A GPU accelerated Barnes-Hut Tree Code for FLASH4

Gunther Lukat*, Robi Banerjee

Hamburger Sternwarte, Universität Hamburg, Gojenbergsweg 112, 21029 Hamburg

Abstract

We present a GPU accelerated CUDA-C implementation of the Barnes Hut (BH) tree code for calculating the gravitational potential on octree adaptive meshes. The tree code algorithm is implemented within the FLASH4 adaptive mesh refinement (AMR) code framework and therefore fully MPI parallel. We describe the algorithm and present test results that demonstrate its accuracy and performance in comparison to the algorithms available in the current FLASH4 version. We use a MacLaurin spheroid to test the accuracy of our new implementation and use spherical, collapsing cloud cores with effective AMR to carry out performance tests also in comparison with previous gravity solvers. Depending on the setup and the GPU/CPU ratio, we find a speedup for the gravity unit of at least a factor of 3 and up to 60 in comparison to the gravity solvers implemented in the FLASH4 code. We find an overall speedup factor for full simulations of at least factor 1.6 up to a factor of 10.

Keywords: gravitation, hydrodynamics, methods:numerical, stars: formation

1. Introduction

Self-Gravity is a key phenomena in many Astrophysical Simulations, hence the solution of Poisson's equation of the gravitational potential $\phi(x)$ for a given density distribution $\rho(x)$

$$\nabla^2 \phi(x) = 4\pi G \rho(x) \quad (1)$$

is one of the main functions in these simulations. The direct method for evaluating the gravitational potential at position x_i of a point mass m_i requires the evaluation of all pairwise interactions in the system.

$$\Phi(x_i) = -\frac{1}{2}G \sum_{\substack{j=1 \\ j \neq i}}^N \frac{m_j}{|\vec{r}_i - \vec{r}_j|}, \quad (2)$$

where, m_j is the mass of the j th particle and $\vec{r}_i$ and $\vec{r}_j$ are the position vectors of particle i and j . While the direct method with it's $O(N^2)$ arithmetic complexity is conceptually simple, it is obviously unsuitable for larger systems. Luckily several heuristic algorithms are available which require fewer operations within acceptable error bounds. In general, one can distinguish between the grid-based algorithms and the tree-based algorithms for calculating the gravitational potential. We use the FLASH code package (Fryxell et al., 2000; Fisher et al., 2008) and three of the implemented algorithms: the multipole solver, the multigrid solver and the tree solver.

The *Multipole Poisson solver* is based on a multipolar expansion of the mass distribution around a certain center of expansion. Both accuracy and runtime can be controlled via the multipole cutoff value l_{max} . The multipole approach is by nature appropriate for systems with spherical mass distributions, such that a spherical harmonic expansion can be expected to reach high accuracy after a small number of terms. (Couch et al., 2013)

The *Multigrid Poisson solver* is a modified version of the direct multigrid algorithm of Huang & Greengard adapted to the FLASH4 grid structure (Ricker, 2008). The *Multigrid Poisson solver* is appropriate for general mass distributions.

The *Tree Poisson solver* is based on the Barnes & Hut tree code where the implemented octree is an extension of the AMR mesh tree down to the individual cells (based on FLASH4.2.2 release ?). The tree Poisson solver is appropriate for general mass distributions. In the following sections we refer to the Barnes & Hut tree Poisson solver as CPU-BH tree solver.

Although the heuristic solvers reduce the computation time for the gravity module, they still require a large amount of computation time, often much more than the integration of the MHD equations.

In this paper, we developed a GPU accelerated Barnes & Hut tree Poisson solver for astrophysical applications based on the CUDA runtime library (NVIDIA; Nickolls et al., 2008). Our GPU accelerated Barnes & Hut tree code implements tree walks as well as tree builds solely on the GPU-Device. By accelerating the gravity module with the GPU accelerated Barnes & Hut tree code, we measured a performance improvement of more than factor 20. In the

*Corresponding author

Email addresses: glukat@hs.uni-hamburg.de (Gunther Lukat), banerjee@hs.uni-hamburg.de (Robi Banerjee)

following sections we refer to our GPU accelerated Barnes & Hut tree Poisson solver as GPU-BH tree solver.

2. Algorithm

The Barnes & Hut tree algorithm for three-dimensions works by grouping "particles"¹ using an octree structure (Barnes and Hut, 1986). A single cube containing all particles surrounds the system. This cube is then recursively divided into eight sub cells with each containing their own set of cells. The tree building method continues down in scale until only one particle is left for every sub cell. The tree construction can be done using a bottom-up i.e. inserting one particle at a time or a top-down approach by sorting the particles across divisions. Both methods take $O(N \log N)$ time. For each tree node, the total mass and the center of mass is calculated. The force on a particle in the system is evaluated by "walking" down the tree. At each level, every node is tested against a test particle if it is distant enough for a force evaluation. If the node is too close, it is "opened" and the 8 children are selected for the same procedure. Various criteria exist to test whether a particle is sufficiently distant for a force evaluation. The most common and simplest criterion is based on the opening angle parameter θ (Barnes and Hut, 1986). If the size of a node is l and the distance of the particle from the cell center of mass is d , the node can be accepted for a force evaluation if

$$d > l/\theta. \quad (3)$$

Smaller values of θ lead to a higher accuracy cause by more node selected for opening. Typically a θ value of 1 leads to errors around 1% (Hernquist, 1987). In some cases, in which the center of mass is near the edge of a node, the basic criterion described in Eq. 3 can produce large errors (Salmon and Warren, 1993) Various alternatives are given in literature to avoid this problem. (Salmon and Warren, 1993; Barnes, 1994) Here, we adopt the opening angle parameter described by Barnes (1994) with

$$d > l/\theta + \delta, \quad (4)$$

where δ is the distance between the center of mass of the node and the geometric center (see fig. 1). This criterion guarantees that if the center of mass is near the node's edge, only positions removed by an extra distance δ use the cell for a force evaluation. In case the center of mass is near the node's center, the old criterion (Eq. 3) is used.

2.1. The Parallel Tree-Algorithm

The parallelization of the Barnes & Hut algorithm is not obvious, since the inhomogeneous distribution of the

Figure 1: The Geometry of the Barnes & Hut opening criterion used in the GPU-BH tree code.

particles in the tree does not lead to simple load balanced domain decompositions. To distribute the work to many independent processors, the particles and the tree must be divided in a balanced way. Possible parallelization methods are described in (Salmon, 1991) and (Dubinski, 1996). The basic workflow of one parallel Barnes & Hut algorithm is outlined in fig. 2. The parallel Barnes & Hut algorithm starts with a domain decomposition to achieve a load balanced particle distribution. The decomposition can be carried out using any suitable algorithm e. g. the method of orthogonal recursive bisection the method of orthogonal recursive bisection (Dubinski, 1996), the Morton space filling curve (Bédorf et al., 2012) or a Peano-Hilbert space filling curve (Bédorf et al., 2014).² Given a load-balanced distribution of particles among different processors, a BH-tree can be constructed using the locally stored set of particles (fig. 2, Construct local trees). If every processor had a copy of the complete domain, the force calculation could be carried out now. Unfortunately, the amount of memory required for a full copy of the domain data is prohibitive and thus, we use a smaller but sufficiently large enough sub-tree.

The *LET (locally essential trees)* approach assumes that only a subset of tree nodes is necessary since the BH-tree can be pruned by using the opening angle criterion (Salmon, 1991). Applying the opening angle criterion to the entire group of particles on a donating processor, a subset of nodes can be selected that are necessary for a successful force evaluation on a different processor. The selected nodes are then sent to the target processor (fig. 2, Exchange tree nodes). On the target processor, the existing local tree is enlarged in this way by a pruned version of the trees resident on other processors (fig. 2, Exchange tree nodes). This *local essential tree* is then traversed and the gravitational potential for the particles is evaluated (fig. 2, Tree walk).

3. The GPU Barnes-Hut Tree Code for FLASH4

The approaches from (Salmon, 1991) and from (Dubinski, 1996) have been modified and implemented into the

¹We use the term "particles" as a proxy for any mass-item, including grid cells.

²The FLASH code implements the Morton space filling curve for domain decomposition.

Figure 2: Flowchart of the parallel Barnes & Hut algorithm as it is implemented with our GPU accelerated gravity solver into the FLASH code. Given an initial domain decomposition, we construct a BH tree covering the locally stored data cells. The local trees are traversed and essential tree nodes are exchanged between the processes. With the Essential Nodes and the local data, we build local essential tree (LET) and traverse it to calculate the gravitational potential. The calculated potential is then ready to use by other FLASH modules and the gravity solver can be called in the subsequent simulation step.

FLASH software (Fryxell et al., 2000). Our implementation of the BH-tree algorithm is written in FORTRAN and C for the code running on CPU’s and in CUDA-C for the code running on the GPU-Device. Accessing the global data fields of the FLASH package and calling message-passing routines is solely done using CPU code. The tree walks and tree builds are carried out on the GPU-Devices using several CUDA kernels.

3.1. The CPU Code

The major part of the CPU Code covers the allocation of data fields as well as accessing the FLASH routines for writing and reading global fields. Performance wise, the most important parts of the CPU code are the calls to message-passing routines. Generally, message-passing takes a large amount of time and therefore plays an important role in optimization procedures. In our implementation, message-passing routines are called to communicate pruned BH-trees and BH root node data. The naive strategy to calculate the forces and communicate data-on-demand requires a two-way communication and may cause a large communication overhead if not programmed carefully. Furthermore, a data-on-demand communication does not follow the aim to calculate the forces on the GPU. A GPU kernel would have to be stopped while waiting for the data and restarted after the data was received. To overcome this limitation, the BH-LET can be constructed prior to the force calculation (Liu and Bhatt, 2000). For each FLASH call to calculate the gravitational potential, every process runs through a series of steps. Figure 3 illustrates a basic flowchart of the algorithm.

Whenever the GPU-BH tree solver is called from inside the FLASH4 environment, the AMR-Nodes are expanded

Figure 3: Basic outline of the algorithm steps used in the GPU-BH tree code. Notes in brackets mark whether the respective routine runs on the CPU or the GPU and if message-passing routines are called (MPI).

and the cell data is read (fig. 3, Collect cell data). As a next step, some basic information about the domain’s size and its orientation is distributed along all processes (fig. 3, Exchange basic data). In step three (fig. 3, Calculate and exchange Essential Nodes), we build the local BH-trees and calculate the LET nodes for each process, which are delivered to the respective processes. The tree builds and the tree walks in this step are implemented solely on the GPU device. In step four (fig. 3, Construct BH-LET and calculate potential), we build the final BH-LET on the GPU-device and calculate the gravitational potential. Finally, the calculated potential is written back to the FLASH internal solution vector (fig. 3, Write forces to solution vector).

3.1.1. Collect cell data and exchange basic data

The FLASH4 simulation software (Fryxell et al., 2000) uses the PARAMESH library (MacNeice et al., 2000; Plewa et al., 2005; Deane et al., 2006) to handle its AMR mesh and the domain decomposition. Whenever the GPU-BH tree solver is called from inside the FLASH4 environment, the AMR-Nodes are expanded and the AMR cell data is read (fig. 3, Collect cell data). To construct a local root node, the Center of Mass of the local data cells is calculated on the fly while parsing the AMR tree. Position data and the respective mass of every data cell is read and stored in arrays with a one-dimensional layout. For later usage on the GPU, the arrays are structured into blocks holding the data of exactly 32 neighboring cells. This does

not only preserve locality of the data cells but matches the *warp*³ size of the GPU, which is beneficial for the tree walks.

A virtual root node (in the following referred to as "VN") located at the center of mass of all local data cells is constructed. Every VN holds position and mass data as well as its size in a double precision floating point variable. Because every process needs a copy of all VNs and the amount of data communicated in this step is comparably small, a single `MPI_Allgather` call realizes the VN exchange.⁴ This way, the owner of every VN can be figured out solely by the index.

3.1.2. Calculate and exchange Essential Nodes

As mentioned before, the overlap of communication time with calculation is an important concept in distributed memory approaches. When using the LET approach, the chance of a communication and computation overlap during the force evaluation is obviously lost. Since the vast majority of communication time is needed when the Essential Nodes are exchanged, the respective calculations can take place during this communication procedure (fig. 3, Calculate and send essential sub-trees).

Figure 5 shows a flowchart of the procedures and how the overlap is implemented. Whenever a process acquires a GPU-device, the device is blocked for other processes and a local BH tree is build. The tree is traversed on the GPU for every imported VN to calculate the respective Essential Nodes. After all Essential Nodes for one VN are calculated; an index array is copied to host memory⁵. We use an index array instead of the floating-point node data to reduce the memory operations during the main loop. The actual array of the node data is based on the index array and constructed on the CPU and is sent to the respective VN owner using non-blocking routines. The message size varies depending on the cell opening criterion and the number of data cells a processor holds. For a process holding n data cells, the number of Essential Nodes may vary between 1 (the process just sends it's root node) and n (the process sends all data cells). Processes holding no data cells do not access the GPU device and do not build up a local tree. Figure 4 illustrates the structure of the index- and the exchange-array.

The next VN can be processed immediately, because non-blocking MPI routines (`MPI_Isend`) are used for the communication in this step. On the receiver's side, the Essential Nodes are received while waiting for a free GPU device. Ideally, the waiting times to access a GPU device and the actual calculation time on the device overlap

Essentials index:

Exchange array:

Figure 4: Structure of the index array and the exchange array. The index array holds pointers to the position and mass data of Essential Nodes and data cells. For N Essential Nodes, N integer values are copied from device.memory to host-memory. The exchange array is built by means of the index array holding $4N$ entries.

Figure 5: Flowchart of the steps carried out when calculating and distributing Essential Nodes. Blue boxes indicate procedures carried out on the GPU device, green boxes indicate CPU procedures. A serial GPU-device access is assumed for all contributing processes $P1...PN$ (red circles). Process $P1$ acquires the GPU and starts building a local BH tree. Other processes ($P2...PN$) have to wait until the GPU device is freed by $P1$. The local BH tree is traversed for each imported VN (blue box (Walk tree)) and the calculated Essential Nodes are sent to the respective processes (green box (send Nodes)). In case all VNs are processed, $P1$ frees the GPU device and $P2$ can acquire it to fulfill the same tasks.

with the communication time. This is true for all but the first processes accessing the GPU device. The actual procedures carried out on the GPU device are outlined in section 3.2.

3.2. The CUDA Kernels

The whole tree building process together with the calculation of the essential sub trees and the gravitational potentials is solely done in device memory. This approach minimizes memory transactions like copying data from the host to the device or vice verse. Only position and mass data is copied to the device and just results are copied back to host memory. The construction of the Barnes & Hut trees and the memory layout is based on the findings reported by (Burtcher and Pingali, 2011). In the following, the data layout and the largest CUDA kernels are

³A *warp* is a set of threads scheduled at the same time on the CUDA device. The maximum number of threads in a warp is 32.

⁴Using `MPI_Allgather`, the block of data sent from the j th process is received by every process and placed in the j th block of the receive buffer.

⁵Host memory is maintained by the CPU in its own separate memory space in DRAM

Figure 6: Illustration of the device array structures. The Node-size l and the distance δ is only used for internal nodes. The array for the gravitational potential is only used for leaf nodes (data cells). Position (x, y, z) and mass (m) arrays are used for both.

described in detail.

3.2.1. Data Layout

We use several aligned scalar arrays and array indices to allow for coalesced memory access. The common fields (three dimensional position and mass) for cell data (leaf nodes) and internal nodes are represented with four floating-point arrays. The leaves are allocated at the beginning and the internal nodes at the end of the arrays. Other fields are only valid for leaf nodes or internal nodes. For the latter, we use one array to store the node size (l) and one array to store the distance (δ) between the node's center of mass and it's geometric center. The node size l is calculated and stored during the tree build (see section: 3.2.2) and the distance δ is evaluated while calculating the centers of mass (see section: 3.2.3). Finally, the array for the gravitational potential is only used for the leaf nodes. Figure 6 outlines the described memory structure.

3.2.2. Tree Building Kernel

The *Tree Building Kernel* implements an iterative tree-building algorithm. Starting with a root node, all local data cells and all Essential Nodes are inserted into the tree. In case of a local BH-tree, as it is used for determining Essential Nodes, the root nodes position data exactly at the geometric center of all local data cells. In case of a LET, the root node is located at the domain's center. For the LET it is important that all processes agree on the root node level, because Essential Nodes are treated as data cells during the tree build. In the *Tree Building Kernel*, data cells are assigned to blocks and threads within blocks

in a round robin fashion⁶. Every thread traverses the tree down to the desired leaf node and inserts an index pointing to it's cell into the free leaf node. For this, the respective leaf node is locked such that other threads attempting to access this node have to wait until the appropriate index is inserted.

3.2.3. Centers of Gravitation Kernel

In the *Centers of Gravitation Kernel* the masses of the internal nodes are calculated, their three-dimensional position is updated to the center of gravity and the distance between the node's center of gravity and it's geometric center is evaluated. During the tree build all internal nodes are initiated with a negative mass indicating that their actual mass needs to be calculated. Since mass data is only available for the leaf nodes (the nodes containing cell data), the masses are summed up in a bottom-up manner. The tree is walked from bottom up where internal nodes are assigned to blocks and threads within blocks. Here, every thread processes one node and it's direct children. For a node N with n children $C_i, i = 1..n$ each with mass m_i and position data with coordinates $r_i, i = 1, .., n$. The mass of the node M is simply the sum of the children's masses $m_i, .., m_n$ and it's position is at R with:

$$R = \frac{1}{M} \sum_{i=1}^n m_i r_i \quad (5)$$

A thread with it's assigned internal node at the very bottom of the tree simply sums up the masses and updates it's node mass. The geometric center and it's distance to R is calculated and stored in a separate array before the node's mass and position is updated. Threads that encounter a child with negative mass wait until it's mass is updated by another thread.

3.2.4. Find Essential Nodes Kernel

To generate the essential sub trees, the local BH-trees must be traversed with regard to the VNs. Nodes and data cells are treated as essential if they pass the opening angel test. We use a slightly modified version of the opening angle criterion described in Eq. 4. The imported VNs are treated as nodes during the tree walk, and the distance between a VN and an internal BH node must be evaluated. Here, we refer to the minimum distance d_c between the nodes center of mass and the VN (see fig. 7). In general, the selection of Essential Nodes can be done using any tree traversal method. A simple algorithm targeted to GPU devices would be to transform a depth-first or breath-first traversal with each VN assigned to one GPU-Thread. To support the earlier described method of overlapping and to minimize the memory footprint, the selection of Essential Nodes is implemented for single VNs.

⁶CUDA blocks are organized into a one-dimensional, two-dimensional, or three-dimensional grid of thread blocks. On current GPUs, a thread block may contain up to 1024 threads (NVIDIA).

Figure 7: The geometry of the modified opening criterion used in the GPU-BH tree code when calculating the Essential Nodes.

Opposed to assigning the VNs to threads, the local tree nodes are assigned to blocks and threads within blocks in a round robin fashion starting from the root node. The *Find Essential BH-Nodes Kernel* generates an index list holding indices for all nodes and data cells treated as essential for the processed VN. The CUDA kernel is called once for each depth and thus block level synchronization is assured. Neither atomic operations nor direct synchronization calls are used. The generated index list is copied back to host memory, where the respective communication buffer is filled with the position and mass data of the Essential Nodes and data cells.

3.3. Calculate Gravitational Potential Kernel

Given a tree with updated centers of gravity and masses, the gravitational potential is calculated for all local data cells. The corresponding CUDA-Kernel performs a stack based depth-first tree traversal. The stack is created in shared memory to guarantee fast memory access. The stack size is aligned to the warp size and only the first thread of a warp (lane 0) is allowed to modify the stack. Local data cells are assigned to blocks and threads within blocks in a round robin fashion. The first thread of the warp traverses the tree and pushes the nodes onto the stack. Every thread reads the top node on the stack and calculates the distance between its data cell and the node in the stack. The opening angle parameter test according to Eq. 4 is applied and the gravitational potential for the data cell is updated in cache. Note that the opening angle test works as a data-dependent conditional branch where threads in a warp could diverge. This is a potential performance bottleneck, because the warp would serially execute each taken branch. To overcome this performance lack, the thread voting function `all()`⁷ is used for the outcome of the test. As a result, every thread in the warp "opens" a node even if just one of the threads in the warp has a respective test outcome. Although this prevents thread divergence and can potentially increase the accuracy, the beneficial effect of the opening angle test is obvious lost

for some threads. In a worst-case scenario every warp traverses the whole tree because the warps data cells do not share the same test outcomes. The number of diverging test outcomes can be reduced if warps process only neighboring data cells, because nearby data cells share the same interaction list (Barnes, 1990). For this, we store the data cells in blocks of exactly 32 (warp size) neighbors such that every warp reads a set of nearby data cells.

4. Results

In this section, we describe the results of three astrophysical motivated test problems solved using our GPU-BH tree algorithm: the gravitational potential of a homogeneous spheroid, a homogeneous spheroid with initial velocity turbulence and a homogeneous oblate spheroid. All simulations were carried out in FLASH4.2.2 (Fryxell et al., 2000).

The simulations were carried out on our institute's cluster where every compute node holds two Intel Xeon Processors with each having six cores with a clock speed of 2.4 GHz and one NVIDIA Tesla C2075 GPU device with 6 GB of GDDR5 memory and 448 CUDA cores with a clock speed of 1.15 GHz. The CPU-Code was compiled using the Intel compilers (ICC v. 10.1, IFORT v. 10.1) with aggressive optimization flags (-O3). The GPU code was compiled using the NVCC compiler (Cuda compilation tools, release 7.0, V7.0.27) with default settings. The runtime of the simulations were measured using different timer functions provided by the FLASH4 API.

⁷The `all()` warp voting function evaluates a predicate for all active threads of the warp and returns non-zero if and only if the predicate evaluates to non-zero for all of them (NVIDIA).

Figure 8: Potential of a MacLaurin spheroid with eccentricity 0.9 as computed with the GPU-BH tree code using $\alpha = 0.25$ and using an adaptively refined mesh with seven levels of refinement. The left image shows the potential and AMR block outlines (grey and black lines) in the x-y plane passing through the center of the ellipsoid. The right image shows the same quantities for the y-z plane.

4.1. Accuracy with the MacLaurin Sphere

To test the accuracy and the influence of the opening angle parameter α , we use a so called "MacLaurin" spheroid. The gravitational potential at the surface of, and inside such a homogeneous spheroid is expressible in terms of analytical functions (Chandrasekhar, 1969). For a point inside the spheroid with density ρ the gravitational potential is

$$\Phi(x) = \pi G \rho [2A_1 a_1^2 - A_1(x^2 + y^2) - A_3(a_3^2 - z^2)], \quad (6)$$

where a_1, a_2 and a_3 are the semi major axes of the spheroid and $a_1 = a_2 > a_3$. Here

$$A_1 = \frac{\sqrt{1-e^2}}{e^3} \sin^{-1} e - \frac{1-e^2}{e^2}, \quad (7)$$

$$A_3 = \frac{2}{e^2} - \frac{2\sqrt{1-e^2}}{e^3} \sin^{-1} e, \quad (8)$$

where e is the ellipticity of the spheroid:

$$e = \sqrt{1 - \left(\frac{a_3}{a_2}\right)^2}. \quad (9)$$

For a point outside the spheroid, the potential is:

$$\begin{aligned} \Phi(x) = & \frac{2a_3}{e^2} \pi G \rho \\ & \times \left[a_1 e \tan^{-1} h - \frac{1}{2} \left((x^2 + y^2) \left(\tan^{-1} h - \frac{h}{1+h^2} \right) \right. \right. \\ & \left. \left. + 2z^2(h - \tan^{-1} h) \right) \right], \end{aligned} \quad (10)$$

where

$$h = \frac{a_1 e}{\sqrt{a_3^2 + \lambda}}, \quad (11)$$

and λ is the positive root of the equation

$$\frac{x^2}{a_1^2 + \lambda} + \frac{y^2}{a_2^2 + \lambda} + \frac{z^2}{a_3^2 + \lambda} = 1. \quad (12)$$

The simulation was setup with a uniform density $p = 1 \text{ g cm}^{-3}$ inside the spheroid and $p = \epsilon \rightarrow 0$ outside the spheroid with an eccentricity of 0.9. The spheroids were centered in a box with unit dimensions. An adaptively refined mesh with seven levels of refinement ($\approx 7.2m$ data cells) was used. Using our GPU-BH tree code and the CPU-BH tree code within FLASH4, we computed potentials with varying α values ($\alpha = 1.0, 0.75, 0.5, 0.25, 0.1$). An example of the potential for $\alpha = 0.25$ computed with the GPU-BH tree code using a maximum refinement level 7, is shown in fig. 8. All simulations were run using 12 CPU-Cores with two GPU devices utilized by the GPU-BH tree code.

We compare the analytical solution for the gravitational potential $\phi_{\text{MacLaurin}}$ to the potential calculated with the GPU-BH tree code ϕ_{GPU} and the CPU-BH tree code ϕ_{CPU} . For the tests, we evaluate the relative error ϕ_{err} with:

$$\phi_{\text{err}} = \left| \frac{\phi_{\text{MacLaurin}} - \phi_{\text{<algorithm>}}}{\phi_{\text{MacLaurin}}} \right|. \quad (13)$$

Figure 9: Maximum relative errors (ϕ_{err}) in the gravitational potential for the MacLaurin sphere calculated with the GPU-BH tree code (—●—) and the CPU-BH tree code (—■—) for different α values. The GPU-tree code produces a higher accuracy for α values between 0.2 and 0.8 while the CPU tree code shows an accuracy advantage for α values > 0.8 . Here, an alpha value of 0.0 refers to the direct $O(n^2)$ algorithm.

For comparison, we calculated the potential with the direct $O(n^2)$ method where a relative error of 2.17×10^{-4} was calculated. Typically, the Barnes & Hut algorithm achieves higher accuracy with lower α values (Barnes and Hut, 1986; Hernquist, 1987). Figure 9 clearly shows that this is true for the GPU-BH tree code and for the CPU-BH tree code. Both, the GPU-BH tree code and the CPU-BH tree code produced relative errors below 1%. For α values near 0.2 both solvers nearly reached the accuracy of the direct summation $O(n^2)$ method. For α values between 0.25 and 0.8 the GPU-BH tree code calculated a more accurate potential than the CPU-tree code. For values above 0.8 the CPU-BH tree achieved a higher accuracy. The differences between the GPU-BH and the CPU-BH accuracies are a direct result of the different strategies when applying the opening angle test.

4.2. Scaling of the GPU-BH tree code with the top-hat sphere setup

In this section, we summarize the results of different performance tests based on the collapse of spherical cloud cores with and without initial velocity fluctuations. The setup parameters of our different simulations can be found in tables 1 and 2.

4.2.1. Simulation setup

We perform the scaling test with a top-hat (TH) sphere setup with the following parameters (see also table 1). We use a mean molecular weight of $\mu = 2.3$ g/mol with an isothermal equation of state (EoS) at a temperature of

20 K. The density profile of the system with a sphere radius of R_0 is described by a step function

$$\rho = \begin{cases} \langle \rho \rangle & \text{for } r \leq R_0 \\ 0.01 \langle \rho \rangle & \text{for } r > R_0. \end{cases} \quad (14)$$

With the density $\langle \rho \rangle = 1.76 \times 10^{-18}$ g cm $^{-3}$ leading to a free fall time of

$$t_{\text{ff}} = \sqrt{\frac{3\pi}{32G\langle \rho \rangle}} = 1.58 \times 10^{12} \text{ s} \approx 50 \text{ kyr}. \quad (15)$$

The initial sphere is highly gravitationally unstable with the Jeans length

$$\lambda_{\text{J}} = \sqrt{\frac{\pi c_s^2}{G\langle \rho \rangle}} = 5736 \text{ au} = 0.28 R_0. \quad (16)$$

We use the top-hat sphere with varying parameters for our performance tests. To investigate the performance influence of varying GPU/CPU ratios, we use different sphere radii (R_0) and refinement levels. One simulation setup with a sphere radius of $R_0 = 3.5 \times 10^{17}$ cm and $l_{\text{max}} = 6$ resulting in $\approx 7 \times 10^6$ data cells (SP-RSS) and one with a radius of $R_0 = 1.73 \times 10^{17}$ cm and $l_{\text{max}} = 7$ resulting in $\approx 20 \times 10^6$ data cells (SP-RSL).

For the weak scaling test simulations (SP-WS), we scale the problem size to the number of processing units by modifying the sphere radius R_0 . Here, we use maximum refinement level of $l_{\text{max}} = 8 \Rightarrow \Delta x = 7.8125 \times 10^{14}$ cm. We also performed collapse simulations with the turbulent TH setups (SP-TL) for a runtime of exactly one week. Here, we initialize the TH-profile with an additional random supersonic velocity field with a Mach number of $M_a = 2.0$. The average turbulence crossing time is $t_{\text{tc}}(R_0) = 1.8 \times 10^5$ yrs which is about 3.5 times larger than the free fall time t_{ff} . An overview of the physical parameters is given in table 1. In these simulations, we set the maximum refinement level to $l_{\text{max}} = 15$ which results in a maximum resolution of 131072 grid cells in one direction, corresponding to $\approx \Delta x \approx 0.4$ au. For the collapse simulations, we applied the Truelove criterion (Truelove et al., 1997) to resolve the Jeans length

$$\lambda_{\text{J}} = \sqrt{\frac{\pi c_s^2}{G\rho_{\text{max}}}} \quad (17)$$

throughout the simulation. Furthermore, we used the sink particle module with the Federrath criterion for sink creation (Federrath et al., 2010). The sink particles have an accretion radius of $r_{\text{acc}} = 3\Delta x$ which leads to the threshold density ρ_{max} of

$$\rho_{\text{max}} = \frac{\pi c_s^2}{4G(3\Delta x)^2} = 2.518 \times 10^{-11} \text{ g cm}^{-3} \quad (18)$$

For the strong scaling tests, we use the SP-TL simulation setup with a resolution of 12 grid cells per jeans

Table 1: Simulation parameters for the turbulent top-hat sphere

Parameter		Value
Simulation box size	L_{box}	8×10^{17} cm
Smallest cell size	Δx	6.1×10^{12} cm
Max. refinement	l_{max}	15
Min. refinement	l_{min}	4
Sink particle accr. radius	r_{accr}	1.83×10^{13} cm
Max. density	ρ_{max}	2.52×10^{-11} g cm $^{-3}$
Opening angle parameter	α	0.5
Sphere radius	R_0	3×10^{17} cm
Total sphere mass	M_{tot}	100 $M_{\odot}$
Mean density	$\langle \rho \rangle$	1.76×10^{-18} g cm $^{-3}$
Max. gas density	ρ_{max}	9.67×10^{-12} g cm $^{-3}$
Sound speed	c_s	0.27×10^4 km s $^{-1}$
Mean free fall time	t_{ff}	1.58×10^{12} s $\approx$ 50 kyr
Turbulent crossing time	t_{tc}	5.5×10^{12} s
Jeans Mass	M_{J}	1.23 $M_{\odot}$
Jeans length	λ_{J}	1.4×10^{17} cm

Physical and numerical simulation parameters for the turbulent top-hat sphere simulations.

length. We run the simulation with our accelerated gravity solver until two collapse regions including several sink particles are formed ($t_s = 1.2 \times 10^{12}$ s $\approx$ 3.8×10^4 yrs). The corresponding checkpoint file is then used as an initial structure for the strong scaling test simulations (SP-TRS) with our accelerated solver, the CPU-BH tree solver and the Gridsolver.

4.2.2. Scaling with different GPU/CPU ratios

In general, FLASH 4 (Fryxell et al., 2000) simulations are setup for distributed-memory machines where several MPI (Gabriel et al., 2004) processes are created. Since every process holds it's own CUDA context and kernels are effectively serialized as they launch on the GPU, several MPI ranks compete for the device and slow down the code execution. Figure 10 shows an example sequence of the access order for 4 processes trying to access the same GPU device (4 MPI Ranks per compute node).

In order to compare the runtimes of the GPU-BH tree code at different CPU/GPU ratios with the runtime of the CPU-BH code, we calculate the respective speedup S with:

$$S = \frac{T_{\text{CPU-BH}}}{T_{\text{GPU-BH}}} \quad (19)$$

where $T_{\text{CPU-BH}}$ is the time for one gravity step evaluated for the CPU-BH tree code and $T_{\text{GPU-BH}}$ is the time for one gravity step evaluated for the GPU-BH tree code. Independent from the problem size (number of data cells), the highest speedup is expected for a one to one GPU/CPU ratio. We analyzed the performance depending on different GPU/CPU ratios with the SP-RSS run using $\approx 7 \times 10^6$

Figure 10: Illustration of the serial access sequence for processes competing for the same GPU device. Direct access to the GPU device is only granted for the first process in the queue (Rank 1). All other processes wait for the GPU device to be available again.

Figure 11: The initial density and AMR block distribution for the SP-RSL setup. The image shows a slice through the box along the X-axis. An adaptively refined mesh with seven levels of refinement was used with a sphere radius of 1.73×10^{17} resulting in $\approx 20 \times 10^6$ data cells.

cells and with the SP-RSL run with $\approx 20 \times 10^6$ cells (see fig. 11).

Figure 12 shows the speedup calculated with Eq. 19 for GPU/CPU ratios of 1/12, 1/8, 1/6, 1/4, 1/2 and 1/1 while maintaining a constant number of 24 cores for each ratio. Here, we reach a maximum speedup of ≈ 45 with the one-to-one GPU/CPU ratio (24 GPU devices and 24 CPU cores) for the simulation with $\approx 20 \times 10^6$ data cells (fig. 12, $\bullet$). Similarly, the lowest speedup factor of ≈ 5 was evaluated for the 1/12 ratio (2 GPU devices and 24 CPU cores) caused by the long GPU device access times following the serialized access pattern illustrated in fig. 10.

One way to circumvent this performance bottleneck is to collect all data on one of the CPU cores sharing one device (fig. 12, $\blacktriangle$).⁸ Although intra node communication

⁸This is only practical for small problem sizes because the GPU-

Table 2: Scaling of the different top-hat sphere simulations

Name	Solver	Steps	R_0 [cm]	$l_{\max}$	Δx [cm]	Resolution [cells]	Cell count	Figure
SP-RSS	CPU-BH	10	3.50×10^{17}	6	3.12×10^{15}	256^3	7×10^6	12
	GPU-BH	10	3.50×10^{17}	6	3.12×10^{15}	256^3	7×10^6	12, 13
SP-RSL	GPU-BH	10	1.73×10^{17}	7	1.56×10^{15}	512^3	20×10^6	11, 12, 13, 14
	GPU-BH	10	4.50×10^{16}	8	7.81×10^{14}	1024^3	2.73×10^6	15, 16, 17
SP-WS	CPU-BH	10	-	8	7.81×10^{14}	1024^3	-	15
			3.60×10^{17}				1.48×10^8	
			4.50×10^{16}				2.73×10^6	
			-				-	
SP-TL	GPU-BH	2	2.61×10^{17}	15	6.10×10^{12}	131072^3	7.42×10^7	18, 19, 20, 20
			3.0×10^{17}				3.5×10^7	
			3.0×10^{17}				4.4×10^7	
			3.0×10^{17}				1.85×10^7	
SP-TS	GPU-BH	2	3.0×10^{17}	15	6.10×10^{12}	131072^3	1.85×10^7	22, 23, 24, 25, 26, 27, 28
			3.0×10^{17}				1.85×10^7	
			3.0×10^{17}				1.85×10^7	
			3.0×10^{17}				1.85×10^7	
SP-TS	CPU-BH	2	3.0×10^{17}	15	6.10×10^{12}	131072^3	1.85×10^7	22, 23, 25, 26, 27, 28
	Gridsolver	2	3.0×10^{17}	15	6.10×10^{12}	131072^3	1.85×10^7	22, 23, 25, 26, 27, 28

Top-hat sphere simulations used for our scaling tests. The cell count for the SP-TL setups refer to the final stage of the simulations. The radii for the the SP-WS simulations are used to scale the cell count for the weak scaling test simulations (see table 3 for a list of radii and their corresponding cell count).

is comparably fast and the number of required tree builds is reduced, the evaluated speedup factors for this approach are generally lower compared to the previously described approach.

Nevertheless, we achieved a nearly linear scaling efficiency considering the count of GPU-devices. In Fig. 13 we show the strong scaling efficiency for different GPU to CPU ratios. Starting with a time t_1 for one gravity step at a GPU/CPU ratio of 1/12, the percentual efficiency at a specific ratio with the respective time t_r is given as:

$$SC_{\text{ratio}} = \frac{t_1 - t_r}{t_1} * 100 \quad (20)$$

The results are shown for the two simulations SP-RSS and SP-RSL. Note, that the base ratio of 1 GPU device for 12 CPU cores (t_1) is achieved using 2 compute nodes with a total of 24 cores and 2 GPU devices and already includes communication times for inter and intra node communication. (The presented speedup and scaling values refer to the gravity solver module only.) Although we measured a speedup factor of ≈ 40 compared to the original CPU-BH tree code (see fig. 12, $\bullet$), the entire wall time for the simulation did not show the same speedup since our gravity solver only makes up less than half of the entire wall time for the related GPU/CPU ratio.

Figure 14 shows the percentual fractions of our GPU accelerated gravity solver and the hydro solver for different GPU/CPU ratios. The presented values are evaluated for

Figure 12: Speedup factors for the gravity module using different GPU/CPU ratios for the GPU-BH tree code in comparison to the CPU-BH tree code. The Y-axis refers to the speedup of the GPU-BH tree code in comparison to the CPU-BH tree code and the X-axis refers to the GPU count for different GPU/CPU ratios using a total of 24 CPU cores. We see the highest speedup with the SP-RSL setup ($\bullet$). With the smaller simulation setup SP-RSS ($\blacksquare$), the speedup is slightly lower especially when using less CPU cores per GPU. The speedup factors with the SP-RSS setup-using the gathered data approach ($\blacktriangle$) show the lowest speedup compared to the other setups. The highest speedup of ≈ 45 was gained for a 1/1 GPU/CPU ratio (24 GPU devices and 24 CPU cores) with the SP-RSL simulation setup holding $\approx 20 \times 10^6$ data cells ($\bullet$).

device memory is comparably small and limited.

Figure 13: The strong scaling efficiency of our GPU-BH tree solver calculated with Eq. 20 for varying CPU/GPU ratios and a constant number of CPU cores. The X-axis refers to the GPU count for different GPU/CPU ratios using a total of 24 CPU cores. We see a nearly linear strong scaling for the SP-RSS simulation (—■—) and the SP-RSL simulation (—●—).

the (SP-RSL) simulation with $\approx 20 \times 10^6$ data cells.⁹ At a ratio of 1/12, we find the runtime share of our gravity solver is the dominating part with $\approx 65\%$ of the total wall time.¹⁰ Since our accelerated gravity solver benefits from a lower GPU/CPU ratio, it's contribution to the entire simulation time is reduced with lower ratios. Hence, we measure lower runtime shares with lower ratios. Already for a ratio of 1/6, our solver contributes only $\approx 45\%$ to the wall time. Finally, we measure only a $\approx 23\%$ contribution for a ratio of 1/1 and find the hydro solver to be the dominating code part (see fig. 14).

4.2.3. Weak scaling with the top-hat sphere

For the weak scaling, the problem size (number of data cells) assigned to each processing unit stays constant. We refer to a processing unit, " P_U " as a unit of six CPU cores plus one GPU device.¹¹ We use the weak scaling efficiency given by:

$$\frac{t_1}{(t_N)} * 100, \quad (21)$$

where t_1 is the wall time of one gravity time step with one P_U and t_N is the time of one gravity step with N processing elements. We use the SP-WS top-hat sphere setup and

⁹Values evaluated for the smaller SP-RSS simulation with $\approx 7 \times 10^6$ data cells only differ for $\approx 2\%$ to the presented values.

¹⁰For the same simulation setup the CPU-BH gravity solver dominated the runtime with $\approx 88\%$.

¹¹We choose six cores to fully load exactly one of the hexacore CPUs on each compute node and use the MPI call to control the respective processor binding affinity.

Figure 14: Runtime fractions of the ppm hydro solver (■) and our accelerated gravity solver (■) for different CPU/GPU ratios. The X-axis refers to the GPU count for different GPU/CPU ratios using a total of 24 CPU cores. Values are evaluated for the SP-RSL simulation with $\approx 20 \times 10^6$ data cells. With a lower runtime of the gravity solver for increasing GPU/CPU ratios, the hydro solver becomes the dominating solver for the simulations. For a CPU/GPU ratio of 6/1 (4 GPU devices and 24 CPU cores), already 55% of the evolution time was spent in the hydro routines.

scale the problem size to the number of P_U by changing the sphere radius R_0 and as a result, the overall cell count. Table 3 gives an overview of the different sphere radii and cell counts used for the weak scaling tests. (The scaling results are normalized to match 1 Leaf-node per core.)

In fig. 15 we show the weak scaling efficiency (Eq. 21) for three different simulation sizes with the SP-WS setup holding $\approx 8.68 \times 10^5$ (—●—), 6.72×10^5 (—▲—) and 4.4×10^5 (—■—) data cells assigned to each CPU core. We see a drop of the weak scaling efficiency with increasing number of processing elements. The main reason for this loss in efficiency is not only based on increased communication times also the algorithm design is a factor. Especially the method of calculating the essential nodes plays a role, since the local tree is traversed once for every process. Although this is beneficial to overlap computation with communication time, the loop size increases with higher numbers of processing elements. As a result, more data is copied from device memory to host memory and more GPU kernels are started. The fractional runtimes of the respective memory operations and kernel executions are outlined in fig. 16. With 16 P_U , the memory operations and kernel executions during the calculation of the Essential Nodes make up $\approx 10\%$ of the total GPU time which is dominated by kernel executions. Note, that the actual runtime per kernel does not change but the kernel is started more often. The large number of kernel calls and memory operations is clearly reflected in the runtime of the the routines for calculating the Essential Nodes. Figure 17 illustrates

Table 3: AMR block count with different sphere radii

R_0	Blocks	Leaf-blocks/Core	cells	P_U
<i>SPWS</i> with $\approx 4.4 \times 10^5$ data cells/cores				
4.50×10^{16}	6089	≈ 888	2.73×10^6	1
9.70×10^{16}	24457	≈ 891	1.10×10^7	4
1.36×10^{17}	44937	≈ 819	2.01×10^7	8
1.70×10^{17}	71945	≈ 874	3.22×10^7	12
1.98×10^{17}	94985	≈ 865	4.25×10^7	16
2.20×10^{17}	118537	≈ 864	5.31×10^7	20
2.40×10^{17}	141257	≈ 858	6.32×10^7	24
2.61×10^{17}	165705	≈ 863	7.42×10^7	28
<i>SPWS</i> with $\approx 6.72 \times 10^5$ data cells/cores				
5.70×10^{16}	8905	≈ 1298	3.99×10^6	1
1.19×10^{17}	36617	≈ 1335	1.64×10^7	4
1.70×10^{17}	71945	≈ 1311	3.22×10^7	8
2.08×10^{17}	105993	≈ 1288	4.75×10^7	12
2.43×10^{17}	144905	≈ 1320	6.49×10^7	16
2.70×10^{17}	178889	≈ 1320	8.01×10^7	20
2.98×10^{17}	217033	≈ 1318	9.72×10^7	24
3.23×10^{17}	253577	≈ 1320	1.14×10^8	28
<i>SPWS</i> with $\approx 8.68 \times 10^5$ data cells/cores				
6.50×10^{16}	11721	≈ 1709	5.25×10^6	1
1.36×10^{17}	44937	≈ 1638	2.01×10^7	4
1.98×10^{17}	94985	≈ 1731	4.25×10^7	8
2.40×10^{17}	141257	≈ 1716	6.32×10^7	12
2.75×10^{17}	183497	≈ 1672	8.22×10^7	16
3.10×10^{17}	232905	≈ 1698	5.25×10^6	20
3.40×10^{17}	277193	≈ 1684	1.24×10^8	24
3.60×10^{17}	330377	≈ 1720	1.48×10^8	28

The AMR block count for different sphere radii used to scale the top-hat sphere simulations *SPWS*_{CPU} and *SPWS*_{GPU}.

Figure 15: Weak scaling efficiency in % of linear for the top-hat simulations with the SP-WS setup (gravity only). The scaling values for the smaller simulations with 4.4×10^5 data cells/cores run with our accelerated solver ($\square$) and with the CPU-BH solver ($\bullet$) show a similar scaling trend for more than 15 P_U . We achieved the best scaling results for our solver with the larger system holding $\approx 8.68 \times 10^5$ data cells/cores ($\bullet$).

Figure 16: Weak scaling of the runtime fractions of GPU device memory operations and GPU kernel executions. The blue and the yellow bars refer to the runtimes of the memory operations and the kernel executions during the calculation of the Essential Nodes. The red and the grey bars show the proportions for the code calculating the final potential. The values are calculated for the SP-WS simulation with $\approx 8.68 \times 10^5$ data cells. One P_U refers to 6 CPU cores and 1 GPU device.

the growing runtime contribution of the respective routines with the larger top-hat setup holding $\approx 1 \times 10^5$ data cells per core. We see that the routines for calculating the Essential Nodes take a maximum of $\approx 15\%$ of the total solver runtime with 28 P_U . Nevertheless, at the same count nearly $\approx 80\%$ of the linear scaling efficiency was reached with the small system and $\approx 88\%$ was reached with the larger system. For the smaller setup, the CPU-BH and the GPU-BH simulations show nearly the same weak scaling efficiency for $P_U \geq 15$.

4.2.4. Comparison of GPU-BH and CPU-BH for fixed runtime (7 days)

We run two simulations of the top-hat sphere setup with turbulence and sink particles (SP-TL). One simulation with our GPU accelerated gravity solver and one with the CPU-BH Tree solver. While the CPU only simulation was run with 72 CPU cores, we used 60 CPU cores and 10 GPU-devices for the GPU accelerated simulation. Both simulations were run for $t_w = 168 \text{ h} = 7 \text{ days}$. With the CPU-BH simulation, we reached a total of 2006 evolution steps and a final simulation time of $t_s \approx 1.0 \times 10^4 \text{ yrs}$. Since the simulation stopped at an early stage of the collapse with a maximum density of $9.12 \times 10^{17} \text{ g cm}^{-3}$, neither the maximum refinement level was reached nor any sink particles were formed.

With the GPU accelerated simulation, 4465 evolution steps were executed and a simulation time of $t_s \approx 1.6 \times 10^4 \text{ yrs}$ was reached with a maximum density of $3.73 \times 10^{16} \text{ g cm}^{-3}$. Similar to the CPU only simulation, the

Figure 17: Weak scaling runtime fractions of different parts of the GPU-BH tree gravity solver for the SP-WS simulation with $\approx 8.68 \times 10^5$ data cells per CPU core. The runtime fraction of the routines for calculating the Essential Nodes grows with increasing number of P_U . One P_U refers to 6 CPU cores and 1 GPU device.

maximum refinement level was not reached and no sink particles were formed. Note, that the GPU-BH tree simulation reached the 2006 step mark already after ≈ 2.6 days. An overview of the simulation parameters for the final stages is given in table 4. For comparison, we show slices of the density distributions of both simulations at the same simulation time ($t_s = 10.2 \text{ kyr}$) in fig. 18. A simple com-

Table 4: Final simulation parameters

Parameter	CPU-BH	GPU-BH
Max. refinement	9	10
Leaf-blocks total	69273	85261
Leaf-blocks/CPU	≈ 962	≈ 1421
Max. gas density	$9.12 \times 10^{17} \text{ g cm}^{-3}$	$3.73 \times 10^{16} \text{ g cm}^{-3}$
Evolution steps	2006	4465
Max. sim. time	$3.2185 \times 10^{11} \text{ s}$	$5.2345 \times 10^{11} \text{ s}$

Simulation parameters in the final stage (after 168 compute hours, i.e. 7 days) of the GPU accelerated simulation and the CPU only simulation.

parison between the number of executed evolution steps suggests a speedup factor for the GPU accelerated simulation of ≈ 2.2 . At $t_s = 10.2 \text{ kyr}$ where both simulations hold approximately equal numbers of leaf blocks, we find a total simulation speedup factor for the GPU-accelerated simulation of ≈ 2.7 ¹². Note that this simulations ran with a 1/6 GPU/CPU ratio, nevertheless, we expect a much larger speedup for larger GPU/CPU ratios.

¹²These estimations do not take into account the workload difference caused by different CPU counts.

(a) GPU-accelerated at $t_{\text{sim}} = 1.02 \times 10^4$ yrs, $t_{\text{wall}} = 62$ h

(b) CPU only at $t_{\text{sim}} = 1.02 \times 10^4$ yrs, $t_{\text{wall}} = 168$ h

Figure 18: Slices of the density for the SP-TL simulation at $t_{\text{sim}} = 1.02 \times 10^4$ yrs. The left plot shows the GPU accelerated simulation while the right plot shows the CPU only simulation. The respective simulation step was reached after 7 days with the CPU only solver and after ≈ 2.6 days with the GPU accelerated solver. The CPU simulation was executed with 72 CPU cores and the accelerated GPU simulation was executed with 60 cores and 10 GPU devices.

Furthermore, in fig. 19 we show the runtime fractions of different modules for both SP-TL simulations. Following our expectations from the previous simulations (14), we find the runtime fractions of our gravity solver and the hydro solver to be roughly the same with $\approx 4\%$ difference (fig. 19, GPU (■, ■, ■), GPU-FIN (■, ■, ■)). With the CPU only SP-TL simulation, the CPU-BH gravity solver is by far the dominating module with over 80% wall-time contribution. In the GPU accelerated SP-TL simulation, the number of AMR blocks increases with the simulation time and as a result, the workload for our GPU-BH tree solver grows. The runtime fractions of the major routines executed during the evolution of the accelerated SP-TL simulation are shown in fig. 20. We see a growing runtime contribution of our gravity solver during the first 1000 steps of $\approx 8\%$ from the initial 27% to 35%. After 1000 steps, the runtime contribution slowly converges to $\approx 36\%$. Furthermore, we show the fractional contributions of the tree-algorithm in fig. 21. Here we see that the fractional increase of the gravity solver (see fig. 20) comes mainly from the additional communication time with the growing number of data cells. The communication dominates the runtime during the early steps ($N < 500$) due to small number of grid cells.

4.2.5. Strong scaling with the turbulent top-hat sphere

We evaluate the strong scaling capabilities of our GPU-BH solver with the turbulent top-hat sphere setup. For this, we run the SP-TL simulation until $t_s = 1.2 \times 10^{12}$ sec $\approx 3.8 \times 10^4$ yr. At this point, two collapse regions and 26 sink particles with a total mass $\approx 0.42 M_\odot$ are formed. We

Figure 19: Runtime fractions of different code parts for the CPU only (CPU) and GPU accelerated (GPU, GPU-FIN) SP-TL (CPU) simulations. The values for CPU and for GPU-FIN refer to the whole simulation time of the CPU only and the GPU accelerated simulation. The values for GPU refer to the GPU accelerated SP-TL simulation at a state near the CPU simulation's final state. For the CPU only simulation, the gravity solver dominated the runtime with more than 80%, while for the GPU accelerated simulations, the hydro solver is the dominating module. (The GPU accelerated SP-TL simulation was run with 10 GPU devices and 60 CPU cores.)

Figure 20: Runtime fractions of the major modules used in the GPU accelerated SP-TL simulation. The fractional time of the accelerated gravity solver (—■—) grows as a result of the increasing number of data cells (---).

Figure 21: Average runtime fractions of selected routines of the GPU-BH tree solver during the evolution of the SP-TL simulation. The routines solely executed on the CPU (gathering data, allocation, deallocation, ...) stay at a constant fraction (—●—), while the fraction of the non overlapped communication time (—●—) is highly affected by the increasing cell count (---).

Figure 22: The SP-TL simulation at $t_s = 1.2 \times 10^{12} \text{ s} \approx 3.8 \times 10^4 \text{ yrs}$. Plots show the column density g cm^{-2} for the lower collapse region at different resolutions. From left to right, the box length is reduced from $8.0 \times 10^{17} \text{ cm}$ over $8.0 \times 10^{16} \text{ cm}$ to $3.0 \times 10^{15} \text{ cm}$. Sink particles appear as white dots in the large image to the right.

used a resolution of 12 data cells per Jeans length to scale the simulation to hold $\approx 1.85 \times 10^7$ cells. Figure 22 and 23 show the two collapse regions at different resolutions. We use the final checkpoint file of this simulation as a starting point for the strong scaling tests (SP-TS). For the strong scaling, the problem size stays constant and we refer to the strong scaling efficiency (as a percentage of linear) with

$$\frac{t_1}{(N * t_N)} * 100\%, \quad (22)$$

where t_1 is the time spent in the gravity unit with one processing element, N is the number of processing elements and t_N is the time spent in the gravity unit with N processing elements. Again, we use six CPU cores and one GPU device as a single processing unit (P_U).

Figure 24 shows plots of the strong scaling efficiency of different code parts of our gravity solver and the main modules run during the simulation. The values are evaluated for 4, 8, 12, 16, 20 and 24 P_U . For each test, we measured the runtime for 2 evolution steps starting from the afore mentioned checkpoint file. The total scaling efficiency of our gravity solver dropped down to $\approx 73\%$ for 24 P_U . The loss in efficiency can be attributed to the routines for calculating the Essential Nodes (fig. 24, —●—), since these show the largest drop in efficiency. We trace this back to the increased loop size and increased memory operation cost as a result of the higher number of cores.

For comparison, we run the same simulation setup with the CPU-BH tree solver and the Gridsolver. Figure 25 shows the evaluated strong scaling efficiency for the respective simulations. We see a super linear scaling efficiency of the CPU-BH tree code (fig. 25, —■—) and sub linear strong scaling of our accelerated solver (fig. 25, —●—) and the Gridsolver (fig. 25, —●—). The Gridsolver and our ac-

Figure 23: The SP-TL simulation at $t_s = 1.2 \times 10^{12} \text{ s} \approx 3.8 \times 10^4 \text{ yrs.}$ Plots show the column density g cm^{-2} for the upper collapse region at different resolutions. From left to right, the box length is reduced from $8.0 \times 10^{17} \text{ cm}$ over $8.0 \times 10^{16} \text{ cm}$ to $3.0 \times 10^{15} \text{ cm}$. Sink particles appear as white dots in the large image to the right.

Figure 24: Strong scaling efficiency in % of linear for different code parts of the GPU-BH tree solver and the major evolution routines executed during the SP-TS simulations. One P_U refers to 6 CPU cores and 1 GPU. The strong scaling efficiency of our gravity solver ($\text{---}\star\text{---}$) drops down to $\approx 73\%$ of the linear scaling.

Figure 25: Strong scaling efficiency for the gravity unit in % of linear with the SP-TS setup for our GPU accelerated gravity solver ($\text{---}\bullet\text{---}$) the CPU-BH tree solver ($\text{---}\blacksquare\text{---}$), and the Gridsolver ($\text{---}\bullet\text{---}$). The CPU-BH tree solver shows a super linear strong scaling behavior ($\text{---}\blacksquare\text{---}$), while our accelerated solver ($\text{---}\bullet\text{---}$) and the Gridsolver ($\text{---}\bullet\text{---}$) show non ideal scaling capabilities. One P_U refers to 6 CPU cores and 1 GPU.

celerated solver show an almost identical scaling behavior with a slight advantage of $\approx 3\%$ towards our solver. Although the strong scaling efficiency of our solver shows a non linear behavior, we find a reasonable performance gain compared to the other solvers (see fig. 26). In comparison to the CPU-BH tree code, we find a high speedup of factor 19 with 4 P_U and a speedup factor of 13 with 24 P_U (fig. 26, $\text{---}\bullet\text{---}$). In comparison to the Gridsolver, our GPU accelerated solver was roughly 5 times faster for all numbers of P_U (fig. 26, $\text{---}\blacksquare\text{---}$).

Following the speedup factor, we see a much higher data throughput with our accelerated solver than with the other solvers (see fig. 27). For our accelerated solver, we evaluated a data throughput of $\approx 2.0 \times 10^5$ data cells/sec with just 4 P_U (fig. 27, $\text{---}\bullet\text{---}$). At the same cores, we find much lower data throughput of $\approx 4.7 \times 10^4$ cells/sec for the Grdsolver and $\approx 1.0 \times 10^4$ cells/sec for the GPU-BH tree solver. We find the maximum data throughput values with $\approx 6.6 \times 10^4$ cells/sec with the CPU-BH solver $\approx 1.7 \times 10^5$ cells/sec with the Gridsolver and $\approx 8.7 \times 10^5$ cells/sec with our accelerated GPU-BH solver for 24 P_U . Note that even with additional 120 cores (20 P_U) both the Gridsolver and the CPU-BH tree solver stayed below the $\approx 2.0 \times 10^5$ data cells/sec mark which our solver reached with 4 P_U . The values reported in figure 26 only refer to a GPU/CPU ratio of 1/6 and we find higher speedup factors when utilizing less CPU core for each GPU. Figure 28 shows the calculated speedup factors for different ratios. We find the highest speedup of factor 63 for our GPU-BH tree solver compared to the CPU-BH tree solver when run

Figure 26: Strong scaling speedup factors of our accelerated gravity solver compared to the CPU-BH tree solver and the Gridsolver. The values are evaluated for the SP-TS simulation for the gravity unit only. One P_U refers to a unit of 6 CPU cores and 1 GPU.

Figure 27: Strong scaling data throughput in data cells/sec of the different gravity solvers for the SP-TS simulations. With the CPU-BH tree solver (—■—) and the Gridsolver (—●—), we reached the maximum data throughput of 5×10^4 respectively 1.16×10^5 data cells per second with 24 P_U . A similar value of $\approx 1.97 \times 10^5$ data cells per second is reached with our accelerated gravity solver (—●—) with only 4 P_U . For our accelerated solver, we reached the maximum of nearly 9×10^5 data cells per second with 24 P_U . One P_U refers to 6 CPU cores and 1 GPU.

Figure 28: Speedup factors of our accelerated gravity solver run on 24 cores with different CPU/GPU ratios compared to the CPU-BH tree solver and the Gridsolver run on 24 and 144 CPU cores. The values are evaluated for the SP-TS simulation. The X-axis refers to the GPU count for different GPU/CPU ratios using a total of 24 CPU cores.

with the same number of CPU cores and a CPU/GPU ratio of 2/1 (12 GPU devices and 14 CPU Cores). Even when utilizing 144 (6 times more) CPU cores to the CPU-BH solver, we find our GPU accelerated solver to be over 9 times faster. We measure the speedup based on the “evolution” time stamp in the FLASH log file. Hence we ignore the initialization which is negligible in production simulations (Cordery et al., 2014). As can be seen in fig. 28, we find the best speedup with a one-to-one CPU/GPU ratio (24 CPU Cores and 24 GPU devices). Here, we see a total speedup of factor 10 compared to the CPU-BH solver and a speedup of 3 compared to the gravity Gridsolver (see fig. 29).

5. Conclusions

We described a GPU accelerated BH tree code as an additional Poisson solver for the FLASH4 software package. Our implementation works only for three-dimensional problems. Furthermore, the tree walk is limited to a maximum number of $\approx 32 \times 10^6$ data cells for each process.¹³ Three different simulation setups were used to test the accuracy and the performance of our novel GPU-BH code. Depending on the setup and GPU/CPU ratio we find a speedup for the gravity unit of at least a factor of 3 and up to 60 compared to the original CPU-BH implementation. For GPU/CPU ratios below 1/6 the runtime of our simulations is no more dominated by the gravity unit but

¹³Assuming a GPU with at least 6GB of memory

Figure 29: Evolution step speedup factors of our accelerated gravity solver with varying CPU/GPU ratios and a constant number of CPU cores compared to the CPU-BH code (—●—) and the Gridsolver (—■—). The X-axis refers to the GPU count for different GPU/CPU ratios using a total of 24 CPU cores. All values are evaluated for the SP-TS simulation run with 24 CPU cores.

by the hydro solver. Hence we find lower speedup factors between 1.6 and 10 for the total application runtime. We have shown that even with a small GPU/CPU ratio an advantageous performance gain can be achieved. The GPU-BH code was written for GPU-devices with a minimum compute capability of 2.0 (Fermi architecture) and we expect further runtime improvements by porting the code onto more modern devices with a higher compute capability. Here, we expect improvements not only from higher clock rates and higher register counts, but from additional features. E.g. with a compute capability ≥ 3.0 the warp shuffle functions `shfl()` could be used to reduce the shared memory usage in the force calculation kernels. Since our GPU code is written in CUDA, only NVIDIA GPUs are supported with the current version. A portable version of the GPU-BH tree code ported to OpenCL will be available in a later release. The GPU-BH tree solver for FLASH4 is available for free from the Hamburg Observatory at www.hs.uni-hamburg.de/gpubh.

6. Acknowledgement

The PARAMESH software used in this work was developed at the NASA Goddard Space Flight Center and Drexel University under NASA’s HPCC and ESTO/CT projects and under grant NNG04GP79G from the NASA/AISR project. The software used in this work was in part developed by the DOE NNSA-ASC OASCR Flash Center at the University of Chicago. This work was funded by the Deutsche Forschungsgemeinschaft under grant BA 3707/4-1. Much of the analysis and data visualization was performed using the yt toolkit by (Turk et al., 2011)

Appendix A. Periodic boundaries

The GPU-BH tree code supports isolated and periodic boundary conditions. In case of periodic boundary conditions, a node’s or data cell’s contribution to the gravitational potential of a data cell d_c changes to

$$\phi = GM f_{\text{EF}}(\vec{r}). \quad (\text{A.1})$$

Here, G is the gravitational constant, M is the mass of the node or data cell, $\vec{r}$ is the position vector between d_c and data cell or node and f_{EF} is the Ewald field. The original Ewald method is a method for computing the gravitational field for problems with periodic boundary conditions in three directions. Using the original method, the gravitational potential is split into two parts,

$$Gm/r = Gm \operatorname{erf}(\alpha r)/r + Gm \operatorname{erfc}(\alpha r)/r \quad (\text{A.2})$$

where α is an arbitrary constant. By applying Poisson summation formula on the erfc terms, the gravitational field at position $\vec{r}$ can be written in the form

$$\Phi(\vec{r}) = -G \sum_{a=1}^N m_a \times \left(\sum_{i_1, i_2, i_3} A_S(\vec{r}, \vec{r}_a, \vec{l}_{i_1, i_2, i_3}) + A_L(\vec{r}, \vec{r}_a, \vec{l}_{i_1, i_2, i_3}) \right). \quad (\text{A.3})$$

The first sum runs over whole computational domain, where mass m_a is at position $\vec{r}_a$. The second sum runs over all neighboring computational domains, which are at positions $\vec{l}_{i_1, i_2, i_3}$ and $A_S(\vec{r}, \vec{r}_a, \vec{l}_{i_1, i_2, i_3})$ and $A_L(\vec{r}, \vec{r}_a, \vec{l}_{i_1, i_2, i_3})$ are short, resp. long-range contributions. The Ewald field is calculated once on startup and stored in a large array representing a hierarchy of nested grids. The evaluation of the Ewald Field at certain points is only needed during the final force evaluation and carried out during the respective tree walk using quadratic interpolation.

Appendix B. Usage

The GPU-BH tree code for FLASH4 implements a set of runtime parameters to control the accuracy and runtime. All the parameters can be set individually using the common *flash.par* file. The opening angle parameter θ can be set with *grv_bh_gpuLimAngle* to a value between 0.1 and 1.0. The gathered data strategy can be selected with setting the runtime parameter *grv_bh_gpu_concat* to 1 where the default value of 0 refers to a serialized GPU device access pattern. In case periodic boundaries are used, the respective Ewald Field is controlled via the runtime parameters *gpuEwaldSeriesN* which controls the range of the indices i_1, i_2, i_3 in (Eq. A.3) and *gpuEwaldSeriesNref* which refers to the number of nested grids in the field array. The number of data points in each grid can be set with the parameters *gpuEwaldFieldNx*, *gpuEwaldFieldNy* and *gpuEwaldFieldNz*. Note, that the Ewald field resides in GPU memory during the hole simulation.

References

- Josh Barnes and Piet Hut. A hierarchical $O(N \log N)$ force-calculation algorithm. *Nature*, 324(6096):446–449, December 1986.
- Joshua E. Barnes. A modified tree code: Don’t laugh; It runs. *Journal of Computational Physics*, 1990.
- Joshua E. Barnes. *Computational Astrophysics*. Springer-Verlag, Berlin, February 1994.
- J Bédorf, E Gaburov, and S. Portegies Zwart. Bonsai: A GPU Tree-Code. *Advances in computational astrophysics: methods*, 453:325, July 2012.
- Jeroen Bédorf, Evghenii Gaburov, Michiko S Fujii, Keigo Nitadori, Tomoaki Ishiyama, and Simon Portegies Zwart. 24.77 Pflops on a Gravitational Tree-Code to Simulate the Milky Way Galaxy with 18600 GPUs. *arXiv*, 1412:659, December 2014.
- M Bartscher and K Pingali. An Efficient CU DA Implementation of the Tree-Based Barnes Hut n-Body Algorithm. *GPU computing Gems Emerald edition*, 2011.
- S Chandrasekhar. ELLIPSOIDAL FIGUEES OF EQUILIBIUM. 1969.
- M J Cordery, Brian Austin, H J Wassermann, C S Daley, N J Wright, S D Hammond, and D Doerfler. *Lecture Notes in Computer Science*, volume 8551 of *Lecture Notes in Computer Science*0302-97431611-3349. Springer International Publishing, Cham, October 2014.
- Sean M. Couch, Carlo Graziani, and Norbert Flocke. An Improved Multipole Approximation for Self-Gravity and Its Importance for Core-Collapse Supernova Simulations. *arXiv*, 778:181, 2013.
- A Deane, Gunther Brenner, David R Emerson, James McDonough, Damien Tremeur-Dervout, N Satofuka, A Ecer, and Jacques Periaux. *Parallel Computational Fluid Dynamics 2005*. Theory and Applications. Elsevier, September 2006.
- John Dubinski. A Parallel Tree Code. *arXiv*, astro-ph, March 1996.
- Christoph Federrath, Robi Banerjee, Paul C Clark, and Ralf S Klessen. Modeling Collapse and Accretion in Turbulent Gas Clouds: Implementation and Comparison of Sink Particles in AMR and SPH. *Audio, Transactions of the IRE Professional Group on*, pages –, January 2010.
- R T Fisher, L P Kadanoff, D. Q. Lamb, A Dubey, T Plewa, A Calder, F Cattaneo, P Constantin, I Foster, M E Papka, S I Abarzhi, S M Asida, P M Rich, C C Glendenin, K Antypas, D J Sheeler, L B Reid, B Gallagher, and S G Needham. Terascale turbulence computation using the FLASH3 application framework on the IBM Blue Gene/L system. *IBM J. Res. & Dev.*, 52(1.2):127–136, 2008.
- B. Fryxell, K. Olson, P. Ricker, F. X. Timmes, M. Zingale, D. Q. Lamb, P. MacNeice, R. Rosner, J. W. Truran, and H. Tufo. FLASH: An Adaptive Mesh Hydrodynamics Code for Modeling Astrophysical Thermonuclear Flashes. *The Astrophysical Journal Supplement Series*, 131(1):273–334, November 2000.
- E Gabriel, GE Fagg, G Bosilca, and T Angskun. Open MPI: Goals, concept, and design of a next generation MPI implementation. *Recent Advances in ...*, 2004.
- Lars Hernquist. Performance characteristics of tree codes. *Astrophysical Journal Supplement Series (ISSN 0067-0049)*, 64:715–734, August 1987.
- Pangfeng Liu Pangfeng Liu and SN Bhatt. Experiences with parallel N-body simulation. *IEEE Trans Parallel Distr Syst*, 11(12):1306–1323, November 2000.
- Peter MacNeice, Kevin M. Olson, Clark Mobarrry, Rosalinda de Fainchtein, and Charles Packer. PARAMESH: A parallel adaptive mesh refinement community toolkit. *Comput. Phys. Commun.*, 126(3):330–354, April 2000.
- J Nickolls, I Buck, M Garland, and K Skadron. Scalable parallel programming with CUDA. *Queue*, 6(2), 2008.
- NVIDIA. NVIDIA CUDA C Programming Guide. URL <http://docs.nvidia.com/cuda/cuda-c-programming-guide>.
- Tomasz Plewa, Timur Linde, and V Gregory Weirs. *Adaptive Mesh Refinement - Theory and Applications*. Proceedings of the Chicago Workshop on Adaptive Mesh Refinement Methods, Sept. 3-5, 2003. Springer, January 2005.
- P. M. Ricker. A Direct Multigrid Poisson Solver for Oct-Tree Adaptive Meshes. *The Astrophysical Journal Supplement Series*, 176(1):293–300, May 2008.
- John K Salmon. Parallel hierarchical N-body methods, 1991.
- John K Salmon and Michael S Warren. Skeletons from the Treecode Closet. *Journal of Computational Physics*, 111(1):136–155, December 1993.
- J. Kelly Truelove, Richard I. Klein, Christopher F. McKee, John H. II Holliman, Louis H. Howell, and Jeffrey A. Greenough. The Jeans Condition: A New Constraint on Spatial Resolution in Simulations of Isothermal Self-gravitational Hydrodynamics. *The Astrophysical Journal*, 489(2):L179–L183, November 1997.
- Matthew J. Turk, Britton D. Smith, Jeffrey S. Oishi, Stephen Skory, Samuel W. Skillman, Tom Abel, and Michael L. Norman. yt: A Multi-code Analysis Toolkit for Astrophysical Simulation Data. *The Astrophysical Journal Supplement*, 192(1):9, January 2011.