



Deposited via The University of Sheffield.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/id/eprint/123163/>

Version: Accepted Version

---

**Article:**

Gyenge, N., Griffiths, M.K. and Erdélyi, R. (2018) MHD code using multi graphical processing units: SMAUG+. *Advances in Space Research*, 61 (2). pp. 683-690. ISSN: 0273-1177

<https://doi.org/10.1016/j.asr.2017.10.027>

---

Article available under the terms of the CC-BY-NC-ND licence  
(<https://creativecommons.org/licenses/by-nc-nd/4.0/>)

**Reuse**

This article is distributed under the terms of the Creative Commons Attribution-NonCommercial-NoDerivs (CC BY-NC-ND) licence. This licence only allows you to download this work and share it with others as long as you credit the authors, but you can't change the article in any way or use it commercially. More information and the full terms of the licence here: <https://creativecommons.org/licenses/>

**Takedown**

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing [eprints@whiterose.ac.uk](mailto:eprints@whiterose.ac.uk) including the URL of the record and the reason for the withdrawal request.

# MHD Code Using Multi Graphical Processing Units: SMAUG+

N. Gyenge<sup>a,b,d,\*</sup>, M. K. Griffiths<sup>c</sup>, R. Erdélyi<sup>a,d</sup>

<sup>a</sup>*Solar Physics and Space Plasmas Research Centre (SP<sup>2</sup>RC), School of Mathematics and Statistics, University of Sheffield, Hounsfield Rd, Sheffield S3 7RH, UK*

<sup>b</sup>*Debrecen Heliophysical Observatory (DHO), Konkoly Observatory, Research Centre for Astronomy and Earth Sciences Hungarian Academy of Sciences, Debrecen, P.O.Box 30, H-4010, Hungary*

<sup>c</sup>*Corporate Information and Computing Services, The University of Sheffield, 10-12 Brunswick Street, Sheffield S10 2FN, UK.*

<sup>d</sup>*Department of Astronomy, Eötvös Lóránd University, Pf. 32, Budapest, H-1518 Hungary*

---

## Abstract

This paper introduces the Sheffield Magnetohydrodynamics Algorithm Using GPUs (SMAUG+), an advanced numerical code for solving magnetohydrodynamic (MHD) problems, using multi-GPU systems. Multi-GPU systems facilitate the development of accelerated codes and enable us to investigate larger model sizes and/or more detailed computational domain resolutions. This is a significant advancement over the parent single-GPU MHD code, SMAUG (Griffiths, M., Fedun, V., and Erdélyi, R. (2015). A fast MHD code for gravitationally stratified media using graphical processing units: SMAUG. *Journal of Astrophysics and Astronomy*, 36(1):197–223). Here, we demonstrate the validity of the SMAUG+ code, describe the parallelisation techniques and investigate performance benchmarks. The initial configuration of the Orszag-Tang vortex simulations are distributed among 4, 16, 64 and 100 GPUs. Furthermore, different simulation box resolutions are applied:  $1000 \times 1000$ ,  $2044 \times 2044$ ,  $4000 \times 4000$  and  $8000 \times 8000$ . We also tested the code with the Brio-Wu shock tube simulations with model size of 800 employing up to 10 GPUs. Based on the test results, we observed speed ups and slow downs, depending on the granularity and the communication overhead of certain parallel tasks. The main aim of the code development is to provide massively parallel code without the memory limitation of a single GPU. By using our code, the applied model size could be significantly increased. We demonstrate that we are able to successfully compute numerically valid and large 2D MHD problems.

*Keywords:* Numerical simulations, magnetohydrodynamics, graphical processing units, Sheffield Advanced Code

---

## 1. Introduction

Numerical simulations are one of the most important tools for studying astrophysical magnetohydrodynamic (or MHD) problems since the birth of computer science. MHD modelling the physical processes of a complex astrophysical observation frequently requires enormous computational efforts with high compute performance because of the often strong inhomogeneous and/or stratified magnetised plasma medium, just to name two challenges. Advances in modern processing unit technology allows us to solve more and more complex physical problems by using faster and higher number of central processing units (CPUs) or accelerators, such as graphical processing units (GPUs). Many implementations of MHD exist using a range of approaches and methods (see, e.g., Stone and Norman, 1992a,b; Stone et al., 1992; Tóth, 1996; Arber et al., 2001; Keppens et al., 2003; Hasan et al., 2005; Shelyag et al., 2008; Stone et al., 2008), however, numerical performance of GPUs is recently exploited for scientific purposes and it is becoming used by an increasing number

of diverse studies for modelling MHD problems (e.g. Wong et al., 2009; Pang et al., 2010; Griffiths et al., 2015).

Multi-GPU (mGPU) systems are able to provide further benefits, such as larger computational domains and substantial compute time savings, the latter resulting in saving of operational costs. Many studies demonstrate the performance effectivity for solving various astrophysical problems with mGPU architecture. For instance, Schive et al. (2011) used a parallel GPU-accelerated adaptive-mesh-refinement method for solving hydrodynamic problems. The mGPU systems allow us to achieve orders of magnitude performance speed-up compared to CPU cores (Wong et al., 2014b,a). However, the mGPU systems also enable to extend considerably the investigated model size or increase the resolution of the computational domain, therefore allowing to obtain more details.

The Sheffield Advanced Code (SAC) is a fully non-linear MHD numerical tool for simulating, in particular, linear and non-linear wave propagation in strongly magnetised plasma with structuring and stratification (Shelyag et al., 2008). SAC is based on the seminal Versatile Advection Code (VAC) (see Tóth, 1996). Griffiths et al. (2015) developed the parallel version of SAC using a single GPU. Griffiths et al. (2015) reported significant speed up, by

---

\*Corresponding author

Email address: n.g.gyenge@sheffield.ac.uk (N. Gyenge)




Figure 1: Flowchart outlining SMAUG+ with the implemented MPI parallelisation technique. The red boxes demonstrate the initial and distributed model configurations. The configuration is equally divided and spread around the different CPUs and GPUs, using MPI and CUDA. The white rectangles show the 'halo' cells, the grey rectangles demonstrate the buffers for storing the exchanged information. The red domains within each process show the actual mesh outline and the blue boundaries mark data, which is stored by different processes. The numerically intensive calculations are sent and performed by the GPUs. The GPUs distribute the subdomains further for calculating the actual numbers. These calculations are performed by the thousands of GPU cores (green rectangles). The figure is an example of a  $2 \times 2$  configuration.

demonstrating simulations of the Orszag-Tang vortex on a single NVIDIA M2070 GPU, being 18 times faster than the same problem running on a single core of an Intel Xeon X5650 CPU.

Our primary aim, here, is to introduce a fast MHD code for gravitationally highly stratified media by further developing SMAUG. Therefore we introduce SMAUG+, running it on a multi-GPU system for allowing increased computational domain, i.e. enabling larger physical model size computations. The developed software provides the opportunity to execute/perform simulations of MHD wave propagation mimicking the strongly magnetised solar atmosphere, in particular, representing the lower solar atmosphere from photosphere to low corona. Such approach is important, as there are a number of high-resolution ground- (e.g. SST - Swedish Solar Telescope, La Palma; DKIST - Daniel K. Inouye Solar Telescope, USA to be commissioned in 2019 or the EST - European Solar Telescope, to be realised by the second half of the next decade) and space-based (e.g. Hinode, SDO - Solar Dynamics Observatory, IRIS - Interface Region Imaging Spectrograph) facilities providing a wealth of earlier unforeseen observational details that need now to be understood.

## 2. Numerical Approach

The SMAUG+ is a numerical finite element solver, which is based on addressing the ideal fully non-linear 3-dimensional MHD equations. In particular, it is suitable to model linear and non-linear wave propagation in strongly magnetised plasma with structuring and stratification. The code is designed to solve multi-dimensional hyperbolic systems of partial differential equations. The governing equations of compressible MHD with stratification in their conservative form are:

$$\begin{aligned} \frac{\partial \rho}{\partial t} + \nabla \cdot (\mathbf{v}\rho) &= 0, \\ \frac{\partial(\rho\mathbf{v})}{\partial t} + \nabla \cdot (\mathbf{v}\rho\mathbf{v} - \mathbf{B}\mathbf{B}) + \nabla p_t &= \rho\mathbf{g}, \\ \frac{\partial e}{\partial t} + \nabla \cdot (\mathbf{v}e - \mathbf{B}\mathbf{B} \cdot \mathbf{v} + \mathbf{v}p_t) + \nabla p_t &= \rho\mathbf{g} \cdot \mathbf{v}, \\ \frac{\partial \mathbf{B}}{\partial t} + \nabla \cdot (\mathbf{v}\mathbf{B} - \mathbf{B}\mathbf{v}) &= 0, \end{aligned}$$

where  $\rho$  is the density,  $\mathbf{v}$  is the velocity,  $\mathbf{B}$  is the magnetic field,  $e$  is the energy density and  $\mathbf{g}$  is the gravity. The total pressure  $p_t$  is defined by:

$$p_t = p_k + \frac{B^2}{2},$$

where,  $p_k$  is the kinetic pressure:

$$p_k = (\gamma - 1) \left( e - \frac{\rho v^2}{2} - \frac{B^2}{2} \right).$$

Fourth-order central differencing method is applied for solving the spatial derivatives and fourth-order Runge-Kutta solver is performed for solving the time derivatives. By virtue of their symmetry, central differencing schemes are conservative, with the desired side effect that the solver conserves the divergence of the magnetic field. We employ hyper-resistivity and hyper-diffusion for increasing the numerical stability of the calculated MHD equation, based on the implementation by e.g., [Caunt and Korpi \(2001\)](#); [Stein and Nordlund \(1998\)](#); [Shelyag et al. \(2009\)](#). The hyper-viscosity coefficient calculations require the estimation of maximum wave speed in the computational domain.

By applying the central difference approximation to the hyperbolic differential governing equations, the solutions are unstable with a spurious oscillatory behaviour. Hence, numerical diffusion is applying for stabilising the code. The primary purpose of the diffusion terms is to compensate for the anti-diffusion from truncation errors arising in the computation of temporal and spatial derivatives. When the diffusion is correctly tuned, the resulting evolution is non-diffusive. In addition, the diffusion terms control the steepness of shocks by becoming large wherever the compression is large.

The MHD equations and the hyper-diffusion source terms are described in details in [Griffiths et al. \(2015\)](#). Numerical instabilities could be generated by the applied central differencing method, hence obtaining the solutions of the shocked systems could be difficult. The hyper-viscosity parameter is the ratio of the forward difference of a parameter to third order and first order. The temporal evolution of the hyper-viscosity term helps to identify numerical noise. If it is necessary, the hyper-viscosity term also helps to smooth the identified noise. For further details of the algorithm and the tests the reader is referred to the papers describing The Sheffield Advanced Code (SAC) ([Fedun et al., 2011a,b,c](#)) and SMAUG, the parallel version of the SAC employing a single GPU ([Griffiths et al., 2015](#)).

### 3. Parallelisation Techniques

MPI and NVIDIA CUDA allow us to develop a hybrid environment with multiple GPUs and multi-core CPUs. The MPI provides a message-passing parallel programming model: data is moved from the address space of one process to that of another process through cooperative operations on each process. We use the MPI methods to spread the initial model configuration.

Figure 1 shows the principals of an example system of architecture. The red boxes represent the computational

domain of the initial model configuration (e.g. for the Orszag-Tang vortex problem). The original grid, however, is now divided into four equal sub-regions, as indicated by the successive serial numbers. Each sub-region is assigned to a CPU. The CPUs are able to communicate with each other using communication fabrics, such as MPI messaging, OMNI-Path technology ([Birrittella et al., 2015](#)).

Figure 1 sketches an example of the CPU-CPU MPI communication. Exchanging information between the sub-domains with 'halo' layers is a common practice in parallel computation on CPUs ([Kjolstad and Snir, 2010](#)). The halo layers are demonstrated by the white rectangle within the computation domains (red rectangles). The data are obtained from (or sent to) the buffer of the top (or bottom) neighbour processors (indicated by grey rectangles). By sending and receiving only 'halo' cells and not the full grid, we reduce the size of the communications. Based on halo messaging technique employed in the SMAUG and SAC code, we demonstrate an example code of the implementation of MPI messaging.

---

```
void exchange_halo(vector v)
{
    //gather halo data from v into gpu_buffer1
    cudaMemcpy(host_buffer1, gpu_buffer1, ...);
    MPI_Isend(host_buffer1, ..., destination, ...);
    MPI_Irecv(host_buffer2, ..., source, ...);
    MPI_Waitall(...);
    cudaMemcpy(gpu_buffer2, host_buffer2, ...);
    //scatter halo data from gpu_buffer2 to
    //halo regions in v
}
```

---

The CUDA platform, however, provides us to access GPU accelerated solutions. Here, the actual numerical methods are performed/applied by the GPUs. With MPI and CUDA we can use multiple GPUs simultaneously for e.g., increasing model resolution or archiving speed-up.

### 4. The HPC facility

We ran the simulations using two different HPC facilities, namely, the University of Cambridge (Wilkes) and the University of Sheffield (ShARC) architectures. The Wilkes computer is based on Dell T620 machines. Each Dell T620 has two NVIDIA Tesla K20c (5 GiB) card and 12 cores Intel Ivy Bridge 2.6 GHz CPU and 64 GByte memory. Overall, the cluster consists of 128 Dell nodes with 256 GPUs. Each Tesla K20c contains 2496 GPU cores, hence the total number of GPU cores is 638976 and the total number of CPU cores is 1536.

The ShARC cluster is based on Dell PowerEdge C4130 units with 16 cores Intel Xeon E5-2630 v3 (2.4GHz) processors and 64 GByte memory. Accessing the GPU nodes provides 8 NVIDIA Tesla K80 (24 GiB) graphical units. The total number of GPU cores is 39936.

Table 1: Timings for 100 iterations for the Orszag-Tang test. The timing results are based on the simulations performed at the Cambridge Wilkes Cluster.

Grid Size	Number GPUs	Time [s]	
		Without HD	With HD
1000 × 1000	1 × 1	34.50	-
1000 × 1000	2 × 2	11.19	-
1000 × 1000	4 × 4	13.70	-
2044 × 2044	2 × 2	41.32	184.10
2044 × 2044	4 × 4	43.39	199.89
4000 × 4000	4 × 4	77.44	360.71
8000 × 8000	8 × 8	61.70	253.80
8000 × 8000	10 × 10	41.00	163.60

## 5. Verification and Validation

### 5.1. The Orszag-Tang vortex

The single GPU code (SMAUG) is already validated using e.g. the Brio-Wu shock tube test (Griffiths et al., 2015) for 1-dimensional MHD. The Orszag-Tang vortex is a common validation test employing two-dimensional non-linear MHD. This test is able to verify that the code employing it can handle MHD shock formations and shock interactions efficiently and cost-effectively. The horizontal and vertical non-dimensional domain sizes are set to  $0 \leq x \leq 1; 0 \leq y \leq 1$ , the adiabatic index  $\gamma = 5/3$ , the gas pressure and the density are constants. The boundary conditions are set to periodic at each boundary. The magnetic field components ( $B_x, B_y$ ) and the initial velocity components ( $V_x, V_y$ ) are defined by:

$$\begin{aligned} B_x &= -B_0 \sin(4\pi y), \\ B_y &= B_0 \sin(2\pi x), \\ V_x &= \sin(2\pi y), \\ V_y &= \sin(2\pi x). \end{aligned}$$

We ran a series of simulations with different simulation box resolution as seen in the Table 1. The timings with and without hyper-diffusion (HD) are distinguished, because we would like to understand the impact of the different code segments on the communication overhead between the GPUs.

Figure 2 is a set of snapshots of an example Orszag-Tang simulation. The panels shows the temporal variation of the density on a linear colormap. As the bottom panel of Figure 2 shows, various string waves pass through each other. This motion creates turbulent flow in different spatial scales. Figure 2 demonstrates that there is a convincing agreement between the results of SMAUG+ and their counterpart output of SMAUG (Griffiths et al., 2015) and SAC (Shelyag et al., 2008).

### 5.2. The Brio-Wu shock tube test

The Brio and Wu shock tube is an excellent test for solving a Riemann-type problem (Roe, 1981). The shock

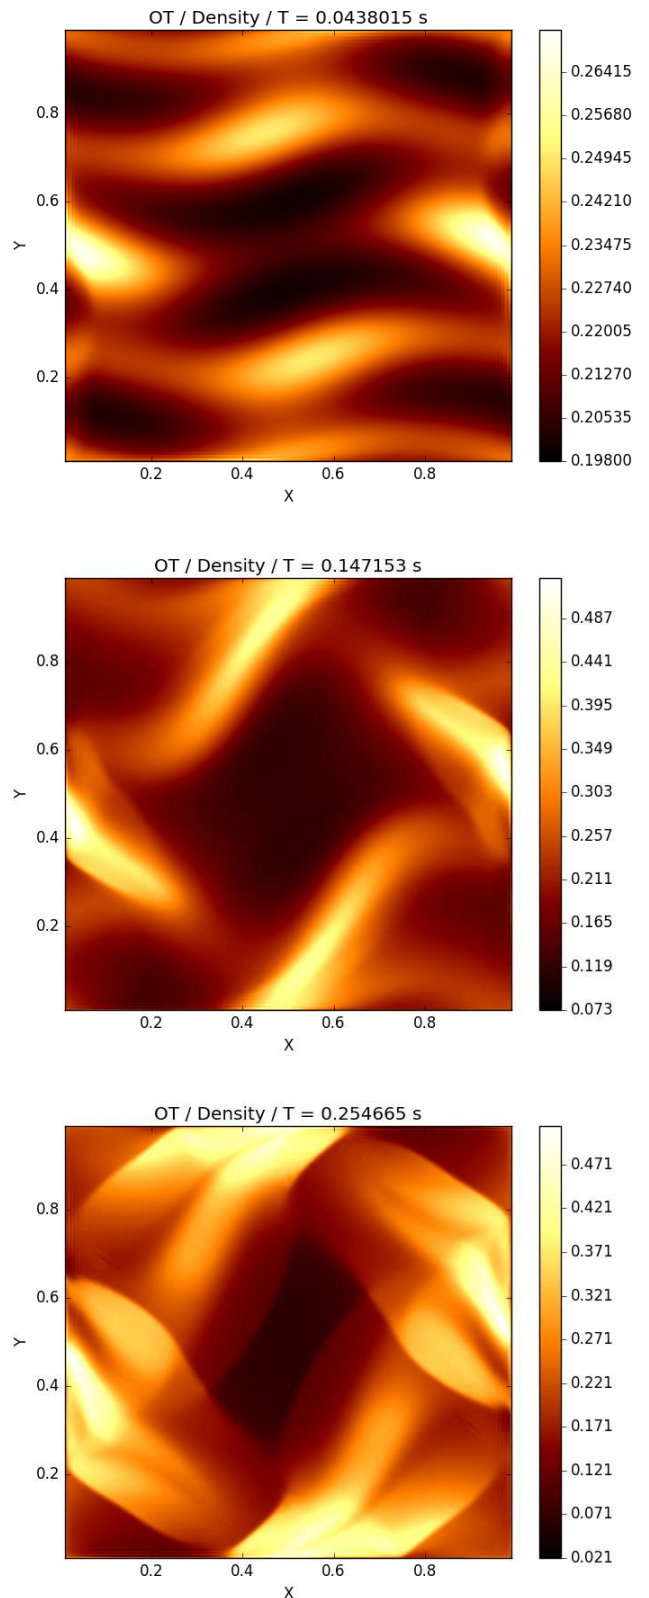


Figure 2: Orszag-Tang vortex results computed by SMAUG+. The initial configuration contains  $1000 \times 1000$  data and is distributed among 4 GPUs ( $2 \times 2$ ). The figure shows the temporal variation of the density for  $t_1 = 0.04$  s,  $t_2 = 0.14$  s and  $t_3 = 0.25$  s. The simulation ran on the ShARC cluster.

tube is a 1D ideal MHD test problem in which the initial conditions of the model feature a discontinuity in the centre of the configuration, i.e., the left and right states are initialised with different values (Brio and Wu, 1988). On either side of the discontinuity, the initial parameters are:  $p_l, \rho_l, B_{yl} = 1, p_r = 0.1, \rho_r = 0.125, B_{yr} = -1$  and  $B_x = 0.75$ . For the method employed by Brio and Wu (1988), the exact solution is approximated by a linearised version, averaged on either side of the discontinuity.

Figure 3 shows various features confirming the accurate run on the code. The slow compound wave, the contact discontinuity, the slow shock and the fast rarefaction can be seen. Running the problem on a numerical domain with 800 grid points, gives an excellent agreement with the original SAC (Shelyag et al., 2008) and SMAUG (Griffiths et al., 2015) results. We distributed the initial configuration among up to 10 GPUs. The tests demonstrate that code is sufficiently robust so that it can handle supersonic MHD turbulence.

## 6. Parallel Performance

The actual parallel performance of the applied models is determined by various factors, such as: (i) the granularity of the parallelizable tasks, (ii) the communication overhead between the nodes, (iii) finally, the load balancing. The load balancing refers to the distribution of the data among the nodes. If the data distribution is not balanced, some of the GPUs with less load must wait until the heavily loaded GPUs finish the job. We always use equally divided configurations, hence all the GPUs are equally loaded. The granularity of the parallel task represents the amount of work that will be carried out by a certain node. The communication overhead is the cost of sending and receiving information between the different nodes. In our case, the overhead is built-up by two components: the MPI node communication and the CPU-GPU information transfer, namely the functions `cudaMemcpy()`, `MPI_Isend()`, `MPI_Irecv()` and `MPI_Waitall()`. The actual running time  $T$  of a parallel job task is,

$$T = T_{comp} + T_{comm} + T_{idle},$$

where,  $T_{comp}$  is the useful computation time,  $T_{comm}$  is the communication overhead and  $T_{idle}$  is the waiting time until all the GPUs finishes a certain step. The  $T_{idle}$  could be minimised by optimal load balance and  $T_{comm}$  could be relatively small by choosing optimal granularity.

Let us define the quantity  $\tau$  for characterising the parallelisation overheads:

$$\tau = \ln \frac{T_{comp}}{T_{comm} + T_{idle}}.$$

Parallel slowdown ( $T_{comm} + T_{idle} > T_{comp}$ ) could be the result of a communication bottleneck. More GPUs must




Figure 3: Numerical solution of Brio and Wu shock tube problem numerical solution taken at time  $t_0 = 0$ ,  $t_1 = 0.015$  and  $t_2 = 0.03$ . The density (purple colour), the tangential magnetic field (yellow colour), the tangential velocity (blue colour) and the (red colour) normal velocity are distinguished among the panels. The initial configuration was equally distributed among 4 GPUs.




Figure 4: Brio-Wu and Orszag-Tang communication overhead test. The vertical axis represents the parameter  $\tau$ , the ratio of useful calculation and communication. The horizontal axis shows the number of applied GPUs. The size of initial configuration is constants (800 data points for Brio-Wu and  $2044 \times 2044$  for the Orszag-Tang test. The increasing parallelism alone cannot increase performance.

spend more time for communication. Sometimes, the communication between the GPUs takes actually more time than the useful calculation. This situation is demonstrated by the Figure 4. The blue values represent Brio-Wu simulations, using up to 10 GPUs. As the number of GPUs increases, the parameter  $\tau$  decreases. We kept the model size constant. The range of the parameter  $\tau$  shows extremely small numbers, e.g  $\tau = -4.5$  means only 1% of the running time was actual calculation. This value is extremely small, because it is likely that much of the communications overhead arises from routines, used for transferring data within the GPU memory. From the GPU memory the data must be transferred to the system memory. From here, the CPU will send the information to another CPU node, finally this node transfers the data to the GPU and so on. This continuous data transfer significantly jeopardise the parallel performance. This is the consequence of using not computationally dense GPUs. The red values in Figure 4 shows similar behaviour but the used model is a  $2044 \times 2044$  Orszag-Tang configuration. In this case, a GPU calculates a larger amount of data, hence the nodes are more computationally dense. Around 30% of the running time is the actual calculation. The total running time of the 9-GPU configurations in both cases is around 4 times slower than the 2-GPUs configuration.

As shown above, choosing the non-optimal configuration could cause massive wasting of computational power. To avoid parallel slowdown the following must be considered: (i) Only increasing parallelism will not provide the best performance (Chen et al., 1990). Increased parallelism with non-optimal data granularity could easily cause parallel slowdown, as seen above. (ii) The amount of exchanged MPI messages must be reduced as much as possible for the best performance (Thakur et al., 2004). It also means that a single GPU could give better performance than multiple GPUs, if the applied model size is the same.

(iii) Task must be enough to overlap the parallel communication overheads. The processes must have a higher task granularity if the number of applied GPUs increases. For avoiding the communications overhead, it is advisable to use always arithmetically dense GPUs. The adjustment of the model size and the decomposition of the initial configuration must be optimal. (iv) It is possible to improve communication performance by using higher-performance communication hardware, but it may turn out expensive. The collective communication, can also improve communication performance, since it optimises the communication based on the hardware, network and topology.

By applying the above principle parallel performance speed-up is possible. In Table 1, the  $1000 \times 1000$  Orszag-Tang test with 4 GPUs is around 3 times faster than the 1 GPU configuration. The  $8000 \times 8000$  test shows 1.5 times speed-up between 64 and 100 GPUs. The simulations with HD show similar properties than the simulations without HD. The timings with HD are significantly longer than without HD which shows that the HD functions are extremely computation intense segments of the SMAUG+.

## 7. Discussion

Future, high-performance computer architectures are enabling an increasing use of large numbers of GPU accelerators. It is crucial to develop codes enabling an the available fast communication between the GPUs. Our approach to exchange information between decomposed domains not only occupies one core in each CPU but also it may become a bottleneck to speed and scaling because of indirect data transfer between GPUs. This issue may be solved by direct data transfer between GPUs introduced by the GPU Direct technology.

However, the primary aim of our approach is to archive extended model size. A single GPU is only able to support limited memory but using our method an mGPU system provides as much memory as the total of GPUs. The only disadvantage is the communication overhead, however, an mGPU system may still be faster and significantly cheaper than a multi-CPU approach. For the same price, a GPU contains orders of magnitude more processing cores than a CPU. Our approach provides affordable desktop high-performance computing.

The challenge with advancing these codes is the identification and application of the current communication standards: different possibilities including NVLINK and various versions of the GPU Direct technology. These implementations are sensitive to the system configuration, making it challenging to successfully implement the code.

## 8. Conclusion

We have demonstrated that SMAUG+ is able to successfully compute numerically valid and large MHD problems by distributing the needed compute tasks across multiple GPUs. We conclude that the speed-up of the MHD

simulations performed depends on the distribution of the initial architectural configuration. In some cases, the running time could be slower if unwanted communication overhead arises between the higher number of GPU. However, and we propose this is a key point, a not over-distributed simulation could show significant speed-up. Further performance enhancements are also feasible through application architecture modification. By using our method an mGPU system essentially is able to provides as much memory as the total of applied graphical accelerators offer. The algorithm has been implemented in 3D testing of 3D models that will be completed over a forthcoming projects.

## Acknowledgments

MG and RE are grateful to STFC (UK), and RE acknowledges The Royal Society (UK) for the support received. This research was made use of SunPy, an open-source and free community-developed Python solar data analysis package (Mumford et al., 2013). The authors are grateful to the University of Cambridge HPC service for providing access to the Wilkes cluster and to Filippo SPIGA, Head of Research Software Engineering (RSE) University of Cambridge for providing assistance with running the the SMAUG+ code on the Wilkes cluster.

## References

### References

Arber, T., Longbottom, A., Gerrard, C., and Milne, A. (2001). A staggered grid, Lagrangian–Eulerian remap code for 3-D MHD simulations. *Journal of Computational Physics*, 171(1):151–181.

Birrittella, M. S., Debbage, M., Huggahalli, R., Kunz, J., Lovett, T., Rimmer, T., Underwood, K. D., and Zak, R. C. (2015). Intel® Omni-path Architecture: Enabling Scalable, High Performance Fabrics. In *High-Performance Interconnects (HOTI), 2015 IEEE 23rd Annual Symposium on*, pages 1–9. IEEE.

Brio, M. and Wu, C. C. (1988). An upwind differencing scheme for the equations of ideal magnetohydrodynamics. *Journal of computational physics*, 75(2):400–422.

Caunt, S. and Korpi, M. (2001). A 3D MHD model of astrophysical flows: Algorithms, tests and parallelisation. *Astronomy & Astrophysics*, 369(2):706–728.

Chen, D.-K., Su, H.-M., and Yew, P.-C. (1990). The impact of synchronization and granularity on parallel systems. *SIGARCH Comput. Archit. News*, 18(2SI):239–248.

Fedun, V., Shelyag, S., and Erdélyi, R. (2011a). Numerical modeling of footpoint-driven magneto-acoustic wave propagation in a localized solar flux tube. *The Astrophysical Journal*, 727:17.

Fedun, V., Shelyag, S., Verth, G., Mathioudakis, M., and Erdélyi, R. (2011b). MHD waves generated by high-frequency photospheric vortex motions. *Ann. Geophys.*, 29:1029–1035.

Fedun, V., Verth, G., Jess, D., and Erdélyi, R. (2011c). Frequency filtering of torsional alfvén waves by chromospheric magnetic field. *The Astrophysical Journal Letter*, 740(2):L46.

Griffiths, M., Fedun, V., and Erdélyi, R. (2015). A fast MHD code for gravitationally stratified media using graphical processing units: SMAUG. *Journal of Astrophysics and Astronomy*, 36(1):197–223.

Hasan, S., Van Ballegoijen, A., Kalkofen, W., and Steiner, O. (2005). Dynamics of the solar magnetic network: Two-dimensional MHD simulations. *The Astrophysical Journal*, 631(2):1270–1280.

Keppens, R., Nool, M., Tóth, G., and Goedbloed, J. (2003). Adaptive mesh refinement for conservative systems: multi-dimensional efficiency evaluation. *Computer Physics Communications*, 153(3):317–339.

Kjolstad, F. B. and Snir, M. (2010). Ghost cell pattern. In *Proceedings of the 2010 Workshop on Parallel Programming Patterns, ParaPLOP '10*, pages 4:1–4:9, New York, NY, USA. ACM.

Mumford, S., Pérez-Suárez, D., Christe, S., Mayer, F., and Hewett, R. J. (2013). SunPy: python for solar physicists. In *Proceedings of the 12th Python in Science Conference*, volume 74-77.

Pang, B., Pen, U., and Perrone, M. (2010). Magnetohydrodynamics on heterogeneous architectures: a performance comparison. *CoRR*, abs/1004.1680.

Roe, P. L. (1981). Approximate riemann solvers, parameter vectors, and difference schemes. *Journal of computational physics*, 43(2):357–372.

Schive, H.-Y., Zhang, U.-H., and Chiueh, T. (2011). Directionally Unsplit Hydrodynamic Schemes with Hybrid MPI/OpenMP/GPU Parallelization in AMR. *ArXiv e-prints*.

Shelyag, S., Fedun, V., and Erdélyi, R. (2008). Magnetohydrodynamic code for gravitationally-stratified media. *Astronomy & Astrophysics*, 486(2):655–662.

Shelyag, S., Zharkov, S., Fedun, V., Erdélyi, R., and Thompson, M. J. (2009). Acoustic wave propagation in the solar sub-photosphere with localised magnetic field concentration: effect of magnetic tension. *Astronomy & Astrophysics*, 501(2):735–743.

Stein, R. and Nordlund, Å. (1998). Simulations of solar granulation. i. general properties. *The Astrophysical Journal*, 499(2):914–933.

Stone, J. M., Gardiner, T. A., Teuben, P., Hawley, J. F., and Simon, J. B. (2008). Athena: A New Code for Astrophysical MHD. *The Astrophysical Journal Supplement Series*, 178:137–177.

Stone, J. M., Mihalas, D., and Norman, M. L. (1992). ZEUS-2D: A radiation magnetohydrodynamics code for astrophysical flows in two space dimensions. III - The radiation hydrodynamic algorithms and tests. *The Astrophysical Journal Supplement Series*, 80:819–845.

Stone, J. M. and Norman, M. L. (1992a). ZEUS-2D: A radiation magnetohydrodynamics code for astrophysical flows in two space dimensions. I - The hydrodynamic algorithms and tests. *The Astrophysical Journal Supplement Series*, 80:753–790.

Stone, J. M. and Norman, M. L. (1992b). ZEUS-2D: A Radiation Magnetohydrodynamics Code for Astrophysical Flows in Two Space Dimensions. II. The Magnetohydrodynamic Algorithms and Tests. *The Astrophysical Journal Supplement Series*, 80:791.

Thakur, R., Gropp, W. D., and Toonen, B. (2004). Minimizing synchronization overhead in the implementation of MPI one-sided communication. *PVM/MPI*, 3241:57–67.

Tóth, G. (1996). General code for modeling MHD flows on parallel computers: Versatile Advection Code. *Astrophysical Letters and Communications*, 34:245–250.

Wong, H.-C., Wong, U.-H., Feng, X., and Tang, Z. (2009). Efficient magnetohydrodynamic simulations on graphics processing units with CUDA. *ArXiv e-prints*.

Wong, U.-H., Aoki, T., and Wong, H.-C. (2014a). Efficient magnetohydrodynamic simulations on distributed multi-GPU systems using a novel GPU Direct–MPI hybrid approach. *Computer Physics Communications*, 185(7):1901–1913.

Wong, U.-H., Wong, H.-C., and Ma, Y. (2014b). Global magnetohydrodynamic simulations on multiple GPUs. *Computer Physics Communications*, 185(1):144–152.