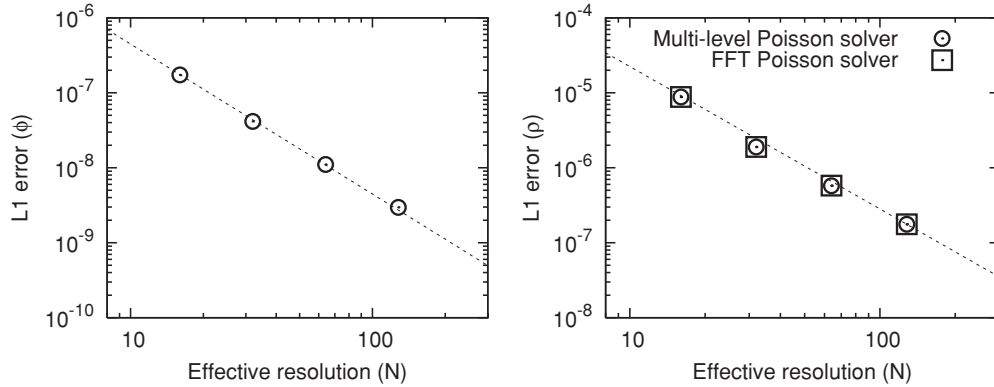
where  $k_J$  is Jean’s wave vector,

$$k_J^2 = \frac{4\pi G\rho_0}{c_s^2}, \quad (44)$$

$c_s = (\gamma P_0/\rho_0)^{1/2}$  is the sound speed, and  $\rho_0$  and  $P_0$  are the background density and pressure, respectively. For  $k <$



**Figure 16.** Poisson solver test. Here we plot the scaled refinement-level solutions divided by the root-level solutions, using 2, 4, and 6 refinement levels. The error bars represent the standard deviations at different data bins. The particle separation is normalized to the minimum grid size in each case. The dashed lines represent the ideal results, in which the scaled refinement-level solutions are equal to the root-level solutions.



**Figure 17.** Jean’s instability test. The left panel shows the  $L1$  error norms of the potential. The right panel shows the  $L1$  error norms of density obtained by the multi-level Poisson solver (circles) and the FFT Poisson solver (squares). Two schemes give nearly identical results.

$k_J$ , a growing-mode solution can be found and the density perturbation grows exponentially as

$$\rho(t) = \rho_0 + \delta\rho \sin(kx)e^{\sqrt{-\omega^2}t}, \quad (45)$$

where  $\delta\rho$  is the amplitude of initial density perturbation.

We simulate the Jean’s instability problem by having  $\delta\rho = 10^{-6}$ ,  $G = 10^{-3}$ ,  $\rho_0 = 1$ ,  $P_0 = 3/5$ , and  $\gamma = 5/3$ . The wavelength is chosen to be equal to the size of the simulation box, which has a  $16^3$  root level with zero to three refinement levels. The effective resolutions are therefore  $N = 16^3$ – $128^3$ . Double precision is adopted to reduce the round-off errors. Since the sinusoidal density distribution has an analytical solution to the Poisson equation, we first compare it to the numerical solutions. Figure 17 shows the  $L1$  error norm of potential. Encouragingly, the second-order convergence is verified. We then perform simulations until the perturbation amplitudes have grown two orders of magnitude, and compare

the numerical results to the analytical predictions. The  $L1$  error norm of density is shown in Figure 17. The error converges as  $L1(\rho) \propto N^{-1.9}$ , which is consistent with the convergence rate observed in the acoustic wave test. For comparison, we also conduct simulations using only the FFT Poisson solver with the same effective resolutions. As shown in Figure 17, the numerical errors observed in these two schemes are nearly identical.

#### 4.6. Spherical Collapse Test

In the spherical collapse test, an initial overdense perturbation with spherical symmetry is placed in the Einstein–de Sitter space ( $\Omega_m = 1$ ). The background pressure is assumed to be negligible. Since the system is gravitational bound, the overdense region first expands with the Hubble flow, and eventually begins to collapse at the “turn around” time. The mass shells at larger radii also collapse at later times, and a strong shock wave propagates outward. Since the turnaround radius provides the



only length scale in this problem, the solution is self-similar. The spherical collapse test is the most comprehensive test for the three-dimensional hydrodynamic code with self-gravity. It simultaneously involves several essential properties of the code, namely, the accuracy of both hydrodynamic and Poisson solvers, the shock-capturing capability, the domain refinement due to the density contrast, and the accuracy of solving a spherically symmetric problem using the Cartesian grid.

An analytical solution to the spherical collapse problem was given by Bertschinger (1985), for both collisional and collisionless components. Here, we only consider the collisional case. Let  $\delta_i$  and  $R_i$  denote the density contrast and the radius of the top-hat overdense perturbation at an initial time  $t_i$ . The turnaround radius, from where the mass shells cease expanding and start to collapse, is given by

$$r_{\text{ta}}(t) = \left( \frac{4}{3\pi} \frac{t}{t_i} \right)^{8/9} \delta_i^{1/3} R_i. \quad (46)$$

Accordingly, the dimensionless radius and the fluid variables are defined as

$$\lambda \equiv \frac{r}{r_{\text{ta}}}, \quad (47)$$

$$\tilde{\rho}(\lambda) \equiv \frac{\rho}{\rho_b}, \quad (48)$$

$$\tilde{v}(\lambda) \equiv v \left( \frac{t}{r_{\text{ta}}} \right), \quad (49)$$

$$\tilde{P}(\lambda) \equiv \left( \frac{P}{\rho_b} \right) \left( \frac{t}{r_{\text{ta}}} \right)^2. \quad (50)$$

From Equations (47)–(50), the Euler equations can be reduced into a set of ordinary differential equations, and the analytical solutions can be calculated by direct integration.

We perform the spherical collapse test by setting  $\delta_i = 0.01$  and  $R_i = 160$  at  $a = 10^{-5}$  in a comoving box, in which the periodic boundary condition is adopted. The ratio of the specific heats  $\gamma$  is chosen to be  $5/3$ , and the background pressure is set to an arbitrarily low value. Since the comoving coordinates are adopted in this test, the initial velocity is set to zero, which corresponds to an unperturbed Hubble flow in the physical coordinates. The size of the root level is set to  $128^3$ , and the maximum refinement level is four, giving  $2048^3$  effective resolution. Note that the radius of the top-hat density distribution is much smaller than the size of the simulation box, so that the image density introduced by the periodic boundary condition should not severely infect the numerical results, and the overdense perturbation can be regarded as an isolated system.

We adopt the amplitude of density as the only refinement criterion. Any cell at level  $\ell$  is marked for refinement if it exceeds the density threshold:  $\rho_\ell > 8^\ell$ . Accordingly, the overdense region is refined to  $\ell = 1$  at the beginning. The size of the flag buffer ( $N_b$ ) is chosen to be 8.

The simulation results at  $a = 0.09$  is shown in Figure 18. We plot the shell averages of the dimensionless variables defined in Equations (47)–(50). The analytical solutions derived by Bertschinger (1985) are also depicted for comparison. In order to test the convergence of the numerical solutions, we perform simulations with zero, two, and four refinement levels, giving  $128^3$ ,  $512^3$ , and  $2048^3$  effective resolutions, respectively. A good agreement between the numerical and analytical solutions is

found. The simulation run with a higher maximum refinement level indeed probes the solutions in a more central region, where only two to four cells in the most central region are incapable of following the rising density and pressure profiles predicted by the analytical solutions. This is due to the second-order TVD scheme that requires three-point interpolation and has the tendency to smooth out the local extreme values. The propagation and the jump conditions of the strong spherical shock are accurately captured.

In order to test the property of spherical symmetry, we also show the standard deviations from the shell-averaged values for the simulation run with four refinement levels. No significant deviations are observed in both the density and the pressure profiles, whereas a relatively large deviation is found in the velocity profile in the post-shock region where the infall velocity abruptly drops to zero. It is consistent with the result given by Teyssier (2002). Nevertheless, the shell-averaged velocity profile agrees well with the analytical solution even in the post-shock region.

The result of domain refinement along the  $x$ -axis for the simulation run with four refinement levels is also shown in Figure 18. Note that neither the density gradient nor the pressure gradient is adopted as the refinement criteria in this test; therefore, the shock front is not necessarily refined to the maximum level. At  $a = 0.09$ , only 0.01% of the region is refined to the maximum level. Due to the self-similar property, the volume of the overdense region increases as the shock propagates outward, and hence the refined region increases as well. However, still only 0.02% of the region is refined to the maximum level at  $a = 0.3$ .


## 5. PERFORMANCE TESTS

The timing measurements described in Section 3 mainly focus on the performances of the hydrodynamic and Poisson solvers. Although there is no doubt that these two solvers are the most time-consuming parts in *GAMER*, it is still not clear whether other parts of the code will become the performance bottlenecks, especially when the computation times of both solvers are highly reduced by using GPUs. For example, the ghost-zone preparation, the domain refinement, and the network bandwidth may also affect the performance. Accordingly, in this section, we perform detailed timing analyses for different parts of the code as well as the overall performance.

The accuracy tests described in Section 4 are inadequate for a persuasive timing analysis, since the profiles of these solutions are too simple compared to realistic astrophysical simulations. Therefore, to have a more comprehensive timing analysis, we measure the performance in purely baryonic cosmological simulations. The initial condition is constructed by using CMBFAST (Seljak & Zaldarriaga 1996) on  $256^3$  grids in a  $\Lambda$ CDM universe at redshift  $z = 99$ . The cosmological parameters are chosen to be  $\Omega_m = 0.3$  and  $\Omega_\Lambda = 0.7$ , and the size of the periodic comoving box is set to  $100 h^{-1}$  Mpc. In the following, we first describe the performance of *GAMER* in a single-GPU system, and then follow it with the multi-GPU performance.

### 5.1. Single-GPU Performance

To fit into the CPU memory of a single node, we set the size of the root level equal to  $128^3$ . The initial condition is obtained by downgrading the  $256^3$  initial condition in order to have consistent results among runs with different spatial



**Figure 18.** Spherical collapse test. The top left, top right, bottom left, and bottom right panels show the shell-averaged dimensionless velocity, mass density, pressure, and refinement level as a function of the dimensionless radius at  $a = 0.09$ , respectively. The circles, triangles, and squares show the numerical solutions using zero, two, and four refinement levels, respectively. The error bars represent the standard deviations to the shell-averaged values for the run using four refinement levels. The solid lines depict the analytical solutions. In the right bottom panel, the refinement level is recorded along the  $x$ -axis.

(A color version of this figure is available in the online journal.)

resolutions. Five refinement levels are adopted to give  $4096^3$  effective resolution. The corresponding spatial resolution is therefore  $24 h^{-1}$  Kpc. We rebuild the refinement map every four steps, which provides the balance between performance and accuracy. Accordingly, the flag buffer size is set to  $N_b = 4$ , so as to prevent the information in the flagged cells from propagating out of the refined regions before the refinement map is rebuilt.


The refinement criterion is similar to the one adopted in the spherical collapse test described in Section 4.6. A cell is labeled as refinable if its local density exceeds a level-dependent threshold:  $\rho_\ell \geq 8^{\ell+1} \rho_b$ . This kind of “quasi-Lagrangian” refinement strategy is often adopted in the AMR cosmological simulations. It roughly fixes the total mass within each cell at different refinement levels. Moreover, when the dark matter particles are included and the standard PM method is used to calculate the gravitational potential, this refinement strategy naturally provides an adaptive soften length according to the level-dependent grid size, thereby minimizing the effect of two-body relaxation.

Figure 19 shows the speed-up ratio of one GPU over one CPU as a function of the redshift. We measure the performances at six different redshifts:  $z = 9.45, 4.12, 1.96, 0.88, 0.21, 0$ . To provide more robust results, we perform the timing analyses in the GraCCA system, NCHC, and CQSE, respectively. In each system, we further compare the performances with and without activating the concurrent execution between CPU and GPU described in Section 3.5. It shows that, as the concurrent execution is enabled, an order of magnitude performance improvements are demonstrated in all three different hardware implementations. The maximum sustained speed-up ratios are 12.19, 10.78, and 10.07 in the GraCCA system, NCHC, and CQSE, respec-

tively, and the speed-up ratios are approximately constant when  $z \leq 4.12$ . It demonstrates the significant performance improvement by using GPU.

We note that, at  $z = 9.45$ , the performance speed-up ratios drop about 40% compared to the maximum sustained performance in each GPU system. This is not surprising since at that time massive halos are not yet formed, and thus the total number of patches at refined levels is relatively small. Consequently, the calculation time of the Poisson solver is dominated by the root-level FFT, which is computed by CPU. Nevertheless, factors of 8.05, 6.73, and 6.95 performance speed-ups are still achieved in the three systems at  $z = 9.45$ .

In Figure 19, we see considerable performance deterioration as the concurrent execution is disabled. The maximum sustained speed-up ratios drop to 7.83, 6.69, and 7.43 in the GraCCA system, NCHC, and CQSE, respectively. It reveals the large computational overheads associated with the preparation steps and the closing steps performed by CPU, for both the hydrodynamic and Poisson solvers. This result is reasonable since in the cosmological simulation, structures form hierarchically, and the number of grids used to resolve a substructure at a given refinement level is generally on the order of  $10^3$ . In GAMER, it approximately corresponds to the size of one patch group. Consequently, nearly all patch groups require spatial interpolations to calculate the ghost-zone values, and it results in a computation time comparable to or even longer than the execution time of each GPU solver. Therefore, it is necessary to simultaneously utilize both the computational power of CPU and GPU, thereby hiding the computation times of the preparation steps and the closing steps by the executions of the GPU solvers. For the cases where the CPU workload is higher, the execution time in GPU



**Figure 19.** Performance speed-up ratios as a function of redshift in different hardware implementations. We measure the performance in purely baryonic cosmological simulations. The speed-up ratios are measured by comparing the runs with single-GPU acceleration to the runs using single CPU only. The abbreviations “async” and “sync” represent the timing results with and without the concurrent execution between CPU and GPU, respectively. The sizes of the root levels are set to  $128^3$  and the maximum refinement levels are set to 5 in all tests.

**Table 3**  
Detailed Timing Analysis of the Performance of *GAMER*

Platform	Hydro. <sup>a</sup> (s)	Poisson <sup>b</sup> (s)	Acc. <sup>c</sup> (s)	Time Step (s)	Hydro. Buffer <sup>d</sup> (s)	Poisson Buffer <sup>e</sup> (s)	Refinement (s)	Total (s)
GraCCA (GPU-async)	11.14	12.82	1.95	0.01	0.25	0.34	0.36	26.88
GraCCA (GPU-sync)	18.54	20.07	2.58	0.01	0.26	0.34	0.40	42.21
GraCCA (CPU)	131.51	191.01	2.81	2.26	0.26	0.34	0.40	328.60
NCHC (GPU-async)	6.58	6.51	1.35	0.00	0.16	0.28	0.27	15.20
NCHC (GPU-sync)	10.49	10.86	2.08	0.00	0.15	0.27	0.28	24.17
NCHC (CPU)	58.95	101.16	1.25	0.99	0.16	0.27	0.27	163.09
CQSE (GPU-async)	10.70	9.74	1.86	0.00	0.24	0.30	0.35	23.19
CQSE (GPU-sync)	15.59	14.51	2.42	0.00	0.24	0.29	0.36	33.41
CQSE (CPU)	96.88	131.49	1.52	1.34	0.24	0.29	0.34	232.10

**Notes.** Timing results shown here are measured in purely baryonic cosmological simulations at  $z = 0.88$ . The abbreviations “async” and “sync” represent the runs with and without the concurrent execution between CPU and GPU, respectively. The computation times of the refinement operation have been divided by 4, since in these runs we reconstruct the refinement map every 4 steps. Also note that four sibling relaxation steps are applied in the end of the Poisson solver in these timing tests.

<sup>a</sup> Hydrodynamic solver (including the preparation and closing steps).

<sup>b</sup> Poisson solver (including the preparation and closing steps).

<sup>c</sup> Gravitational acceleration.


<sup>d</sup> Preparing the buffer patches for the hydrodynamic solver.

<sup>e</sup> Preparing the buffer patches for the Poisson solver.

can also be hidden behind the CPU computation. The considerable performance improvements as the concurrent execution is activated also verify the good concurrent efficiencies as shown in Figures 10 and 11.

Table 3 shows the timing results of different parts in *GAMER* at  $z = 0.88$ . We perform the timing measurements in three different hardware implementations, each of which includes the results using GPU with and without the concurrent execution, and the results using CPU only. First we note that, even when the concurrent execution is enabled, the speed-up ratios of both the hydrodynamic and Poisson solvers are still lower than the

results given in Tables 1 and 2. It implies that the computational overheads associated with the preparation step and the closing step of both solvers dominate the computation time. In other words, these computational overheads are the performance bottlenecks in the current version of *GAMER*, and therefore any further optimization of the GPU solvers alone will not improve the overall performance at all. However, we must emphasize that the three-dimensional relaxing TVD scheme described in Section 2.2 has a relatively low arithmetic intensity, compared to other high-order shock-capturing schemes, e.g., the third-order PPM method. Therefore, it is promising to further improve



**Figure 20.** Domain refinement ratios as a function of refinement levels in a purely baryonic cosmological simulation. We show the results at six different redshifts. The lines of the refinement ratios at different redshifts are slightly shifted in the horizontal direction in order to be visually distinguishable.

(A color version of this figure is available in the online journal.)

both accuracy and performance by implementing an alternative higher-order hydrodynamic scheme in the code.

In Table 3, we see that the overall performance is still dominated by the two solvers. The computation times of the gravitational acceleration, the data copies for the buffer patches, and the domain refinement are about 8%, 2%, and 1% of the total simulation time. The calculation time of the CFL condition is negligible. Note that in the CQSE and NCHC systems, implementing the evaluation of the gravitational acceleration into GPU turns out to provide lower performances than using CPU. It reveals the trickiness of using GPU, where poor performance may be obtained if the arithmetic intensity of the GPU kernel is not high enough and hence the performance degrades due to the additional communication time in the PCI Express bus. Also note that here the preparation of the buffer-patch data do not involve any network communication. Finally, since in these tests we reconstruct the domain refinement in every four steps, the timing results of the refinement operations recorded here are the average values in order to compare with the computation times of other operations.

Figure 20 shows the refinement ratios at different redshifts as a function of the refinement level. The refinement ratio is defined as the ratio of the volume of space refined to a given level to the volume of the entire simulation box. At  $z = 9.45$ , only 10.86% and 0.01% of the simulation box are refined to the level one and two, respectively. It verifies that the performance drop at  $z = 9.45$  shown in Figure 19 is due to the lack of patches at refined levels. In comparison, when there are a sufficient number of patches at higher levels after  $z \leq 4.12$ , the performance approaches the maximum sustained speed-up ratio. Moreover, note that the total computation time of the simulation is dominated by the evolution at lower redshifts, when lots of substructures formed and, correspondingly, the total number of patches at refined levels is high enough for exploiting the


computational power of GPU. The timing experiment shows that the accumulated computation time before  $z > 4.12$  is less than 0.03% of the total computation time over  $0 \leq z \leq 99$ . Therefore, the overall performance improvement during the entire simulation does not suffer from the relatively low performance at higher redshifts. Totally 2608 steps are required to reach  $z = 0$ , and it took only 16 hr of wall-clock time by using only one GPU and one CPU of the GraCCA system.

## 5.2. Multi-GPU Performance

Utilizing multiple GPUs across different nodes requires the network communication. Since the computation time is highly reduced by an order of magnitude as demonstrated in the previous subsection, correspondingly, the communication time becomes more critical. The data transfer time may easily become the performance bottleneck if it is not properly optimized. Therefore, in this section, we investigate the importance of network communication in *GAMER*.

To test the multi-GPU performance, we perform purely baryonic cosmological simulations with the same initial condition adopted in the single-GPU case, and measure the performance using 1–32 GPUs ( $N_{\text{GPU}} = 1\text{--}32$ ) in the GraCCA system. In order to have consistent results, the simulation parameters and the refinement criterion for the multi-GPU tests are the same as the ones adopted in the single-GPU test, except that the sizes of the root levels are set to  $128^3$  and  $256^3$  for the runs with  $N_{\text{GPU}} = 1\text{--}8$  and  $16\text{--}32$ , respectively. The maximum refinement levels are chosen to be  $\ell_{\text{max}} = 5$  in both cases, giving  $4096^3$  and  $8192^3$  effective resolutions. Accordingly, the highest comoving spatial resolution is  $12 h^{-1}$  Kpc. As an illustration, Figure 21 shows a two-dimensional snapshot of the simulation results at  $z = 0$ , in which we plot the density distribution and the corresponding refinement map. Also note that for  $N_{\text{GPU}} = 1\text{--}16$ , we





**Figure 21.** Two-dimensional slice of the gas density distribution in the logarithm scale. The image is obtained in a purely baryonic cosmological simulation at  $z = 0$ , and the image size is  $18.75 \times 18.75 h^{-1}$  Mpc in a  $(100 h^{-1} \text{ Mpc})^3$  simulation box. The refinement map is also shown for illustration, in which each grid represents a single patch. In this snapshot, the densest region is refined to  $\ell = 4$ , giving  $4096^3$  effective resolution.


(A color version of this figure is available in the online journal.)

use only one GPU per computing node, whereas for  $N_{\text{GPU}} = 32$  we use two GPUs per node in the GraCCA system. Therefore, in all timing tests (except for the case  $N_{\text{GPU}} = 1$ ) the network communications are involved.

Figure 22 shows the performance speed-up ratio versus redshift for  $N_{\text{GPU}} = 1$ –32. The ratios are measured by using the same number of CPUs and GPUs in each test. For example, the speed-up ratio for  $N_{\text{GPU}} = 16$  is measured by comparing to the

timing result using 16 CPUs. The concurrent execution between CPU and GPU is activated in all runs. As expected, for the test runs with the same size of the root level and the same maximum refinement level (and hence the computational workloads are the same), the speed-up factors slightly decrease with the number of GPUs. It is reasonable since the network communication accounts for a larger percentage of total simulation time when using GPUs. Nevertheless, the performance enhancement still exceeds 10.47 after  $z \leq 4.12$  when using 16 GPUs. This result is encouraging since it shows that the total simulation time is still dominated by the computation even with the use of GPUs. Moreover, note that in the GraCCA system different nodes are connected by Gigabit Ethernet, which has a relatively low bandwidth compared to other high-bandwidth interconnections, for example, the InfiniBand and Myrinet. Timing analyses show that the network communication accounts for 8%–11% of the total simulation time for  $N_{\text{GPU}} = 16$ .

We also note that a relatively low speed-up ratio is observed in the case  $N_{\text{GPU}} = 32$ , in which two GPUs in the same computing node are both activated. This performance drop is mainly caused by the decrease of data transfer bandwidth between one CPU and one GPU. In the GraCCA system, although there are two individual PCI Express buses dedicated for the two GPUs, the total bandwidth actually measured is only marginally above the bandwidth observed in the single-GPU case (Schive et al. 2008). In addition, the GeForce 8800 GTX GPUs installed in the GraCCA system do not support the capability of asynchronous memory copies, and hence the data transfer time in the PCI Express bus cannot be overlapped with the GPU kernel executions. However, 8.14–9.24 speed-up ratios are still demonstrated during  $0 \leq z \leq 4.12$  in the  $N_{\text{GPU}} = 32$  case. We note that the difficulty with communications via PCI Express bus can potentially be lifted by using the latest motherboard, which supports dual PCI Express  $2.0 \times 16$ .



**Figure 22.** Performance speed-up ratios as a function of redshift using 1–32 GPU(s) in the GraCCA system. We measure the performance in purely baryonic cosmological simulations. The speed-up ratios are measured by comparing the runs with GPU(s) acceleration to the runs using CPU(s) only. The concurrent execution between CPU and GPU is activated in all tests. The sizes of the root levels are set to  $128^3$  and  $256^3$  for the runs using 1–8 and 16–32 GPU(s), respectively, and the maximum refinement levels are set to 5 in all tests.

## 6. CONCLUSIONS AND FUTURE WORKS

We have presented a novel GPU-accelerated AMR code named *GAMER*, which is dedicated to astrophysical simulations. The AMR implementation is based on constructing a hierarchy of mesh patches with an oct-tree data structure, in which each patch is restricted to contain a fixed number of cells. The hydrodynamic solver is based on a three-dimensional relaxing TVD scheme, which is second-order accurate in both time and space. The gravitational potential is solved by using a multi-level Poisson solver, and the SOR algorithm is adopted to solve the Poisson equation for each mesh patch. The potential error caused by the coarse–fine boundaries is diminished by eliminating the effect of pseudomass sheets, which are introduced from the discontinuity of resolution across different refinement levels. The computational performance of *GAMER* has been highly improved by an order of magnitude by utilizing the GPU computing power in several parts of the code.

*GAMER* is a parallel code that can be run in a GPU cluster system. The parallelization strategy is based on a rectangular domain decomposition and the concept of using the buffer patches to enclose the computational sub-domain of each CPU/GPU. This method ensures that different patches can be calculated independently and in an arbitrary order, and therefore multiple patches can be computed by GPUs simultaneously. The amount of data transfer has been carefully minimized to avoid the bottleneck of network communication, and a hierarchical search algorithm has been adopted so that the global search for patches adjacent to the boundaries of sub-domain is only necessary at the root level. Currently, a more advanced domain decomposition method using the space-filling curve is being implemented in order to further improve the load balance.

Two GPU kernels have been developed for speeding up the hydrodynamic and Poisson solvers, respectively. For the GPU hydrodynamic solver, the data of each patch group are decomposed into a set of data columns, each of which can be stored in the low-latency shared memory and can be calculated efficiently and simultaneously by GPU. For the GPU Poisson solver, we have developed an elaborate data-exchange algorithm, in which the data of odd and even cells are stored in either the shared memory or the per-thread registers. No potential data transfer between the global memory and the shared memory is required during the SOR iterations, and hence the computational performance is highly improved. In both GPU solvers, in order to reduce the data transfer time in the PCI Express bus, we have utilized the capability of asynchronous memory copies, by which the data transfer time is overlapped with the GPU kernel executions. To reduce the computational overhead associated with the preparation of the ghost-zone data, the eight nearby patches are always grouped into a patch group before sending into GPU. Furthermore, we have exploited the concurrent execution between CPU and GPU, by which the computation times in CPU and GPU can be overlapped with each other.

We have measured the performances of individual GPU solvers in different hardware implementations, including the GraCCA system and the GPU-equipped nodes in CQSE and NCHC. In each performance test, we have compared the performance using GPU acceleration as opposed to the performance using CPU only. Maximum speed-up factors of 15.3 and 24.1 are demonstrated for the hydrodynamic and Poisson solvers, respectively. We have also measured the efficiency of the concurrent execution between CPU and GPU, and the maximum

overheads are found to be less than 4.8% for the hydrodynamic solver and 12.5% for the Poisson solver.

The accuracy of *GAMER* has been verified by performing several standard test problems, including the shock tube test, Sedov–Taylor blast wave test, and spherical collapse test. The agreement with analytical solutions, the good shock-capturing capability, and the convergence of numerical solutions have been demonstrated. The aspherical scatter is found to be small when solving spherically symmetric problems on the Cartesian grid. We have also shown the significance of the potential correction by comparing the potential solutions obtained by the adaptive-mesh method to those one obtained by the high-resolution uniform-mesh method.

We have measured the performance of the complete *GAMER* code in purely baryonic cosmological simulations, in which we have adopted effective resolutions  $4096^3$  and  $8192^3$  when using 1–8 and 16–32 GPU(s), respectively. An order of magnitude performance speed-ups have been observed when using 1–16 GPU(s), which demonstrates the remarkable performance of the code. A relatively low speed-up factor (still exceeding 8) has been observed when using 32 GPUs in the GraCCA system, which is due to the insufficient bandwidth between CPU and GPU when using two GPUs in the same node. This bottleneck should, however, be able to get highly reduced by using the latest motherboard supporting dual PCI Express 2.0  $\times 16$  and by utilizing the capability of asynchronous memory copies enabled in the latest GPUs. In these tests, we have also demonstrated the importance of exploiting the concurrent execution between CPU and GPU, by which speed-up factors of 1.4–1.6 have been measured when compared to the results without the CPU and GPU overlap.

In the present implementation, the performance bottleneck of *GAMER* lies in the computing time of the ghost-zone data for each solver. It is because this calculation is performed by CPU and also because the number of cells in the ghost zone is comparable to that of the interior cells in each patch group. Adopting a higher-order and more arithmetic-intensive algorithm, for example, the third-order PPM scheme, can almost certainly eliminate this bottleneck. Also note that, in the current implementation of *GAMER*, the computation time of the Poisson solver is approximately equal to that of the hydrodynamic solver. This is due to the fact that we apply 4–8 SOR operations in one step in order to improve the accuracy of gravitational potential. Clearly, investigating a more precise Poisson solver with higher arithmetic intensity can lead to a further performance improvement.

Another approach to further improve performance is to implement more elements of the code in GPU, for example, the interpolation and the oct-tree data structure. However, it will certainly lead to substantial inflexibility of the code. The current GPU implementation in *GAMER* aims at balancing the performance and the flexibility. The GPU kernels are only applied to individual solvers, and the main AMR data structure is still controlled by CPU. Therefore, by simply modifying the GPU kernels, *GAMER* can be applied to different applications and include various physics straightforwardly.

We note that the *GAMER* code can serve as an extremely high-performance and general-purpose tool for astrophysical simulations. Although at present only hydrodynamics and self-gravity are included, we believe that we have developed the framework for exploiting the enormous GPU computing power with an astrophysical AMR code. Future works in *GAMER* will focus on including more physics for various simulations. For

example, one of the most straightforward extensions is to include the dark matter particles and calculate the corresponding density by the standard CIC method. Another promising extension is to include the magnetohydrodynamics, which has been successfully implemented in the AMR framework in several works (e.g., Ziegler 2005; Fromang et al. 2006; Collins et al. 2009). We can further include several physics in *GAMER*, for example the gas cooling and the feedback mechanism, in order to simulate the galaxy formation. Currently, *GAMER* has been modified to address the detailed halo profile found in the large-scale structure simulation with an extremely light dark matter model (Woo & Chiueh 2009).

We thank Shing-Kwong Wong for stimulating discussions. Portions of the simulations in this work were performed in the National Center for High-Performance Computing of Taiwan and the Center for Quantum Science and Engineering of National Taiwan University. We are grateful to Jyh-Shyong Ho and Ting-Wai Chiu for helping conduct simulations in these two institutions. We also thank the referee for constructive comments that greatly improve the work. This work is supported in part by the National Science Council of Taiwan under the grants NSC97-2628-M-002-008-MY3 and NSC98-2119-M-002-001, and also by NTU-CQSE (Nos. 98R066-65 and 98R0066-69).

#### REFERENCES

- Agertz, O., et al. 2007, *MNRAS*, **380**, 963
- Aubert, D., Amini, M., & David, R. 2009, in ICCS Conf. Proc., Part I, LNCS 5544, ed. G. Allen, J. Nabrzyski, E. Seidel, G. D. van Albada, J. Dongarra, & P. M. A. Sloot (Heidelberg: Springer), 874
- Bagla, J. S. 2002, *JA&A*, **23**, 185
- Barnes, J., & Hut, P. 1986, *Nature*, **324**, 446
- Belleman, R. G., Bédorf, J., & Portegies Zwart, S. F. 2008, *New Astron.*, **13**, 103
- Berger, M. J., & Colella, P. 1989, *J. Comput. Phys.*, **82**, 64
- Berger, M. J., & Olinger, J. 1984, *J. Comput. Phys.*, **53**, 484
- Bertschinger, E. 1985, *ApJS*, **58**, 39
- Bryan, G. L., & Norman, M. L. 1997, in ASP Conf. Ser. 123, Computational Astrophysics, ed. D. A. Clarke & M. J. West (San Francisco, CA: ASP), 363
- Campbell, P. M., Devine, K. D., Flaherty, J. E., Gervasio, L. G., & Teresco, J. D. 2003, Williams College Department of Computer Science Tech. Rep. CS-03-01
- Colella, P., & Woodward, P. 1984, *J. Comput. Phys.*, **54**, 174
- Collins, D. C., Xu, H., Norman, M. L., Li, H., & Li, S. 2009, *ApJS*, submitted (arXiv:0902.2594)
- Couchman, H. M. P. 1991, *ApJ*, **368**, L23
- Dubinski, J., Kim, J., Park, C., & Humble, R. 2004, *New Astron.*, **9**, 111
- Efstathiou, G., Davis, M., Frenk, C. S., & White, S. D. M. 1985, *ApJS*, **57**, 241
- Frigo, M., & Johnson, S. G. 1998, in ICASSP Conf. Proc. Vol. 3 (Washington, DC: IEEE), 1381
- Fromang, S., Hennebelle, P., & Teyssier, R. 2006, *A&A*, **457**, 371
- Fryxell, B., et al. 2000, *ApJS*, **131**, 273
- Gaburov, E., Harfst, S., & Portegies Zwart, S. 2009, *New Astron.*, **14**, 630
- Gingold, R. A., & Monaghan, J. J. 1977, *MNRAS*, **181**, 375
- Godunov, S. K. 1959, *Math. Sb.*, **47**, 271
- Hallman, E. J., O'Shea, B. W., Smith, B. D., Burns, J. O., & Norman, M. L. 2009, *ApJ*, **698**, 1795
- Hernquist, L., Bouchet, F. R., & Suto, Y. 1991, *ApJS*, **75**, 231
- Hockney, R. W., & Eastwood, J. W. 1981, *Computer Simulation Using Particles* (New York: McGraw-Hill)
- Huang, J., & Greengard, L. 2000, *SIAM J. Sci. Comput.*, **21**, 1551
- Jin, S., & Xin, Z. 1995, *Commun. Pure Appl. Math.*, **48**, 235
- Klypin, A. A., & Shandarin, S. F. 1983, *MNRAS*, **204**, 891
- Knebe, A., Green, A., & Binney, J. 2001, *MNRAS*, **325**, 845
- Kravtsov, A. V., Klypin, A. A., & Khokhlov, A. M. 1997, *ApJS*, **111**, 73
- Landau, L. D., & Lifshitz, E. M. 1987, *Fluid Mechanics* (2nd ed.; Oxford: Pergamon)
- Lucy, L. B. 1977, *AJ*, **82**, 1013
- MacNeice, P., Olson, K. M., Mobarrry, C., de Fainchtein, R., & Packer, C. 2000, *Comput. Phys. Commun.*, **126**, 330
- Martin, D., & Cartwright, K. 1996, Tech. Rep. UCB/ERL M96/66 (Berkeley, CA: Univ. California)
- Merz, H., Pen, U., & Trac, H. 2005, *New Astron.*, **10**, 393
- Miniati, F., & Colella, P. 2007a, *J. Comput. Phys.*, **224**, 519
- Miniati, F., & Colella, P. 2007b, *J. Comput. Phys.*, **227**, 400
- Mitchell, N. L., McCarthy, I. G., Bower, R. G., Theuns, T., & Crain, R. A. 2009, *MNRAS*, **395**, 180
- NVIDIA 2008, *CUDA Programming Guide*, 2.1 (Santa Clara, CA: NVIDIA Corp.), [http://www.nvidia.com/object/cuda\\_develop.html](http://www.nvidia.com/object/cuda_develop.html)
- O'Shea, B. W., Nagamine, K., Springel, V., Hernquist, L., & Norman, M. L. 2005, *ApJS*, **160**, 1
- Pen, U., Arras, P., & Wong, S. 2003, *ApJS*, **149**, 447
- Plewa, T., & Müller, E. 2001, *Comput. Phys. Commun.*, **138**, 101
- Press, W. H., Teukolsky, S. A., Vetterling, W. T., & Flannery, B. P. 2007, *Numerical Recipes. The Art of Scientific Computing* (3rd ed.; Cambridge: Cambridge Univ. Press)
- Regan, J. A., Haehnelt, M. G., & Viel, M. 2007, *MNRAS*, **374**, 196
- Ricker, P. M. 2008, *ApJS*, **176**, 293
- Schive, H., Chien, C., Wong, S., Tsai, Y., & Chiueh, T. 2008, *New Astron.*, **13**, 418
- Seljak, U., & Zaldarriaga, M. 1996, *ApJ*, **469**, 437
- Sod, G. A. 1978, *J. Comput. Phys.*, **27**, 1
- Springel, V. 2005, *MNRAS*, **364**, 1105
- Springel, V. 2010, *MNRAS*, **401**, 791
- Springel, V., Yoshida, N., & White, S. D. M. 2001, *New Astron.*, **6**, 79
- Strang, G. 1968, *SIAM J. Numer. Anal.*, **5**, 506
- Tasker, E. J., Brunino, R., Mitchell, N. L., Michielsen, D., Hopton, S., Pearce, F. R., Bryan, G. L., & Theuns, T. 2008, *MNRAS*, **390**, 1267
- Teyssier, R. 2002, *A&A*, **385**, 337
- Teyssier, R., et al. 2009, *A&A*, **497**, 335
- Trac, H., & Pen, U. 2003, *PASP*, **115**, 303
- Trac, H., Sills, A., & Pen, U. 2007, *MNRAS*, **377**, 997
- van Leer, B. 1974, *J. Comput. Phys.*, **14**, 361
- van Leer, B. 1979, *J. Comput. Phys.*, **32**, 101
- Woo, T., & Chiueh, T. 2009, *ApJ*, **697**, 850
- Xu, G. 1995, *ApJS*, **98**, 355
- Ziegler, U. 2005, *A&A*, **435**, 385